LUDWIG-MAXIMILIANS-UNIVERSITÄT

DOCTORAL THESIS

Essays on the Economics of International Migration

Joop Age Harm ADEMA

A thesis submitted in fulfillment of the requirements for the degree of Doctor oeconomiae publicae

2025

Essays on the Economics of International Migration

Inaugural-Dissertation
zur Erlangung des Grades Doctor oeconomiae publicae (Dr. oec- publ.)
an der Ludwig-Maximilians-Universität München

2025

vorgelegt von
JOOP AGE HARM ADEMA

Referent: Prof. Panu Poutvaara
Koreferent: Prof. Fabian Waldinger
Promotionsabschlusssberatung: 29. Januar 2025

Acknowledgements

I am grateful for the never-ending support and enthousiasm of Panu, who transmitted both the joy to do research as well as the rigor to actually do it properly. I learned a lot from you in the past years and enjoyed our conversations about research and beyond. Without your wise advice, I would have never gone beyond Europe for a research visit in the US. Fabian and Joachim, thank you for serving as second and third supervisors.

I started the PhD journey in 2021 after defending my master thesis at TUM, moving from the HiWi-basement to a real office and from Dietersheim to München. However, the Covid-19 pandemic was in its lock-down stages. Many Zoom calls with Panu and Cevat paved the way for the first paper of this dissertation. However, doing research from home is not very stimulating and definitely not good for creativity. However, after the pandemic passed, meeting fun, bright and creative minds (Matteo, Johannes and Gabriele) during conferences was a boost for creativity, which is as contagious as Covid-19.

The second paper of this thesis followed naturally from the first, in a quest to understand *how* the internet impacts migration. On a conference night in a hotel room in Braga, I googled the rollout dates of Duolingo courses, and well, that became my single authored paper. Thank you Cevat and Panu, for developing this in a full-fledged single-authored paper and your suggestions on where to take it.

Return to the office created a more stimulating environment, making work much more enjoyable. All the inhabitants of the Villa, thank you very much for the good lunch times in the garden. With the Italian migration wave came the good coffee and the word of the week. Leander, Alessandro, Michael, Sarah, Felicitas, Lasha, Anna, Stella, and all other (temporary) inhabitants of the Villa, thank you for the pleasant times and interesting discussions.

Although not all researchers like writing policy papers, my experience has been different, not in the last place because of the people I got to work with. Marcel, Timo, and Phillip, thank you for working on interesting reports and being always up for a joke and a chat. Clara, Tanja and Mirely, it was a pleasure to work with you. Tanja, especially thank you for hiring me as a HiWi to ifo exactly six years ago. Beyond those I have worked with directly, I'd like to thank everyone at ifo who helped to shape an open environment with many interesting exchanges on research and policy which led me into doing a PhD in the first place. I'd also like to thank the support staff at ifo who helped me with business trips, printing and numerous photo shoots.

A defining, shocking and sad moment that happened during my PhD was the invasion of Ukraine, leading to the largest refugee movement in Europe since WW II. This shifted our course of work to studying Ukrainian refugees who remain in limbo for the foreseeable future. This led to the last paper of this dissertation, in which I had the joy to work with Yvonne. Thank you for your thoughtful advice and discussions about research. Lasha, Panu and Yvonne, our long discussions on survey questions resulted in insightful work and much useful material to better understand refugee integration in the upcoming years.

Behind the scenes of the chapters of this dissertation there is a lot of tedious data cleaning by RAs and interns. In particular, I would like to thank Tuncay, Lara and Padma for their great help and interesting discussions. I hope you learned something from me.

Although we all flocked out of Groningen, it is always great meeting up and hanging out with you guys, Oscar, Han, Jan and Sybren. Sil, thanks for always up for chilling in Groningen and going places, and for a lot of fun and philosophical discussions about life and the beautiful game. Although I left physics after coming to Munich, it was a pleasure studying and hanging out with you guys, Goku and David. I hope you can make it to my defense this time as well, and hopefully no pandemic breaks out this time.

I am extremely grateful to my parents, who had to endure the distance of having their only child living far from home. You always supported me in my education, even though I deviated a bit in unexpected directions. Coming home is always a pleasure and I hope I make you proud.

Maja, thanks you for being the flower in my garden. Thank you for all the time I get to spend with you, even if I ramble about new research ideas. I know I bored you many times with countless interruptions about regressions, replications, rent-seeking and refugees, but it helped me understand the matter better. Unfortunately, for you, I fear that I will keeping doing this in the years to come. Every end is also a beginning.

Contents

List of Figures

List of Tables

Introduction

Despite globalization and development aid to low-income countries, the rate of economic convergence between countries is low (Johnson and Papageorgiou, 2020). As income differences across countries remain vast, international labor mobility has the potential to considerably improve global welfare (Clemens, 2011; Benhabib and Jovanovic, 2012). Prior studies have confirmed that individuals' income increases considerably after moving to higher-income countries (Hendricks and Schoellman, 2018) and migrants even enjoy higher incomes after returning to their country of origin (Amanzadeh, Kermani and McQuade, 2024). Yet, differences in immigrants' earnings are not only driven by place-based differences in productivity, but also by variation in human capital (Hendricks and Schoellman, 2018), which may be driven by workers' language skills and cultural background (Adserà and Pytliková, 2016; Ek, 2024). However, mobility of people does not only affect involved individuals, but also their old and new country of residence, as well as the ties between the old and new home.

Origin countries. Although low-income migrant-sending countries lose qualified workers due to mobility, migration could provide incentives to invest in education, turning the *brain drain* into a *brain gain* (Beine, Docquier and Rapoport, 2001, 2008). Higher expected returns to education increase educational investments, which benefit the sending country as long as some of the marginally educated stay. Migration opportunities create human capital investments in the origin: For example, increased migration opportunities for nurses to the US increased nursing school enrollment by much more than total nurse emigration in the Philippines (Abarcar and Theoharides, 2024). Once settled in the host country, many migrants from developing countries send money to family members back home or spend money at home after return. Such remittances can spur investments in education and enable subsequent migration (Yang and Choi, 2007; Yang, 2008). Emigrants do not only remit monetary means, but may also transmit cultural and political norms (Docquier et al., 2016; Barsbai et al., 2017). Whether these effects are beneficial depend on host country norms. Moreover, after acquiring valuable skills or attaining sector-specific knowledge in employment, return migration may increase productivity in the sending country, which could mitigate brain drain (Dustmann, Fadlon and Weiss, 2011). Nevertheless, emigration of high-skilled workers may have pronounced negative effects on sending countries, such as the reduction in firm creation (Anelli et al., 2023).

Destination countries. The arrival of immigrants impacts host country labor markets through increased labor supply. If native and migrant labor are interchangeable, this exerts a downward pressure on natives' earnings opportunities. However, studies of natural experiments and quasi-experimental (Instrumental Variables) approaches tell a more complex story. The first widely studied natural experiment is the arrival of Cuban migrants in Miami in 1980 (the Mariel Boatlift), first studied by (Card, 1990) and re-analyzed by Borjas (2017), Peri and Yasenov (2019), and Clemens and Hunt (2019), finding no or negative effects only on the bottom of the skill distribution. Altonji and Card (1991) and Card (2001) use local migrants inflows to study labor market consequences, again only finding negative effects on the lowest skilled natives. Borjas (2003) introduced the skill-cell approach, which identifies the effect of skill-specific migration (based on education and experience) on relative labor market outcomes between skill cells. Borjas (2003) find a large negative effect on relative wages from immigration in the US.

However, immigrant and native labor may be complementary. Hence, increased supply of immigrant labor may increase the marginal product of native labor, which may explain positive wage effects for natives (Ottaviano and Peri, 2012; Manacorda, Manning and Wadsworth, 2012). This is supported by the observation that many immigrants perform different tasks than natives (Ottaviano, Peri and Wright, 2013). Low-skilled refugee migration increased native salaries through reduced employment in manual task intensive jobs (Foged and Peri, 2016) and Peri and Sparber (2009) and Peri (2012) suggest that through specialization of native workers low-skilled migration increased total factor productivity. Moreover, natives may respond by avoiding areas with cheap migrant labor (Dustmann, Schönberg and Stuhler, 2017). Not only do natives adjust in the labor market, changing factor prices for different types of labor can improve natives' schooling outcomes (Hunt, 2017). Altogether, the average effect of immigration on natives' wages on employment is zero or slightly positive, but has negative effects on low-skilled natives who are most likely to compete with immigrants (Edo, 2019). In settings where immigrants are culturally and linguistically close to natives, wage effects may be more negative, particularly in the informal sector (Caruso, Canon and Mueller, 2021).

Low-skilled migration could affect natives by worsening the fiscal balance by increasing transfers more than collected taxes (Storesletten, 2000). However, taking into account the general equilibrium effects discussed above, fiscal effects may be more positive and even low-skilled migration is likely to be fiscally positive (Colas and Sachs, 2024).

Host societies' resistance to migration is of all times (Tabellini, 2020). Despite the moderating role of contact (Allport, 1954), migration often induces negative attitudes and increases support for parties on xenophobic and far-right populist platforms (Steinmayr, 2021; Halla, Wagner and Zweimüller, 2017; Edo et al., 2019; Hangartner et al., 2019), particularly in rural places (Dustmann, Vasiljeva and Damm, 2019). This effect depends strongly on the skill composition of migrants, and be even reversed for high-skilled immigration (Mayda, Peri and Steingress, 2022). By and large, the opposition to migration is not driven by economic concerns (such as fear of job displacement or negative fiscal effects), but rather by cultural concerns (Alesina and Tabellini, 2024). The backlash to migration has far-reaching consequences for host societies. Migration reduces preferences for redistribution (Alesina, Miano and Stantcheva, 2023), which leads to increased (decreased) vote shares for parties favoring redistribution decreases with high-skilled (low-skilled) migration (Moriconi, Peri and Turati, 2019). The election of populist leaders itself reduces the presence of educated natives at the local level (Bellodi et al., 2024) and has large detrimental effects on economic performance on the national level (Funke, Schularick and Trebesch, 2023).

Bilateral ties. Migration flows may strengthen economic ties between sending and hosting countries. Migration between countries has been shown to increase trade flows by reducing trade frictions and shift geographic demand patterns (Gould, 1994; Parsons and Vézina, 2018; Bonadio, 2023). As many economic activities rely on tacit knowledge that is hard to transfer without interpersonal contact, migration enables the sharing of productive knowledge (Bahar and Rapoport, 2018). Hence, high-skilled migration may benefit both sending and hosting countries (after return) through knowledge diffusion (Fackler, Giesing and Laurentsyeva, 2020; Prato, 2022). Ties between countries may last for many generations, and increase cross-border capital flows on the long-run (Burchardi, Chaney and Hassan, 2019). More recent work has focused, among others, on transmission of leadership practices, political preferences (Spilimbergo, 2009; Barsbai et al., 2017) and cultural change (Rapoport, Sardoschau and Silve, 2021), which impact welfare in various ways.

The previous sections demonstrate that international mobility has complex consequences on migrants and their home and host countries. Consequently, reaping the large potential gains from migration requires understanding of its repercussions on all aspects of society. Using this newfound knowledge, policy makers can design sound policies, such as visa requirements and tax regulation. In this dissertation, I contribute to the study of two particular aspects of international migration: the drivers of migration decisions and the integration of migrants or refugees. All three chapters study the former, whereas the latter two study the latter.

Drivers of migration. The first quantitative studies of migration by Ravenstein (1876, 1885, 1889) pinned down several stylized facts about migration patterns in 19th century Europe. Among others, Ravenstein documented the role of distance in migration decisions. Economic theories of international migration have emphasized the earnings motives of prospective migrants. Sjaastad (1962) regarded the migration decision as an investment decision, incurring costs upon moving. Harris and Todaro (1970) suggested that migrants maximize expected utility from migration, explaining the fact that many internal migrants move to urban places with high unemployment. Nevertheless, these models do not predict that migration flows go in opposing directions. Adoption of the individual-destination specific random utility framework addressed this issue (McFadden, 1978), giving rise to "gravity"-like models of migration (Beine, Bertoli and Fernández-Huertas Moraga, 2016; Bertoli and Fernández-Huertas Moraga, 2013). Using this framework, Ortega and Peri (2013) show that migrants strongly respond to changes in income in the destination.

Borjas (1987) applies the Roy (1951) model to migrant selection in terms of education and earnings, showing that selection of migrants is driven by the relative dispersion of earnings in source- and host-country and the correlation between individual-specific earnings in both countries. This model rationalizes why migrants can outperform natives in the U.S. despite originating from lower income countries. Grogger and Hanson (2011) show that a model of *absolute* income-maximizing migrants explains the empirically observed pattern of positive educational selection and sorting patterns of migrants. Nevertheless, theories solely based on income maximization fail to explain why emigration rates from low-income countries is low (Zelinsky, 1971). However, emigration rates are an increasing function of income per capita at low levels, which can be explained by changing skill composition (Clemens, 2020), macro-economic conditions (Dao et al., 2018), and by liquidity constraints (Bazzi, 2017). As immigrants spend a considerable share of income in their country of origin, they have a disproportional preference to move to high productive, high-cost locations (Albert and Monras, 2022). Moreover, the level of welfare benefits may also influence migrants' location decisions (Borjas, 1999; Agersnap, Jensen and Kleven, 2020).

Although aforementioned models of migration focus on individual decisions, the deciding unit is often the household Mincer (1978). However, heterosexual couples often put higher weight on men's earnings (Munk, Nikolka and Poutvaara, 2022). Migration can diversify a household's risk from income shocks (e.g. bad harvests). Despite migrants often being more risk-tolerant themselves (Jaeger et al., 2010), more risk-averse households are more likely to send a migrant to diversify risk (Dustmann et al., 2023).

Migration decisions are often not permanent. Migrants often return home after some time, go forth- and back several times, or move onward to other destinations (Dustmann and Fabbri, 2003; Görlach, 2023). Migrants solve a complex lifecycle problem where they jointly decide on human capital decisions, labor supply and the timing of return migration (Dustmann, Fadlon and Weiss, 2011). Those planning to stay longer have stronger incentives to invest in host-specific human capital, which impacts integration outcomes (Adda, Dustmann and Görlach, 2022). Migration experience may increase the future decision to migrate, through reduced migration costs or acquired country-specific skills. For example, temporary international migrants are more likely to move abroad later in life (Parey and Waldinger, 2011).

Prior migrant networks play a crucial role in migration decisions as they transmit information and facilitate integration (Beine, Docquier and Özden, 2011). Although purely economic drivers have been successful in explaining both the extent as well as selection into migration, many other factors have been identified as important determinants of migration flows, including linguistic (Adserà and Pytliková, 2016) and political factors. Mayda (2010) show that destination country economic conditions and immigration policies, origin country demographic structure and migrant networks shape migration flows.

New technologies change the environment in which prospective migrants make decisions. Aviation reduced travel costs at larger distances, telecommunication enabled contact with prospective employers. Perhaps, the most impactful technology of the past decades is the internet, which may have profound effects on international migration. The internet spurred labor demand (Cooke and Shuttleworth, 2018; Hjort and Poulsen, 2019), and increased productivity relatively more for skilled workers (Akerman, Gaarder and Mogstad, 2015). Beyond the labor market effects, the internet radically changed the way prospective migrants can collect information about potential destinations and prepare for migration. However, not much is known how the internet affected migration decisions.

Migrant integration. Integrating newcomers into a host society is crucial for migrants' earnings, fiscal sustainability and for maintaining societal support for migration. In particular, language skills have been shown to crucial for economic integration of first generation immigrants and their children (Bleakley and Chin, 2004; Chiswick and Miller, 2015; Adserà and Pytliková, 2016; Heller and Mumma, 2023). Moreover, co-ethnic networks improve earnings upon arrival, but these benefits fade out over time (Battisti, Peri and Romiti, 2022). Timing also matters: Barsbai, Steinmayr and Winter (2024); Fasani, Frattini and Minale (2022) show that economic conditions upon arrival have a large impact on subsequent labor market outcomes. Modern communication technologies enable to sustain social contact with people in the origin. This may exert a negative influence on linguistic and social integration of migrants (Yarkin, 2024). However, new digital technologies can also improve the integration of immigrants through the acquisition of skills and the provision of information.

Across host countries, refugees have lower employment rates than other immigrant groups shortly after arrival and never fully catch-up to other immigrant groups and natives (Bratsberg, Raaum and Røed, 2017; Brell, Dustmann and Preston, 2020; Fasani, Frattini and Minale, 2022). As refugees have been forcibly displaced and often have to move unexpectedly, they lack the right formal qualifications and demanded skills (Anger, Bassetto and Sandner, 2022). Moreover, they may have experienced traumatic experiences that worsen (mental) health, further deteriorating their productivity. As refugee and family migrants in high-income countries are most likely to stay longer (Bratsberg, Raaum and Røed, 2017), addressing the limited integration of refugees has large potential benefits. The limited labor market success of refugees is partially driven by employment bans during the asylum process (Fasani, Frattini and Minale, 2021), which have long-run negative effects on economic assimilation.

Refugee-receiving countries have various policy levers beyond employment restrictions to improve refugee integration. Refugees are often centrally located within countries. Studies have shown that refugees' labor market integration strongly depends on local labor markets and natives' attitudes, suggesting large benefits from more sophisticated allocation mechanisms (Bansak et al., 2018; Aksoy, Poutvaara and Schikora, 2023). Not only economic integration is impacted by location. Bailey et al. (2022) show that there are strong regional differences in social integration among Syrian refugees in Germany, which is partially driven by the supply of integration courses. Many host countries provide courses to spur integration, which can be roughly divided into language learning programs and active labor market policies (ALMPs) (Bahar, Brough and Peri, 2024). Language learning programs have been shown to boost employment rates by several percentage points (Lochmann, Rapoport and Speciale, 2019; Foged, Hasager and Peri, 2024). Refugees also benefit from customized labor market plans with supervision and job search assistance (Sarvimäki, 2017; Dahlberg et al., 2024). On the other hand, reductions in benefits increase employment rates, but have negative side-effects on immigrants' children schooling and employment outcomes (Dustmann, Landersø and Andersen, 2024*a*,*b*). Granting refugees citizenship improves economic integration (Fasani, Frattini and Pirot, 2023). On the other hand, stricter regulation for obtaining permanent residency does not strongly incentivize refugees to integrate faster (Arendt, Dustmann and Ku, 2023).

Migrants' return intentions and integration are intricately related. A migrant planning to stay in the host country has stronger incentives to invest in host-country specific skills. On the other hand, better-integrated migrants are economically and socially more attached to the host country, which in turn may reduce return migration (plans). This integration-return nexus makes it hard to study the effect in either direction in isolation. Home-country conflict worsens return prospects, increasing integration efforts and outcomes. However, conflict shocks may not only affect return intentions, but also induce trauma, which may have an opposing effect on integration. Recent studies have examined this, finding support for a dominant role of the first channel. Bassetto and Freitas Monteiro (2024) find that sudden conflict in migrants' origin countries reduces return intentions among migrants in Germany and increases their job finding integration efforts, Zaiour (2023) show that drug-related violence in Mexico increases Mexicans' naturalization rates and social integration in the US, and Aksoy et al. (2024) find that school grades among Syrian children in Turkey improve when their home regions become more violent.

Uncertainty about staying prospects reduce refugees' investments in host-country skills, for example due to uncertainty faced in the asylum process (Brell, Dustmann and Preston, 2020). As most refugees in high-income countries often plan to stay in their host country, they have nevertheless strong incentives to invest in host-country skills and attachment to the labor market. In contrast to prior refugees in Europe, most Ukrainian refugees who fled after the large-scale Russian invasion in February 2022 ultimately plan to return. Conflict in Ukraine is continuing, and lasting longer than anticipated by Ukrainian refugees. Hence, these refugees face large uncertainty about their long-term residence, which could curb future investments in the host country. This motivates the study of how shocks in conflict in Ukrainian refugees' home region translate to return intentions, integration, and host country-specific skill investments.

This dissertation. The chapters of this thesis contribute to the understanding of migration decisions and migrant integration in various ways. The first two papers study how internet technologies have transformed the landscape for international migrants. Reduced costs of communicating across borders, improved information provision through social media and search engines and the ability to use applications to learn foreign language skills. The third paper studies the role of home region conflict in return migration and integration among Ukrainian refugees across Europe.

The emergence of the internet has drastically transformed the information landscape prospective migrants face. In this chapter, based on work with Cevat Giray Aksoy and Panu Poutvaara, I study the global effect of the arrival of mobile internet on individuals' aspirations and intentions to migrate. Prior examinations of the relation between internet and emigration relied on country-level variation uptake of the internet or on correlational evidence (Grubanov-Boskovic et al., 2021). We advance on this literature by combining fine-grained mobile network coverage information with large-scale surveys with information on the subnational location of individuals. Our empirical analysis combines survey data on 617,402 individuals with data on 3G mobile internet rollout from 2008 to 2018. Exploiting temporal variation in 3G rollout from 2,120 subnational districts in 112 countries, we show that an increase in mobile internet access increases the desire and plans to emigrate. Using lightning incidence as an instrument for network expansion, we provide additional evidence that the effects are causal. In line with our theory and recent work (Bertoli, Moraga and Guichard, 2020; Porcher, 2020), an important mechanism appears to be that access to the mobile internet lowers the cost of acquiring information on potential destinations. A case study using municipal-level emigration data from Spain shows that 3G rollout also increased actual emigration flows.

The internet not only enables the provision of information, it also provides new opportunities to learn skills that are otherwise costly to acquire. A prime example are language learning skills, which are crucial for migrants' success in the destination. The second chapter studies the rise of language learning applications, which has reduced the cost of learning host-country specific skills considerably. Although this plausibly increases language skills among those who would have anyways migrated, it may lead to non-trivial effects of selection into migration. I study the effects of the rollout of 84 courses on the influential language learning application *Duolingo* on learning and language skills, migration aspirations and flows and the selection and integration of migrants. As Duolingo courses enable learning a specific target language from a specific source language, this generates rich variation in the availability of low-cost language learning across migration corridors. The roll-out of these courses was plausibly supply-constrained and thus exogenous to trends in language learning and migration. First, I establish that course rollout increases online search interest in the target language and that courses enabling English learning improve TOEFL test scores in reading and listening, but not writing and speaking. Second, I find that the introduction of a course between the languages of a migration corridor strongly increases the stock of individuals aspiring to migrate across that corridor, but I find no conclusive evidence that actual migration flows increased considerably. Third, I find that the availability of low-cost language learning before arrival boosted language skills among immigrants in the EU, but not in the US. Availability after arrival increased language skills both in the EU and US. Moreover, employment rates increased considerably, but the educational attainment of migrants seem to have lowered. These results suggest that basic levels of language skills enable low-skilled migrants to find employment that does not require considerable language skills. These findings shed new light on the impact of the internet on migrant integration (Yarkin, 2024), as well as on the changing quality of more recent immigrant cohorts in the US (Borjas, 2015).

The third chapter changes directions and studies the return intentions and integration of refugees. Contrary to prior refugee groups in Europe, Ukrainian refugees by and large plan to return to Ukraine. This brings them to a highly uncertain situation, which may hamper their integration. Moreover, there is a lack of systematic evidence on how refugees' return intentions change over time and how these intentions align with actual return behavior. Additionally, the impact of conflict in refugees' home regions on their return plans and integration outcomes remains underexplored. To address these gaps I launched a longitudinal survey of Ukrainian refugees across Europe in June 2022 together with Cevat Giray Aksoy, Yvonne Giesing and Panu Poutvaara, and combined it with geocoded conflict data covering respondents' home municipalities. Exploiting conflict intensity after refugees have fled Ukraine, we can causally study the effect of home country conflict on return (intention) and investments in host country integration. Refugees are 5 percentage points more likely to return to Ukraine if their home district is liberated. High conflict intensity redirects return to safer regions, without reducing total returns. It also increases plans to settle outside Ukraine, as do more pessimistic war expectations. Those planning to settle outside Ukraine invest more in host-country-specific skills, but despite shifting return intentions, more intense home-region conflict does not increase the probability of working or language learning. These results differ from that of other recent studies that find that conflict in immigrants' home country increases integration efforts (Zaiour, 2023; Bassetto and Freitas Monteiro, 2024).

Chapter 1

Mobile Internet and the Desire to Emigrate

This chapter is based on joint work with Cevat Giray Aksoy and Panu Poutvaara

1.1 Introduction

The internet and mobile phones have changed how people live, work, connect, and exchange information. The number of internet users has increased about twelve-fold since 2000 - from 410 million in 2000 to nearly 4.9 billion in 2021, and is expected to continue double-digit growth (ITU, 2021). A vast majority of internet users have access to mobile internet: there were more than 4 billion mobile internet subscribers at the end of 2021 (GSMA, 2023).[1] At the same time, the average desire to emigrate across the world is on the rise. In the Gallup World Poll, the average desire to emigrate increased with 1.7 percentage points within the time period from 2008 to 2018, calculated for each country using the difference between the first and the last year the country was included and averaged across countries. Not only are desires to emigrate on the rise, actual migration is as well. The share of world population living outside their country of citizenship increased from 3.2 to 3.6 percent of world population between 2010 and 2020 (Batalova, 2022).

In this paper, we investigate how 3G mobile internet rollout causally affects desire and plans to emigrate.[2] To do so, we combine two unique data sets: Gallup World Polls (GWP) and Collins Bartholomew's Mobile Coverage Explorer, which allow us to study 617,402 individuals living in 2,120 subnational districts in 112 countries between 2008 and 2018. We exploit three different measures of migration aspirations and intentions from GWP: (1) whether an individual would like to move permanently to another country, if he or she had the opportunity; (2) whether an individual is planning to move permanently to another country in the next 12 months; and (3) whether an individual is likely to move away from his or her current city or area in the next 12 months, without a restriction to the migration being permanent or to another country. Previous research has already established that desire and plans to emigrate are strongly linked to subsequent actual migration flows (Tjaden, Auer and Laczko, 2019). We provide further evidence on correlation between desired (and planned) emigration and actual emigration flows in section 1.4.2.

To derive causal effects on desire and plans to emigrate, we exploit variation in subnational district 3G mobile internet coverage over time. We control for two-way (subnational district and year) fixed effects (TWFE), linear district-level time trends, as well as various individual, district and country-level characteristics. This implies that the estimates are identified by exploiting within-district variation in 3G coverage that has been stripped of any influence of constant and linearly changing district-level characteristics. As an additional test to dispel concerns about endogeneity, we instrument 3G mobile internet coverage with the intensity of local incidence of lightning strikes interacted with a linear time trend.

We find that 3G coverage has a sizable impact on the desire and plans to emigrate: a 10 percentage point increase in 3G mobile coverage leads to a 0.27 percentage point increase in the desire to emigrate permanently, and a 0.09 percentage point increase in plans to emigrate permanently over the ensuing 12 months. Given the average increase in 3G coverage of 36% from 2008 to 2018, this implies that the rollout of 3G internet access increased the desire to emigrate by 1 percentage point.

Although our main specification controls for subnational district and year fixed effects, as well as district-level linear time trends, it does not dispel all endogeneity concerns. We deal with

[1] Perhaps surprisingly, more households in developing countries own a mobile phone than have access to electricity or clean water, and nearly 70% of the poorest quintile of the population in developing countries own a mobile phone (The World Bank, 2016). However, not all mobile phone owners have a smart phone with access to the internet. Using the Gallup World Poll between 2016 and 2018, out of all people owning a mobile phone in Sub-Saharan Africa (Europe), 44% (88%) report having access to the internet.

[2] 3G mobile networks enable high-speed data transmission (i.e., advanced access to multimedia, social media, and search engines) that sets it miles away from 2G networks. Whereas regular 2G technologies (GPRS) enable data transfer of up to 5 kB per second, most current-day available 3G technologies (HSPDA) enable data transfer of up to 1 MB per second.

these concerns in four distinct and complementary ways. First, we find qualitatively similar results when we employ an instrumental variables (IV) strategy inspired by Manacorda and Tesei (2020) and Guriev, Melnikov and Zhuravskaya (2021). We use the incidence of lightning between 2005 and 2011 to predict slower 3G mobile network rollout from 2008 to 2018. Thix suggest that our findings are causal. Second, we use the alternative estimator by de Chaisemartin et al. (2024) to address concerns regarding the two-way fixed effects estimator. Using this alternative estimation method, we find qualitatively and quantitatively similar results. Moreover, this estimator enables us to assess pre-trends on a larger segment of our sample (as the estimator allows for continuous treatments) than a traditional event study focusing only around large increases in 3G coverage. We find that prior to receiving 3G internet coverage districts display similar trends in the desire to emigrate to districts that only receive 3G coverage later in time. Third, our results are robust to controlling for alternative time-varying measures of regional economic development. This dispels concerns that the reported effect is actually driven by a spurious relation between mobile internet and migration intentions, as swifter developing subnational districts may experience faster 3G rollout and migration intentions develop differently in these districts. Fourth, following the method proposed by Oster (2019), we show that our results are unlikely to be driven by the unobserved variation that is potentially related to omitted factors.

To further establish robustness, we show that our results do not change when including an extensive set of additional controls (such as satisfaction with life and local amenities). In addition, our estimates are robust across a variety of specification checks. Firstly, we show that higher levels of 3G coverage is not associated with regional income levels and only weakly associated with demographic characteristics using a balancing test, net of controls, fixed effects and time trends. Secondly, we use leads as treatments to assure that future increases in 3G coverage do not predict previous changes in the desire to emigrate. Thirdly, we show that our results are robust when using different survey weights, when we exclude potentially bad controls, and when omitting district-level time trends. In addition, our results remain statistically significant when clustering standard errors on the country level and after correcting for multiple hypothesis testing. Moreover, we show that our results are robust to several additional tests and that our results hold up in specific subsamples.[3]

In addition, we use 2G network expansion as a placebo treatment, which suggests that enabling texting and calling are not driving our results. Although we focus on 3G mobile internet, many individuals were previously exposed to fixed-line broadband internet. To investigate whether our results differ for individuals who have access to broadband internet, we interact our treatment with a variable capturing internet access at home and find no significant negative interaction effect. This suggests that the effect of 3G internet goes beyond the effect of regular internet access.

Using respondents' preferred destination we also find that preferred migration flows are redirected to less popular destinations. Both overall increase in international migration and changes in preferred destinations can have major long-term societal consequences. Migration affects productive capacities and income levels and can also boost innovation (Docquier and Rapoport, 2012; Alesina, Harnoss and Rapoport, 2016) and shape politics in both destinations (Dahlberg, Edmark and Lundqvist, 2012; Halla, Wagner and Zweimüller, 2017; Dustmann, Vasiljeva and Damm, 2019; Edo et al., 2019; Alesina, Miano and Stantcheva, 2023) and origins (Barsbai et al., 2017; Karadja and Prawitz, 2019).

We also explore the mechanisms behind our results. We begin by showing that the effect of 3G coverage on the desire to emigrate is only present for the individuals with no prior network

[3]Our results are robust on a sample (1) without districts and countries with possibly poor-quality 3G coverage data, (2) only including countries and districts that are surveyed in all years between 2008 and 2018, (3) excluding telephone interviews and countries with any telephone interviews altogether, (4) omitting years and global regions one at a time to show that our reported result is not driven by few influential observations, and (5) without top 10 refugee-origin countries as well as districts with very high or very low average desire to emigrate.

abroad, while we do not find effects for those who already have a network abroad. This suggests that the internet substitutes for personal networks as a means of acquiring information on opportunities abroad. We also show that 3G coverage does not improve the financial situation of respondents (e.g., household income) but has a negative effect on the perceived material well-being as well as trust in their national governments, which potentially shape emigration desires. Finally, using municipal level data from Spain, we show that 3G expansion not only alters desire to emigrate, but also increased actual emigration of home-country nationals.

The remainder of the paper is structured as follows. Section 2 reviews related literature and expands on our contributions to it. Section 3 introduces a theoretical framework we use to derive testable predictions. Sections 4 and 5 describe our data and empirical strategy. Section 6 presents the results. Section 7 explores the mechanisms. Section 8 presents evidence on how 3G coverage affected actual emigration from Spanish municipalities. Section 9 concludes.

1.2 Related Literature and Our Contributions

Our analysis connects to several strands of literature. First, there is the literature on the determinants of migration intentions. McKenzie and Rapoport (2010), Docquier, Peri and Ruyssen (2018) and Manchin and Orazbayev (2018) show that networks abroad are a major driving force of international migration intentions. In addition, Ruyssen and Salomone (2018) used the GWP to show that women are more likely to desire to emigrate from countries with high level of gender discrimination. Bertoli et al. (2022) find mixed evidence on the question whether weather shocks affect migration intentions in Western Africa. Böhme, Gröger and Stöhr (2020) provide suggestive evidence that the internet is important for international migration. They show that the intensity of destination-specific Google search queries semantically related to immigration can predict international migration flows. At the least, this suggests that prospective migrants use the internet to obtain information about prospective destinations. Pesando et al. (2021) provide descriptive evidence that there is a positive relation between internet access and migration using data on migration intentions from GWP and Arab Barometer, data on actual migration from the Italian Statistical Institute, as well as data from the Sant'Anna Cara reception center in Southern Italy. Across both levels of analysis, the authors find a positive *association* between internet access (measured as a percentage of the population using the internet) and both the willingness to emigrate as well as actual migration. Grubanov-Boskovic et al. (2021) show using GWP that not only migration aspirations are positively correlated to internet access, but also preparations to migrate are, after controlling for socio-demographic characteristics. We contribute to this literature by providing new *causal* evidence on the impact of internet access on migration aspirations and intentions, and by identifying the underlying mechanisms at work.

Second, we build on the recent literature on the impact of broadband internet and mobile communication technologies on economic and political outcomes. Hjort and Poulsen (2019) find that the arrival of broadband internet has a positive effect on employment in Africa. Zuo (2021) finds that employment probabilities of poor households and their earnings increased after obtaining broadband internet access in the United States. Falck, Gold and Heblich (2014) show that increased broadband internet availability reduced voter turnout in Germany. The authors relate this finding to a crowding-out of TV consumption and increased entertainment consumption. Campante, Durante and Sobbrio (2018) find that broadband internet access had a substantial negative effect on voter turnout in parliamentary elections in Italy until 2008, but this pattern has reversed since. Manacorda and Tesei (2020) use a granular data set for the entire African continent on 2G network coverage combined with geo-referenced data from multiple sources on the occurrence of protests. They find that while mobile phones are instrumental to mass mobilization, this only happens during economic downturns. In the most closely related study, Guriev, Melnikov and Zhuravskaya (2021) analyze the political implications of 3G internet rollout using GWP. They find that 3G expansion increases awareness of government corruption and reduces trust in political institutions. The authors further show that the effect is present only when the internet is not censored, and it is stronger when the traditional media are censored. We complement these studies by showing how 3G internet access also affects non-political outcomes | aspirations and intentions to emigrate | and by showing how internet access affects high stake individual actions | actual emigration from Spain.

Third, there is work on the income-related correlates of migration. Borjas (1987), Grogger and Hanson (2011), and McKenzie, Gibson and Stillman (2013) show that earning potential in the destination country shapes migration behavior. However, Dustmann and Okatenko (2014) show that the relationship between the intention to move (both domestically and internationally) and proxied wealth is non-monotonic. That is, the likelihood to move increases with personal income for those individuals living in the poorest global regions (Sub-Saharan Africa and Asia), while this relationship is absent for those living in relatively richer regions (Latin America). The inverse U-shaped relation between income and migration is also documented in Clemens (2014).

1.3 Theoretical Framework

There are two countries, denoted by 0 and 1. We analyze the decisions of residents of country 0 on whether to invest in acquiring information on opportunities abroad and whether to migrate to country 1 if mobile. We denote by vector $\boldsymbol{x_j}$ individual j's characteristics (such as age, gender, experience and family situation) that can influence earnings, the cost of acquiring information on opportunities abroad and migration costs in the case of being mobile. Vector $\boldsymbol{x_j}$ has a constant term that is used to capture potential earnings as well as information acquisition and migration costs of a reference person and n individual-specific components, given by $\boldsymbol{x_j} = \left(1, x_{j,1}, ..., x_{j,n}\right)$. In addition to individual characteristics, $\boldsymbol{x_j}$ also includes the 3G coverage in the region in which j lives inside country 0, denoted by $x_{j,3G}$. We denote the vector giving after-tax returns to individual characteristics in country k, $k \in \{0,1\}$, by $\boldsymbol{\beta_k}$, giving as potential disposable earnings in country k $\boldsymbol{\beta_k} \cdot \boldsymbol{x_j}$. As in Grogger and Hanson (2011), we divide education into primary, secondary and tertiary, and allow both returns to education and migration costs vary according to the level of education.

Potential mobility also has a stochastic component and acquiring information about opportunities abroad can be costly. This is inspired by Bertoli, Moraga and Guichard (2020), who analyzed costly information acquisition, in a setting with several potential destinations. We present a simpler model with a binary choice for information acquisition as GWP has no questions on the number of destinations from which respondents have gathered information. The information costs could be related to such issues as whether one could obtain a visa as well as job and housing opportunities abroad, with cost vector $\boldsymbol{\alpha}$ that specifies how information costs depend on individual characteristics. The total cost of information acquisition is $\boldsymbol{\alpha} \cdot \boldsymbol{x_j}$. Our main variable of interest is regional internet coverage, the effect of which is denoted by α_{3G}. As mobile internet access makes finding information easier, $\alpha_{3G} < 0$. If being internationally mobile and deciding to migrate, individual j faces migration cost c_j, which also includes the expected post-migration cost of communicating with family and friends left behind. The migration cost depends on individual characteristics $\boldsymbol{x_j}$ with a cost vector γ and an unobservable individual-specific component ϵ_j, capturing individual-specific taste for living abroad that is unobservable to researchers:

$$c_j = \boldsymbol{\gamma} \cdot \boldsymbol{x_j} + \epsilon_j. \tag{1.1}$$

Cost vector γ includes a component related to 3G coverage denoted by γ_{3G}, with $\gamma_{3G} < 0$ as a 3G network facilitates communication. Individual-specific component ϵ_j follows a continuous distribution with density function $\phi(.)$ and differentiable cumulative distribution function $\Phi(.)$, and obtains negative values for those with a preference for living abroad. For simplicity, we assume that those who invest in information acquisition learn with certainty whether they are mobile or not, and that the probability of being mobile is individual-specific, denoted by θ_j. Individual probability to be able to migrate θ_j depends on external constraints, such as immigration rules in the destination, and on psychological and social components, such as the effect of family members. It is individually optimal to invest in information acquisition if

$$\theta_j \left(\boldsymbol{\beta_1} \cdot \boldsymbol{x_j} - \boldsymbol{\beta_0} \cdot \boldsymbol{x_j} - \boldsymbol{\gamma} \cdot \boldsymbol{x_j} - \epsilon_j\right) > \boldsymbol{\alpha} \cdot \boldsymbol{x_j}. \tag{1.2}$$

In equation (1.2), the term in parentheses on the left-hand side gives the utility gain from migration, multiplied by the probability of being able to migrate. This equals the expected benefit from acquiring information on one's mobility and migrating if being able to do so. The right-hand side gives the cost of information acquisition. It is optimal to acquire information if the expected benefit from migration multiplied by the probability of being able to migrate exceeds the cost of finding out whether one could migrate. Those with too small or even negative gains from potential migration remain rationally uninformed on their mobility status, in line with Bertoli, Moraga and Guichard (2020). Equation (2) allows deriving the maximum individual-specific component $\hat{\epsilon}_j$ with which individual j would find it optimal to acquire information:

$$\hat{\epsilon}_j = \left(\boldsymbol{\beta_1} - \boldsymbol{\beta_0} - \boldsymbol{\gamma} - \boldsymbol{\alpha}/\theta_j\right) \cdot \boldsymbol{x_j}. \tag{1.3}$$

Denoting the probability of individual j investing in information acquisition by p_j, we have

$$p_j = \Phi\left(\left(\boldsymbol{\beta_1} - \boldsymbol{\beta_0} - \boldsymbol{\gamma} - \boldsymbol{\alpha}/\theta_j\right) \cdot \boldsymbol{x_j}\right). \tag{1.4}$$

In the individual components of vectors $\boldsymbol{\beta_0}$ and $\boldsymbol{\beta_1}$, we use superscripts for country indices, implying that $\beta^0_{3G,0}$ is the effect of 3G coverage in the region of origin on wage level in that location, and $\beta^1_{3G,0}$ is the effect of 3G coverage in the region of origin on wage level in the other country, if any. The effect of regional 3G coverage on the probability of individual j investing in information acquisition is given by:

Proposition 1. $\frac{\partial p_j}{\partial x_{3G}} = \left(\beta^1_{3G,0} - \beta^0_{3G,0} - \gamma_{3G} - \frac{\alpha_{3G}}{\theta_j}\right) \phi\left(\left(\boldsymbol{\beta_1} - \boldsymbol{\beta_0} - \boldsymbol{\gamma} - \boldsymbol{\alpha}/\theta_j\right) \cdot \boldsymbol{x_j}\right).$

Proof. *Follows by differentiating equation (1.4).*

The effect of 3G coverage on the probability of acquiring information is the product of two terms. The first term, $\left(\beta^1_{3G,0} - \beta^0_{3G,0} - \gamma_{3G} - \frac{\alpha_{3G}}{\theta_j}\right)$, is positive if the effect of 3G coverage on wages is sufficiently low. However, if 3G coverage would sufficiently boost wages in the region of origin, then an increase in 3G coverage could reduce migration. The second term, $\phi\left(\left(\boldsymbol{\beta_1} - \boldsymbol{\beta_0} - \boldsymbol{\gamma} - \boldsymbol{\alpha}/\theta_j\right) \cdot \boldsymbol{x_j}\right)$, is a scaling factor depending on the density of the individual-specific component at the cutoff point. As long as density is not zero, the sign of the effect of 3G coverage on the probability of acquiring information is determined by the first term. We assume that $\beta^1_{3G,0} = 0$, implying that 3G coverage in the region of origin has no effect on wages in the destination region. As $\alpha_{3G} < 0$ and $\gamma_{3G} < 0$, our main testable hypothesis is:

Hypothesis 1. *An increase in 3G coverage increases subsequent desire to emigrate, at least if it does not boost local wages substantially.*

Our model predicts that only a fraction $\bar{\theta}$ of those acquiring information can migrate, in which $\bar{\theta}$ is the average value of θ_j over all individuals who acquire information on their mobility status. Therefore, migration plans increase in the desire to emigrate but at a rate lower than one, giving a second testable hypothesis:

Hypothesis 2. *An increase in 3G coverage increases subsequent plans to emigrate, at least if it does not boost local wages substantially. The increase in plans to emigrate is smaller than the increase in the desire to emigrate.*

Both testable hypotheses are derived with the caveat that there is no substantial direct effect of 3G coverage on local wages. In the empirical analysis, we estimate the net effect of 3G coverage and, if positive, it already implies that the effect on boosting local wages is probably not very strong. A negative effect of 3G coverage on the desire to emigrate, instead, would suggest, as a potential explanation within the model, that the 3G coverage may have boosted local wages. In section 1.7 we show that 3G coverage does not increase per capita household income.

1.4 Data and Descriptive Statistics

The main data used in this work are survey data from Gallup World Polls (The Gallup Organization, 2022), mobile network data from Collins Bartholomew's Mobile Coverage Explorer (Collins Bartholomew, 2018) and geocoded lightning incidence data from the World Wide Lightning Location Network. We complement these data with additional subnational, country-level country-pair level data from a variety of sources.

1.4.1 Data

Gallup World Polls

Our primary data on emigration aspirations and intentions originate from the 2008-2018 GWP. These nationally representative surveys are fielded every year and interview approximately 1,000 individuals in each country on a wide range of topics.[4] Our resulting main sample includes 617,402 respondents from 112 countries living in 2,120 districts.[5]The questions of interest related to international migration in the GWP ask respondents about their desire and plans to emigrate.[6] The main outcomes of interest, their time span, the wording of the underlying question and possible responses are:

1. Desire to Emigrate (2008 – 2018): *Ideally, if you had the opportunity, would you like to move permanently to another country, or would you prefer to continue living in this country?* **Yes/No/Don't know/Refused to answer**

2. Plans to Emigrate (2008 – 2015): *Are you planning to move permanently to another country in the next 12 months?* (asked only of those who are desiring to move to another country) **Yes/No/Don't know/Refused to answer**

[4]If countries have sufficient telephone network coverage, households are drawn from a phone number database or on the basis of dialling random digits. If not, face-to-face interviews are conducted, with a 'random route' methodology of selecting households. Importantly, only after finding a household and identifying all of its members aged 15 or above, a household member is selected at random and up to three attempts to interview the selected member are made. If unsuccessful, a new household is approached to prevent a selection bias within a household's hierarchy. The coverage of countries, number of respondents, language of survey and method of conducting can be found here: `https://www.gallup.com/file/services/177797/World_Poll_Dataset_Details_052920.pdf`

[5]To construct subnational districts that are constant within a country over time, we harmonize the regions by country used in GWP across years to the least granular level available throughout 2008-2018. As for some countries the indicated region does not correspond to a administrative division, these regions are manually mapped to an administrative divisions. In most cases, the ultimate subnational districts correspond to the first or second administrative division of the country.

[6]The GWP contains multiple questions regarding migration intentions that do not fully overlap and, hence, we combine them when possible to not lose observations. This is especially important for question (2). The relevant constructed variables and exact underlying questions are all documented in Online Appendix Table 1.E.1.

Those replying “Yes" to Question (1) are potential migrants. Empirically, most of them do no realize their aspirations, either due to high migration costs or hurdles preventing actual migration (Docquier, Özden and Peri, 2014). Question (2) reveals more concrete intentions to migrate and has been shown to be strongly related to actual migration flows (Tjaden, Auer and Laczko, 2019). The emphasis on a relatively short time window of 12 months makes it plausible that mostly individuals with serious and developed migration plans who have the means to migrate answer affirmatively (Dustmann and Okatenko, 2014; Migali and Scipioni, 2019). This pattern is also revealed in Appendix Table 1.E.2: the share of respondents who actually plan to move abroad in the next 12 months (less than 3%) is substantially lower than the share of those who reported having desire to emigrate (22%).[7] To create binary indicators for both outcomes, a positive answer is recoded to 1, a negative answer is recoded to 0, and set to missing for the two residual options.

There is significant heterogeneity in emigration aspirations within and across countries. Panel a) of Figure 1.1 shows the averaged levels of the desire to emigrate in the 2008 – 2011 period, and Panel b) for the 2015 – 2018 period. Panel c) of Figure 1.1 displays the changes in average reported desire to emigrate between early (before the median year of all observations in the subnational district) and late years (during or after the median year). Notable patterns can be summarized as follows: (i) less than 20% desires to emigrate in most developed countries; (ii) less than 10% in many East and South Asian countries; (iii) there is substantial variation in the share of people desiring to emigrate within global regions over time | in Africa (an increase from 29% in early (2008 – 2011) to 32% in late (2015 – 2018) years), in Asia excluding the former USSR, Japan and South Korea (a decrease from 17% in early to 13% in late years), in Europe (an increase from 19% in early to 22% in late years), in the former USSR (an increase from 19% in early to 21% in late years), in Middle and South America (an increase from 27% in 2010 to 32% in late years) and in high-income non-European countries (an increase from 14% in early to 16% in late years); and (iv) there is substantial regional variation within countries. Between the first and last year a country is included in the Gallup World Poll, the desire to emigrate increased by 1.7 percentage points, from 22.6% to 24.3%.[8]

If a respondent desires to emigrate, they are also asked which country they *desire* to move to, which we use to validate the relevance of migration intentions for actual migration behavior in section 1.4.2 and to consider how 3G changes preferred destination countries in section 1.7.[9] The GWP also provides detailed information on respondents' demographic characteristics (age, gender, educational attainment, marital status and urban/rural residence), labor market outcomes, household income, satisfaction with local amenities and social networks abroad. This allows us to directly control for many relevant factors at the individual level. We proxy the level of development on the district-year level development level by the average of per capita household income of all people in a district (excluding the respondent). Furthermore, to control for the

[7] Additionally, GWP contains a question on the self-assessed likelihood to move away from the are someone lives within the next 12 months. We discuss the effects on this outcome and intersections with the emigration questions in Appendix 1.A.1. Furthermore, the GWP includes a question on *preparations to emigrate within 12 months*, which is asked to respondents indicating plans to emigrate within 12 months between 2009 and 2015 for a subset of subnational districts and answered positively by only 1.7% of respondents worldwide. We find no statistically significant effect on this outcome, which is unsurprising given the low statistical power of the test. By simulating a hypothetical treatment effect of 20% we find that a test of 5% significance only rejects the null hypothesis in 9% of cases, A detailed description of the power calculation is provided in Appendix 1.A.3.

[8] The change in average desire to emigrate of 1.7 percentage points is calculated as the difference in the desire to emigrate between the respondents interviewed in the first and last year a country is included in the Gallup World Polls between 2008 and 2018. We weight all observations by nationally representative weights provided by Gallup, which sum to one for all observations in a country-year.

[9] The proportion of individuals answering positively to Question (1) that do not mention a destination country is less than 7%. Similarly, less than 4% of those answering positively to Question (2) do not mention a destination country. Although respondents can choose not to mention a specific destination, the vast majority does.

time-varying age structure of the country, we calculate the share of respondents aged under 30 using the age distribution of GWP respondents.

Collins Bartholomew's Mobile Coverage Explorer

The information of 2G and 3G mobile network coverage around the world is obtained from Collins Bartholomew's Mobile Coverage Explorer.[10] The data provide information on signal coverage at 1x1 kilometer grid level, as submitted by network operators to the GSM Association. That is, we know whether or not a given 1x1 kilometer grid cell has a 2G or 3G signal (or both). However, we do not observe any information about the strength of the signal. The network coverage data is available on the yearly level, but the timing of data collection differs. Between 2011 and 2017, data is provided for the month December, whereas in 2007, 2008 and 2009, it is provided in the first quarter of the year.[11] We use the reported coverage in year $t-1$ to represent the network coverage in year t.[12]

To calculate the share of population that is covered by the 2G and 3G, we use 1x1 kilometer population data from the Gridded Population of the World (GPW) for 2005, which is distributed by the Center for International Earth Science Information Network.[13] We first calculate each grid point's population coverage and then aggregate this information over the subnational districts as provided in the GWP. The constructed population-weighted coverage of 3G networks is our main treatment variable.[14]

Figure 1.2 illustrates the increase in 3G internet coverage at the subnational district level over time.[15] In particular, Panel a) of Figure 1.2 shows the granular 3G internet coverage in 2008 and 2018, and Panel b) shows the relevant variation in population-averaged 3G coverage between 2008 and 2018 on the subnational district level. Perhaps not surprisingly, the levels of 3G internet coverage are highest in developed and densely populated countries, mostly achieving coverage levels of more than 75% of the population in 2018. Conversely, many Latin American and Sub-Saharan African countries have coverage levels of below 25%. Nevertheless, several non-OECD countries have showed expansions in excess of 25% over the 11-year period that we study. This offers relevant variation in 3G internet coverage on a global scale in the period studied. In the subnational districts present in our sample, 3G coverage increased from 16% to 52% on average, weighted by Gallup's nationally representative weights.

WWLLN Lightning Incidents Data

[10] For more information, please see: https://www.collinsbartholomew.com/mobile-coverage-maps/mobile-coverage-explorer/

[11] Due to the change in data provision, data between the first quarter of 2009 and December 2011 are missing. We overcome this challenge by linearly interpolating the missing information using the data from 2009 and 2011.

[12] As around 70% of the GWP interviews are conducted in July or earlier in the year, using the network data from previous December (for the interviews in 2012 up to 2018) is more informative of the actual network coverage during the interview.

[13] Since 2012, data on 4G network coverage has also been recorded in a subset of countries. As it is technically possible for an area to be covered by 4G but not by 3G, we might underestimate the share of population covered by mobile internet. We investigate this possibility and find that some urban areas in Czechia and India have 4G infrastructure without having 3G coverage. Across the whole sample in 2018, only less than 1% of the sample population is covered by 4G and not by 3G, which is not likely to bias our results.

[14] The data are not available for several populous countries such as Algeria, Angola, Argentina, Bangladesh, China, Ethiopia, Iran, Iraq, Kazakhstan, Myanmar, Morocco, Pakistan, Peru and Yemen.

[15] The data availability is somewhat limited for some countries. Data for some countries with large migration aspirations, intentions and flows in the Middle East and North Africa (MENA) region are absent from the final data set.

FIGURE 1.1: Desire to Emigrate around the World over Time

(A) Country-averaged desire to emigrate (2008-2011)

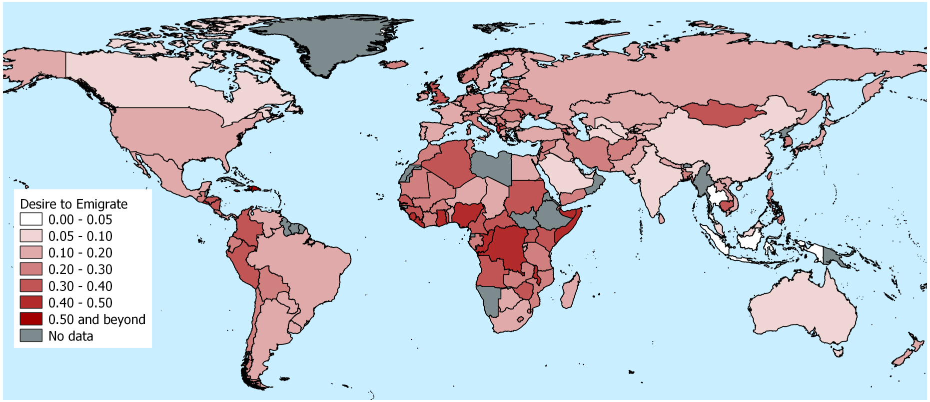

(B) Country-averaged desire to emigrate (2015-2018)

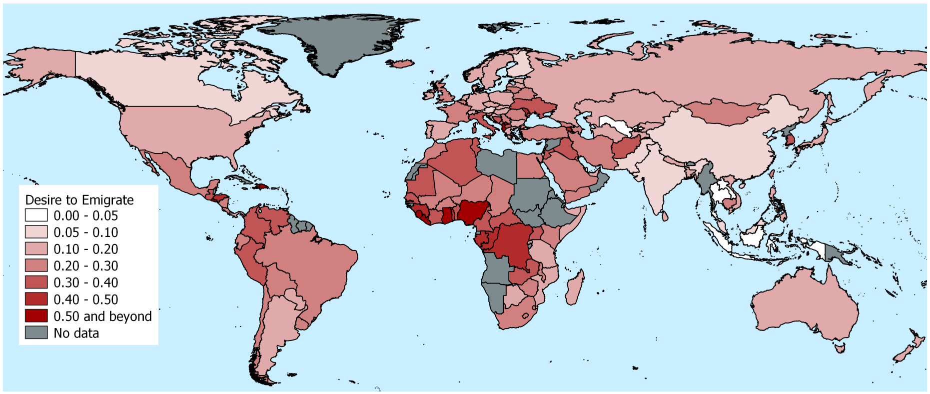

(C) Change in district-level desire to emigrate between early and late years

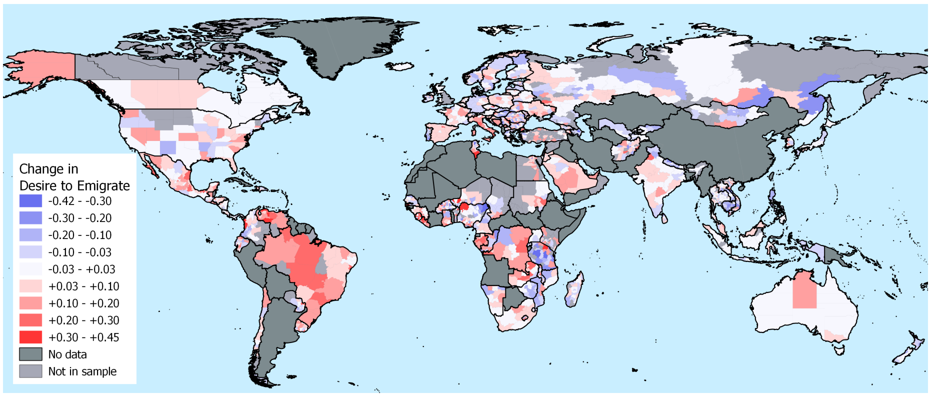

Notes: Panel a) shows the average desire to emigrate on the country-level between 2008 and 2011 and Panel b) between 2015 and 2018. Because not all countries are covered each year, we take the country-level average between 2008 and 2011 for the early years, and between 2015 and 2018 for the late years. The lowest number of observations in Panel a) (Panel b)) is 439 (474). Panel c) shows the difference between the share of respondents desiring to emigrate in the earlier time period (defined as all years before the median year for all observations in a subnational district) and the later time period (defined as all years equal to or exceeding the median year), on the subnational level, in all districts with 3G coverage data and at least 25 GWP respondents per subnational district for both the early and the late time period. All districts with less than 25 observations are shown in light gray and countries without data on 3G coverage are shown in dark gray. All intervals in the legends are closed on the left and open on the right, except for the last bin in Panel c).

FIGURE 1.2: Increase in 3G Coverage Around the World between 2008 and 2018

(A) 3G coverage in 2008 and 2018

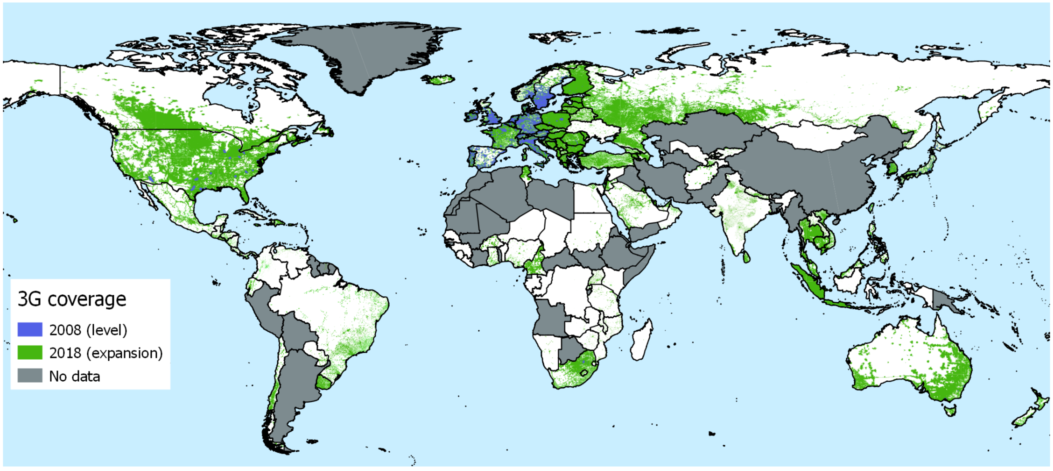

(B) Population-averaged 3G rollout (2008-2018)

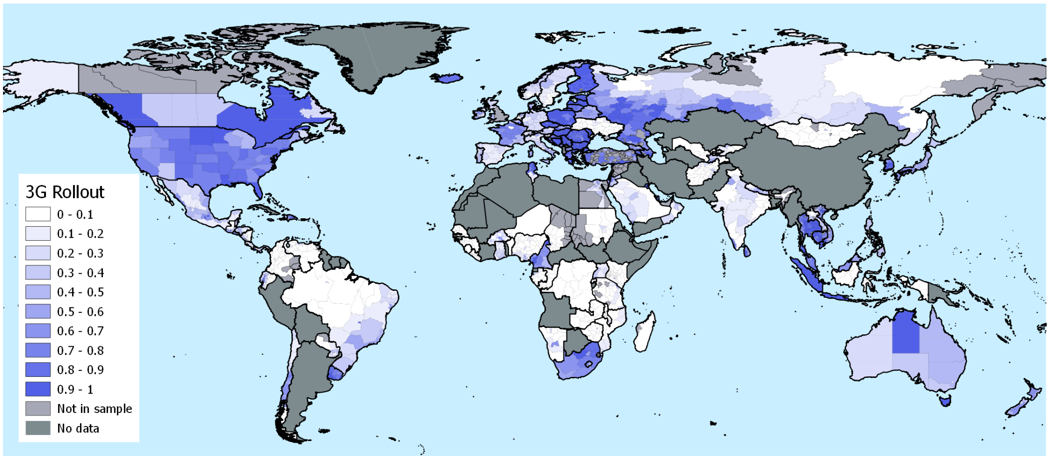

Notes: Panel a) shows the granular 3G coverage in 2008 in blue and the expansion between 2008 and 2018 in green for all data available in the Collins Bartholomew's Mobile Coverage Explorer. Panel b) shows the change in population-averaged 3G coverage between 2008 and 2018 on the subnational district in shades of blue, for those districts available in the GWP. Regions without GWP coverage within countries that are in the 3G data and at least once in the GWP data are colored light grey. Countries fully absent in either the GWP or 3G data are colored dark grey. Districts marked in different shades of blue or remaining white represent the relevant variation in subnational 3G coverage exploited in our empirical strategy. The intervals in Panel b) are closed on the left and open on the right, except for the last bin.

We obtain global data on geo-coded lightning strikes from the World Wide Lightning Location Network (WWLLN).[16] In particular, we use these data to construct an IV following Manacorda and Tesei (2020) and Guriev, Melnikov and Zhuravskaya (2021). The intuition is that cloud-to-ground (CG) lightning is likely to damage the electrical equipment of mobile network towers, which implies a cost of reparation as well as the cost of using additional lightning-protection hardware. Importantly, WWLLN has a good detection efficiency for cloud-to-ground (CG) lightning. This is advantageous over space-based optical detection of lightning, which is most sensitive to intra-cloud (IC) lightning.[17]

Additional Datasets

- Nighttime Light Density: To control for district-level economic development, we follow Henderson, Storeygard and Weil (2012) by using nighttime light density (that is, luminosity at night from satellite images) data as an alternative measure. These data come from Defense Meteorological Program Operational Line-Scan System (DMSP-OLS) and Visible Infrared Imaging Radiometer Suite (VIIRS) instruments.[18] We calculate the district-year-level nighttime light density by weighting light intensity from DMSP-OLS and VIIRS with the GPW population density in 2005. The DMSP-OLS data span until 2013. The VIIRS data are available from 2015 onwards, requiring the year 2014 to be linearly interpolated between the 2013 DMSP-OLS and the 2015 VIIRS datapoint at the district level. As the nighttime light density data come from different sources (and thus are not directly comparable), we normalize each value to a 0 – 1 range within each year.

- OECD: To compare bilateral rates of migration aspirations and intentions with actual migration flows, we obtain migration flow data between 2008 and 2018 (from more than 200 origin countries to 35 OECD countries) from the OECD. In particular, we use the inflows of foreign population by nationality.

- The World Bank: To control for country-level development, we obtain real gross domestic product based on purchasing power parity (GDP (PPP)) per capita per year, expressed in constant 2011 US dollars. We also use country-level population data to construct population weights, as well as the country-level data on broadband subscriptions (per 100 people).

- Center for Systemic Peace: To control for political regime characteristics, we use the Polity2 variable from the Polity IV data set. Polity score ranges from -10 to +10, with -10 to -6 corresponding to autocracies, -5 to 5 corresponding to anocracies, and 6 to 10 to democracies.[19]

- CEPII GeoDist database: To control for dyadic pair-level factors in a gravity framework, we obtained data on shared languages and pairwise weighted distances.[20]

[16] The WWLLN network detects lightning not through optical, but very low frequency (VLF) signals, which has the advantage of carrying further than optical signals and thus requiring fewer detectors. The WWLLN uses only 20-40 detectors worldwide, as the VLF radiation in the kHz range is detectable thousands of kilometers away. Nowadays, the detection efficiency of powerful (discharges exceeding 30kA) lightning strikes is around 30% and the typical spatial accuracy is in the order of a few kilometers. The detection efficiency of CG lightning by WWLLLN improved during the time span 2005 – 2011 from 4% to 10% due to an increase in the number of VLF sensors (Abarca, Corbosiero and Galarneau Jr., 2010).

[17] IC and CG lightning are not very strongly correlated, and the IC-to-CG ratio varies greatly over latitude (Prentice and Mackerras, 1977).

[18] See details at these links: `https://ngdc.noaa.gov/eog/dmsp/downloadV4composites.html` and `https://ngdc.noaa.gov/eog/viirs/download_dnb_composites.html`

[19] For more details on the Polity IV project, see: `https://www.systemicpeace.org/polityproject.html`

[20] For more details on GeoDist, see Mayer and Zignago (2011).

FIGURE 1.3: Bilateral Desire and Plans to Emigrate are Strongly Related to Actual Migration

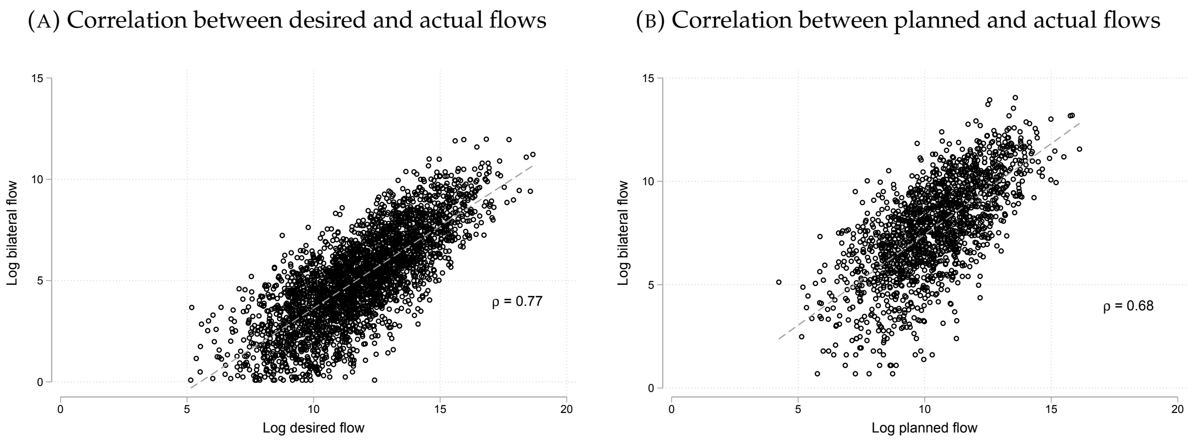

Notes: Scatterplot of the log of the cumulative migration flow versus the log of the estimated population desiring (Panel a); N = 3,055; 2008 – 2018) and planning (Panel b); N = 1,590; 2008 – 2015) to emigrate from 155 origin countries to 35 OECD countries. The constructed flow between an origin and destination is calculated as follows: we multiply the origin-level population in 2008 by the average share of GWP respondents desiring to emigrate to the particular destination country, based on all GWP responses in the origin country between 2008 and 2018. Every data point is an origin-destination pair. Dyads with no actual or desired (planned) flows are omitted in Panel a) (b)).

- United Nations Population Division: To assess the importance of prior dyadic origin-destination country-level stocks of migrants in a gravity framework, we obtained the estimated dyadic stock of migrants from the United Nations Population Division in 2005.[21]

The resulting data set contains 617,402 individuals from 2,120 subnational districts in 112 countries over 11 years of data. There are 13,073 district-by-year cells in the data, implying an average number of 47 respondents in a subnational district in a given year. Although the data is unbalanced, 83% of all districts are present in the data for at least 5 years.

1.4.2 Evidence that Our Treatment and Outcome Variables Convey Meaningful Information

Key to the interpretation of our results is whether our treatment variable (3G) and outcome variables convey meaningful information. To provide evidence of this, we first examine the effects of 3G internet expansion on the individuals' probability of having access to the internet in the full sample.[22] Appendix Table 1.E.3 shows that a full rollout of district-level 3G coverage leads to a statistically significant 4.9 percentage point increase in the likelihood of having access to the internet | thus, the effect of full 3G rollout is about 18% of the baseline average (in 2008, 28% of all GWP respondents reported having access to the internet). This effect is probably an underestimation of the effect of 3G coverage on internet access, as prior to 2016 the question about internet access probes access *at home* only.[23]

[21]For more details on the UN migrant stock data, see United Nations Population Division (2023).

[22]The 'access to the internet' variable is constructed using the following two GWP questions: *Does your home have access to the internet? (2008-2015)* and *Do you have access to the internet in any way, whether on a mobile phone, a computer, or some other device? (2016 – 2018)*

[23]In Appendix Table 1.E.4 we restrict the sample to the subset of individuals that answered the question prior to 2016 about internet access at home. We find no significant effect of 3G coverage on internet access at home, indeed suggesting that the question does not capture mobile internet. The results are similar whether one controls for broadband subscription rate or not.

Second, we check to what extent our outcome variables are statistically significantly associated with actual migration flows. We make use of the fact that we observe individuals' most desired destination as well as the destination country they are planning to move to. We use these data to construct the number of people desiring and planning to migrate between any origin and destination country.[24] We then match our *desired* and *planned* migration-flow matrix with data on actual migration flows to OECD countries between 2008 and 2018 and calculate the log of actual, desired, and planned flows. We omit all dyads with zero actual (166 dyads), desired (1,353) or planned (2,825) flows.[25] The results presented in Figure 1.3 confirm that our outcome variables are strongly associated with the official gross migrant flow data. The correlation on the origin-destination level between the log of actual migrant flows and the log of the number of respondents desiring to migrate from a specific origin to a specific destination is 0.77. The raw correlation between migration flow rates and the number of respondents planning to migrate from a specific origin to a specific destination is 0.68. The latter correlation is weaker, which can arise because the much lower number of GWP respondents planning to migrate than desiring to migrate makes the measure of planned flows noisier.[26] Thus, taken as a whole, we find that our outcomes are strongly positively related to actual migrant flows and, hence, very likely to deliver meaningful information on cross-border movements of people.

Overall, these results suggest that both our treatment and outcomes capture relevant variations in internet access and migration.

1.5 Empirical Strategy

In this section we describe the three complementary estimation strategies that we use to study the effect of 3G coverage on the desire and plans to emigrate. Our main empirical specification is a Two-Way Fixed Effects (TWFE) regression with a continuous treatment that allows exploiting all available variation in mobile network coverage. We complement the TWFE approach by a new estimator by de Chaisemartin et al. (2024) that is robust to heterogeneous and dynamic treatment effects, but does not allow to exploit all available variation in mobile network coverage. Finally, we consider an instrumental variable strategy to dispel concerns about endogeneity of the rollout of mobile internet, using lightning incidence as an instrument to predict subsequent mobile internet rollout.

1.5.1 Main Estimation Model

We estimate the effect of mobile internet access on individuals' migration aspirations and intentions using a difference-in-differences methodology. Our models take the following form:

$$Outcome_{idt} = \beta 3G_{dt} + \boldsymbol{\alpha}'\mathbf{X_{idt}} + \phi_d + \theta_t + \gamma_d \cdot t + \epsilon_{idt} \tag{1.5}$$

where i indexes the individual, d the subnational district, and t the year.

We use answers to questions (1), (2) and (3) from Section 1.4: (1) whether an individual would like to move permanently to another country; (2) whether an individual is planning to move abroad permanently in the next 12 months; and (3) whether an individual is likely to move away

[24]Bilateral flows are constructed by weighting observations within the origin country using Gallup weights to make the data representative at the country level.

[25]For planning to migrate, we only use the time period 2010 – 2015. For 2008 and 2009, the destination country for *planning* to migrate was not allowed to be different from the previously indicated destination country for *desiring* to migrate. As ideal and realistic destination countries may differ, we omit the data from 2008 and 2009.

[26]For a more detailed discussion about the relation between migration aspirations, intentions and realized migration, we refer the reader to Tjaden, Auer and Laczko (2019). In addition to a cross-sectional analysis, they use the GWP and OECD data to show that also temporal variation in migration intentions is predictive of subsequent bilateral migration.

from the city or area in which he or she lives in during the next 12 months. Responses to all three questions are coded as dummy variables, with 1 representing a positive answer and 0 representing a negative answer. Additionally, we use constructed outcomes displayed graphically in Figure 1.A.1 corresponding to the union of outcomes (1) and (3), the intersection between (2) and (3), and the set difference of (3) and (1). We estimate linear probability models for ease of interpretation.

To measure 3G internet coverage, our treatment variable, we follow Guriev, Melnikov and Zhuravskaya (2021) and calculate the share of the district's territory covered by 3G networks in a given year, weighted by population density at each 1x1 kilometer grid-level.[27] This captures an intention-to-treat effect of mobile internet, which includes the direct effect of individuals adopting mobile internet as well as spillover effects.

The vector of controls, X_{idt}, include:

- individual-level demographic characteristics (age and age-squared, a male dummy, an urban dummy, as well as dummy variables for marital status, presence of children in the household, educational attainment and not born in the country of interview);
- log of per capita income of the household;
- satisfaction with life and local amenities; and
- district-year-level average income and country-year-level share of respondents under 30, political regime as measured by Polity2 and log of GDP per capita.

Of course, one might worry that some of the control variables (such as household income or satisfaction with local amenities) are themselves affected by 3G-related economic shocks. In Table 1.1, we dispel concerns about "bad controls" (Angrist and Pischke, 2008) by adding these characteristics gradually. Doing so barely changes the point estimate for our variables of interest. Nevertheless, we keep these controls in our main specification to alleviate concerns related to omitted variable bias.[28]

In all models, we include year dummies, θ_t, (to capture the impact of global shocks that affect all countries simultaneously), district dummies, ϕ_d, (to control for time-invariant variation in the outcome variables caused by factors that vary across districts) and district-specific linear time trends, $\gamma_d \cdot t$, (to remove distinctive trends in outcome variables in various districts that might otherwise bias our estimates if they accidentally coincided with 3G internet-related changes). In the fully saturated models, the estimates are identified by exploiting within-district variation that has been stripped of any influence of constant and linearly changing district-level characteristics.

We two-way cluster standard errors by country-year and subnational district and use sampling weights provided by Gallup to make the data representative at the country level. For all outcomes related to "plans to migrate", we restrict our sample to those who are adults or become adults within one year ($\geq$17 years) as minors usually do not have the ability and/or legal right to decide on migration within 12 months.

Adressing Threats to Identification

One can imagine several potential threats to identification. We address these as follows:

1. To alleviate concerns that the parallel trends assumption may not hold around an increase in 3G coverage, we check whether districts display similar pre-trends in terms of outcomes. We compare the trend between districts that (i) are about to get treated with 3G coverage

[27] As, for the years 2011 to 2018, coverage data is updated until December, we use the known coverage in December $t-1$ to represent the 3G coverage in year t. For further discussion about the 3G data and its timing, see Section 1.4.

[28] We omit smaller subgroups of the included controls in Appendix Table 1.E.6 to show that separate omission of being able to count on friends, satisfaction with local amenities and life satisfaction does not alter the results.

and (ii) are not yet or will not be treated. We provide evidence by two event studies: one around any first increase in 3G coverage and one around large increases (at least 50 percentage points in one year) in 3G coverage.[29] The results indicate parallel trends prior to 3G adoption.[30] In addition, we also show that leads (that is, future levels of 3G coverage) do not affect current migration aspirations.

2. We also include district-specific linear time trends, which remove variation in within-district movements in migration intentions and desires due to factors that are district-specific and that evolve linearly over time. In the fully saturated models, the identification comes from 3G expansions that entail deviations from pre-existing district-specific trends (see Besley and Burgess (2004) for a similar application). As suggested by Angrist and Pischke (2008), after including a parametric trend, the identification hinges on there being a marked change in the outcome in the year of the treatment. Following Autor (2003), we also conduct an F-test of the hypothesis that the district-specific trends are jointly zero. This hypothesis is strongly rejected by the data (the p-value for this test of joint significance is 0.00). We, therefore, keep linear trends in our specifications.[31]

3. Several other factors could potentially affect 3G internet access and migration aspirations simultaneously, net of a linear local time trend. We, therefore, control for a wide range of observable factors (such as the economic development level of districts) as listed above as well as fixed effects to address potential omitted variables concerns.

4. Although we fully saturate our specifications with fixed effects and linear trends, there could still be other omitted variables that are correlated with 3G internet access. To address this concern, we use the methodology developed by Oster (2019). The results suggest that our findings are unlikely to be driven by omitted variables bias.

5. Another concern is that the expansion of 2G infrastructure can affect individuals' migration behavior. As 2G technology only allows for calling, texting and a very limited internet connectivity, it is different from 3G technology. However, as 2G and 3G networks rely on similar infrastructure, expansion of both types of networks may coincide. To ensure that our results are not driven by simple communication technologies, we show that 2G network coverage has no impact on our outcomes.

6. We also conduct multiple hypothesis testing by employing a randomization inference technique recently suggested by Young (2019). In particular, this adjusts for the fact that we are testing multiple hypotheses simultaneously and controls for the tendency for false positives. The method builds on repeatedly randomizing the treatment variable in each estimation and comparing the pool of randomized estimates to the estimates derived via the true treatment variable. The results presented in Table 1.E.13 in Online Appendix 1.B show that our findings remain robust both for the individual coefficients and the joint tests of treatment significance.

All of these and additional identification-related issues are addressed in more detail in the Results section and in Online Appendix 1.B.

[29] In the second event study, we focus on the subnational districts with a large increase in 3G coverage, although these constitute only around 25% of the sample. The vast majority of the remaining 75% of districts shows a more gradual increase in 3G coverage, where testing pre-trends is more challenging. The event study around such large increases is nevertheless complementary to the first event study. It allows us to study whether districts obtaining large hikes in 3G coverage are on different trends in migration aspirations than those that receive a large increase later.

[30] We present the first event study in section 1.6.2, where we do not reject the null hypothesis of joint insignificance of 4 years of pre-trends (p = 0.63). We also do not reject the null hypothesis for the second event study presented in Online Appendix 1.B (p = 0.22).

[31] In Appendix Table 1.E.17, we also show that our results are robust to *not* including district-specific trends.

1.5.2 An Alternative to Two-Way Fixed Effects Estimators

TWFE models are suitable for estimating average treatment effects on the treated (ATT) in the case of homogeneous and non-dynamic treatment effects. By decomposing the TWFE estimator under various assumptions, a recent literature has shown that the TWFE estimator is problematic in the presence of heterogeneous[32] and dynamic[33] treatment effects (Goodman-Bacon, 2021; Borusyak, Jaravel and Spiess, 2024; Callaway and Sant'Anna, 2021; de Chaisemartin and d'Haultfoeuille, 2020; Sun and Abraham, 2021).

To enable the estimation of treatment effects on the treated in the presence of heterogeneous and dynamic treatment effects, one needs to carefully select treatment and control groups. The estimators of Callaway and Sant'Anna (2021) and de Chaisemartin and d'Haultfoeuille (2020) use both never treated and not yet treated units to assess the contemporaneous and dynamic treatment effect.[34] de Chaisemartin et al. (2024); Callaway and Sant'Anna (2021); Borusyak, Jaravel and Spiess (2024); Sun and Abraham (2021) implement alternative estimators that identify the ATT by calculating treatment effects using appropriate control groups. The estimator of de Chaisemartin et al. (2024) is most suitable for our purpose as it allows to study non-binary treatments.

We discuss the implementation of the de Chaisemartin-D'Haultfœuille estimator, which is based on pairwise difference in differences, in Online Appendix 1.C. Importantly, the estimator calculates DiD_l, the treatment effect after obtaining treatment for the first time l periods ago, using a weighted average of the elementary building blocks $DiD_{g,l}^{ini}$ (where g refers to the unit receiving treatment – in our case the subnational district). This is the covariate-adjusted difference between treated and appropriate control units in the differences in outcome over l periods for treated units that obtained first treatment at time F_g and where treated and control units have initial treatment ini. In other words, treated units are only compared to control units with the same initial treatment ini. In a similar fashion, we calculate the pre-treatment difference in differences DiD_l^{pl}, which allows us to assess pre-trends between the same treatment and control units between l periods and one period before receiving treatment. We have to make the following two approximations to be able to calculate DiD_l and DiD_l^{pl} for a sizable part of our sample:

- **Define a threshold $\Delta 3G$, below which treatment between two consecutive years is stable.** As many districts show small increases over time, at the end of the sample period in 2018, most districts saw some increase in 3G coverage. Thus, to have sufficient number of control units for calculation of all $DiD_{g,l}^{ini}$, we need to consider units that have received minimal treatments as untreated.

- **Divide the sample into initial treatment groups *ini*.** The initial treatment is the level of 3G coverage in 2008. However, 3G coverage is continuous, which means that, apart from the districts not yet treated in 2008, all other districts have a unique level of initial treatment. To

[32] In the case of heterogeneous treatment effects, the problem arises because the estimated $\hat{\beta}_{TWFE}$ is a weighted average of group time-level average treatment effects, where the weights are unequal over groups and time, and may be negative. In a general design, weights are more likely to be negative for periods in which many groups are treated and to groups treated for many periods (de Chaisemartin and d'Haultfoeuille, 2020). In a staggered adoption design (a setting where units can move into, but not out, of a binary treatment with heterogeneous timing between groups), this implies that weights on later time periods are more likely to be negative (Borusyak, Jaravel and Spiess, 2024).

[33] When considering a setting with two time periods and one treatment (treatment status changes by one unit) and one control group (treatment status is unchanged), the possibility of dynamic effects requires one to account for the prior path of treatment and control group. Intuitively, a TWFE regression does not control for the complete past treatment history, and is thus not robust to dynamic effects. Similarly, Sun and Abraham (2021) show that the pre- and post-event effect estimates in the canonical event study setting may mix, leading to incorrect estimates of pre-event trends, as well as the instantaneous and dynamic effect of treatment.

[34] A unit that received treatment previously may carry dynamic treatment effects and may thus be unsuitable as a control unit.

be able to match treatment and control units to calculate $DiD_{g,l}^{ini>0}$, we bin initial treatments in groups $ini = 0$ and $ini > 0$. As the estimator of de Chaisemartin et al. (2024) performs covariate adjustment within initial treatment bins, we need the number of observations per bin to be sufficiently large. Therefore, we cannot allow for splitting up the $ini > 0$ bin in multiple bins.[35]

The estimator computes treatment effects in the outcome Y for all l periods after obtaining first treatment (DiD_l). In a similar fashion, we can calculate DiD_l for the treatment variable itself, which simply tells us how much the treatment increased l periods after a district received treatment for the first time. Using the DiD_l for both the treatment as well as for the outcome variable, we can calculate an average effect size of an increase in treatment from 0% to 100%, $\hat{\delta}^L$. In the absence of treatment heterogeneity, dynamic effects and the approximations discussed above, this corresponds to the point estimate β of the TWFE estimator.

However, there is a trade-off as TWFE has an advantage of using all information available in a continuous treatment while the estimator by de Chaisemartin and d'Haultfoeuille (2020) focuses on changes in outcome around first increases in 3G coverage. Furthermore, to find suitable control groups, one needs to define a threshold of stable treatments, which disregards some of the information available in our treatment. Therefore, we consider TWFE as our main specification and use the de Chaisemartin and D'Haultfœuille estimator as a complementary approach.

1.5.3 Instrumental Variable Strategy

Lightning Incidence and Slower 3G Rollout

A potential concern about our main analysis is the endogeneity of 3G network and that decisions on where and when 3G infrastructure is built may be correlated with factors that directly influence migration aspirations. It is plausible that 3G network is rolled out earlier in districts with better economic prospects, more competent administration, and better infrastructure that reduces construction costs. As better economic prospects and infrastructure reduce push factors to emigrate, this could result in an underestimate of the effects of 3G coverage on the desire to emigrate. To address such concerns, we use an IV strategy following Manacorda and Tesei (2020) and Guriev, Melnikov and Zhuravskaya (2021).[36] In particular, Manacorda and Tesei (2020) use spatially differential incidence rates of lightning strikes as a source of exogenous variation in mobile network expansions to study the role of mobile communication in political mobilization in Africa.[37] In the global context, Guriev, Melnikov and Zhuravskaya (2021) adopt a similar instrument using worldwide lightning data from very low frequency (VLF) radiation detectors on a 1x1 km resolution from the WWLLN project.

The intuition of the lightning-based instrument is that electromagnetic discharge due to lightning in or around a base transceiver station (BTS) can damage the antenna and telecommunications equipment, thus requiring repair. Appropriate earthing and shielding of electrical equipment and the use of power surge-protection devices can mitigate this, but come at a substantial cost. Both the cost of repair and the cost of protective measures increase the cost of operating

[35] Covariate adjustment in the estimator of de Chaisemartin et al. (2024) is performed within an initial treatment bin, by performing a regression on the sample of groups where the treatment did not increase (by more than the threshold) yet. Therefore, we need to have at least as many (not yet treated) group-time periods as we have covariates and time periods (as we include time fixed effects). For more details, see Online Appendix 1.C.

[36] Instruments for traditional cable internet connections are often based on the positioning of main ('backbone') internet cables that offer large bandwidth (Hjort and Poulsen, 2019).

[37] Manacorda and Tesei (2020) use optical detection-based NASA data, which is available on a 55x55 km spatial resolution, but this is unavailable for higher latitudes.

mobile networks. As the expected likelihood of lightning in a given region is known, it is plausible that investments in mobile network coverage by operators are deterred in areas with a higher incidence of lightning.

Following Guriev, Melnikov and Zhuravskaya (2021), we focus on lightning strikes from WWLLN between 2005 and 2011 to alleviate concerns of lightning patterns in later time periods being affected by climate change.[38] The WWLLN project documents lightning at the single geo-coded and time-stamped lightning strike level, which we weight by population density and aggregate to the subnational district level. We first determine whether a lightning strike occurred in a one square kilometer box in the grid of the GPW population density data. $L_{box,day,d}$ is a dummy variable indicating whether a lightning strike occurred in a one square kilometer grid cell in a given *day* in a given year in a subnational district d. $P_{box,d}$ is the population in the one square kilometer grid cell in district d in 2005 and $P_d = \sum_{box} P_{box,d}$ is the total population of the district. Then, we aggregate the lightning incidence over all days of the years 2005 to 2011 and all one square kilometer cells in the district:

$$\mathcal{L}_d = \frac{1}{P_d} \sum_{box} \sum_{day} L_{box,day,d} P_{box,d}$$

Assuming that protection measures largely mitigate the damage of lightning strikes, the cost of lightning for a given location is a concave function of lightning strike intensity, which we operationalize by assuming a logarithmic relation (all districts are large enough to have at least one lightning strike during our period of analysis). While Guriev, Melnikov and Zhuravskaya (2021) divide districts into global below and above median lightning districts, we use a log transformation to have a continuous measure of lightning strike intensity. This captures that lightning strikes are very unevenly distributed, with the median of strikes recorded by WWLLN at 0.2 lightning strikes per square kilometer per year, lower quartile at 0.07, upper quartile at 0.82, and the highest percentile at 8.2. We interact $\log(\mathcal{L}_d)$ with a linear time trend to construct our instrument: high lightning frequency districts expand 3G networks more slowly because of the expected additional cost of power surge protection and repairs from lightning damage. Exploiting the differential response of countries with different levels of development, we differentiate between countries with below- and above median GDP per capita, as in Guriev, Melnikov and Zhuravskaya (2021).[39]

A drawback of this IV approach is that we cannot include district-level time trends. Instead of these trends, we include interactions of a large set of district-level geographic and demographic variables with a linear time trend into $\mathbf{X_{dt}}$. Importantly, as initial 3G coverage strongly predicts subsequent coverage, we include interactions between a linear time trend and (i) the level of 3G coverage in the district in 2008, (ii) a dummy variable for zero 3G coverage in the district in 2008, and (iii) a dummy variable for zero 3G coverage in the country in 2008. This prevents our instrument from capturing the effect of initial 3G coverage on the subsequent expansion of 3G networks. Moreover, lightning patterns are likely to be correlated to geography and demography, both of which plausibly impact mobile network expansion.[40] Therefore, it is necessary to control for the effect on 3G expansion of factors such as population density, and the share of land covered by deserts and mountains. As both the instrument as well as 3G coverage vary at the

[38] As the sign of the effect of climate change on global lightning rates is subject to academic debate (Finney et al., 2018) and thus plausibly not anticipated by mobile network operators, it is most likely that network operators base such decisions on historical patterns.

[39] To be transparent we use exactly the same division as Guriev, Melnikov and Zhuravskaya (2021). They calculate the average GDP per capita between 2008 and 2017 using all countries present in the World Bank data and include the median country in the above-median group.

[40] For an overview of the effects of geography and demography on 3G and 4G network expansion in the United Kingdom, see: `https://www.ofcom.org.uk/__data/assets/pdf_file/0027/146448/Economic-Geography-2019.pdf`.

subnational district level, we run the IV at the district-year level. To make it comparable to our baseline estimates, we weight each subnational district by year observation with the number of individuals and the weights provided by GWP. We estimate the following first stage:

$$3G_{dt} = \gamma_1 \text{Below-median GDP pc}_c \cdot \log(\mathcal{L}_d) \cdot t + \gamma_2 \text{Above-median GDP pc}_c \cdot \log(\mathcal{L}_d) \cdot t + \boldsymbol{\alpha}'\mathbf{X_{dt}} + \phi_d + \theta_t + \epsilon_{dt} \quad (1.6)$$

The notation follows that of Equation 1.5. In addition, Below-median GDP pc_c and Above-median GDP pc_c are binary indicators for above- and below-median GDP per capita in line with Guriev, Melnikov and Zhuravskaya (2021), $\mathcal{L}_d$ denotes the natural logarithm of the population-averaged number of lightning strikes per square kilometer and t denotes the year. Using the OLS estimates obtained by estimating Equation 1.7, we obtain the predicted 3G coverage ($\hat{3G}_{dt}$). We use the following second stage to obtain our IV estimates:

$$Outcome_{dt} = \beta \hat{3G}_{dt} + \boldsymbol{\alpha}'\mathbf{X_{dt}} + \phi_d + \theta_t + \epsilon_{dt} \quad (1.7)$$

For the exclusion restriction of our IV approach the hold, the interaction of the log of lightning density and time should not be correlated to the desire to emigrate through other ways than mobile internet expansion, after conditioning on the controls and fixed effects. A potential threat to this exclusion restriction is that the instrument is correlated to other push factors of migration. For this to be the case, places with a higher *level* of lightning incidence should have differential *trends* in those push factors. We tackle these potential concerns by three types of controls: technological, geographical and demographic. First, by interacting the three measures of initial network expansion with linear trends, we capture the long-term effects of the initial delay of communication technologies due to higher lightning density. Second, by interacting geographic factors with linear time trends, we capture potential differential trends in areas with different geographic characteristics that are correlated to lightning intensity. Third, by interacting quintiles of population density factors with linear time trends, we capture differences in trends between areas with different population densities, which may have arisen because of lightning or related meteorological phenomena.

The construction of instruments separately for two income groups to identify a local average treatment effect (LATE) is important for two reasons: relevance and monotonicity.

Relevance: As the potential financial benefits from extending 3G coverage are greater in wealthier districts than in poorer districts, a higher level of anticipated lightning-induced cost is less likely to lead to lower investment in 3G network in wealthier districts than in poorer districts. Constructing separate instruments thus improves the relevance of the instrument.

Monotonicity: Allowing the effect of lightning to vary for various groups is important for satisfying the monotonicity assumption to identify a LATE (Angrist and Pischke, 2008). For example, before the start of our sample in 2008, wealthier countries may have expanded 3G coverage predominantly in districts with lower lightning frequency. Therefore, high lightning frequency districts may even see a stronger increase to catch up to the surrounding districts, given all other characteristics. It is thus important to allow the slope of the instrument to differ for different groups in the first stage. This can be partially controlled for by controlling for initial 3G coverage in 2008 interacted with a linear time trend, but low initial 3G coverage may be driven by many unobserved factors that are more fundamentally limiting factors than high lightning frequency.

1.6 Results

In this section, we present four sets of results. First, we present our baseline results on the effects of 3G rollout on migration aspirations and intentions using the OLS estimator with two-way fixed effects. Thereafter, we focus on the desire to emigrate and we present results for the de Chaisemartin-D'Haultfœuille estimator for a non-binary treatment, and an IV strategy. Ultimately, we analyze heterogeneity of the effects.

1.6.1 Main Results

Table 1.1 reports estimates of Equation 1.5 for the three main outcomes. The dependent variables are binary variables indicating that the respondent "if he/she would have the opportunity, would like to move permanently to another country" (first panel) and that the respondent "is planning to move permanently to another country in the next 12 months" (second panel), Column (1) reports estimates with district and year fixed effects and district-specific time trends. Column (2) adds the demographic characteristics, Column (3) adds controls related to life satisfaction and logarithm of household income per capita and average district-level household income per capita (to control for regional development), Column (4) adds country-level controls, Column (5) fully saturates the specification with country-by-income quintile and country-by-educational-attainment fixed effects to control non-parametrically for all potentially omitted variables that can vary across countries and income quintiles, and countries and educational attainment levels.[41]

All Columns show a positive, statistically significant relationship between 3G mobile internet expansion and desire and plans to emigrate. The most conservative estimates in Column (5) restrict all variation to within-country income quintile and within-country educational attainment. These estimates are similar in magnitude to the first 4 Columns. In our preferred model in Column (4), we find that a 10 percentage point increase in 3G coverage leads to a 0.27 percentage point increase in desire to emigrate, and a 0.09 percentage point increase in plans to emigrate in the next 12 months. Given that the mean levels of these outcome variables are 22.0% and 2.8%. the effects are sizable: the implied aggregate effects of a move from 0% to 100% coverage are about 12% of the average of desire to emigrate permanently and 33% of the average of plans to emigrate permanently. Our estimates suggest that an increase in 3G coverage of 36% (corresponding to the average increase in 3G coverage between 2008 and 2018 across the 2,120 subnational districts in the sample, weighted by Gallup's nationally representative weights) leads to an increase in the desire to emigrate permanently by 0.98 percentage points (95% confidence interval: 0.21 to 1.81). Correspondingly, such an increase in 3G coverage leads to an increase in the share of population planning to emigrate permanently by 0.32 percentage points (95% confidence interval: 0.00 and 0.64).

To corroborate that the results are driven by an increase in the proportion of individuals covered by 3G networks, we show the non-parametric effect of the 3G rollout on the desire to emigrate, net of all baseline controls and fixed effects, in Figure 1.4. The figure suggests that the effect is relatively homogeneous in the intensive margin of 3G coverage. This is in line with the way the 3G coverage variable is constructed, as it represents the share of population covered by 3G.

As some individuals already had broadband access prior to receiving 3G coverage and others did not, we can consider whether 3G rollout has a differential effect if one already has internet access at home. This allows us to study whether 3G access is a substitute to broadband internet

[41] Instead of controlling for the country-year level share of population under 30 years, we can control for its district-level counterpart (calculated using the reported age in GWP). This changes the point estimate by less than 0.0003. In addition, controlling for the district-year average unemployment rate also changes the point estimate by less than 0.0003, but would require us to drop 36,992 observations.

TABLE 1.1: The Effects of 3G Rollout on Desire and Plans to Emigrate

	(1)	(2)	(3)	(4)	(5)
Outcome:	Desire to emigrate (I-IV)				
3G	0.027**	0.026**	0.028***	0.027**	0.026**
	(0.012)	(0.012)	(0.011)	(0.011)	(0.011)
Observations	617,402	617,402	617,402	617,402	617,402
R^2	0.12	0.16	0.19	0.19	0.19
Average dependent variable	0.220	0.220	0.220	0.220	0.220
First year	2008	2008	2008	2008	2008
Last year	2018	2018	2018	2018	2018
Outcome:	Plans to emigrate in the next 12 months (I+II)				
3G	0.008**	0.009**	0.009**	0.009**	0.008*
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Observations	379,703	379,703	379,703	379,703	379,703
R^2	0.06	0.07	0.07	0.07	0.08
Average dependent variable	0.028	0.028	0.028	0.028	0.028
First year	2008	2008	2008	2008	2008
Last year	2015	2015	2015	2015	2015
District and year fixed effects	✓	✓	✓	✓	✓
District-year trends	✓	✓	✓	✓	✓
Demographic controls		✓	✓	✓	✓
Life satisfaction-related controls			✓	✓	✓
Income controls			✓	✓	✓
Country-year-level controls				✓	✓
Country×income quintile fixed effects					✓
Country×education fixed effects					✓

Notes: OLS regressions. Standard errors in parentheses. This table reports the results of Equation 1.5 using the questions on desires and plans to emigrate questions in GWP. The demographic controls include: male dummy, age, age squared, dummy variables for marital status (with partner, separated/divorced/widowed, singles serve as the reference group), the presence of children in the household, living in an urban area, educational attainment (secondary educated, tertiary educated, primary educated serve as the reference group) and a dummy for whether the respondent is not born in the country. Life satisfaction-related controls include: satisfaction with housing, healthcare, education, roads, transportation, city, life and whether the respondent can count on family or friends, whether the respondent believes they will be financially better off in five years, whether the respondent has sufficient means for food and shelter, and whether the respondent had something stolen in the past year. Income controls include the log of household income per person on the individual level and the log of the average of household income per person on the subnational region year-level. Country-year-level controls include: the log of real GDP per capita, polity2 score and the share of respondents aged under 30. The standard errors are clustered two-way on the country-year and district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

or whether 3G enables new ways of interaction with the internet that could be important for migration aspirations, such as social media. For this analysis, we rely on the GWP question about internet access at home (see section 1.4.2 on the discussion of this question, which predominantly captures broadband internet) available prior to 2016. Columns (1) to (3) of Appendix Table 1.E.5 show that both 3G and internet access at home are significantly associated to the desire to emigrate. In Column (4) we include the interaction between these two and do not find a significant interaction effect. This suggests that the effect of 3G on the desire to emigrate is not driven by only those newly connected to the internet, but also by those already connected through a broadband connection.

We also follow the method proposed by Oster (2019) to investigate whether our results could be driven by unobservable factors. Oster's δ indicates the degree of selection on unobservables, relative to observables, needed for our results to be fully explained by omitted variable bias. Appendix Figure 1.D.2 shows Oster's δ for different values of the maximally admissible variation R^2_{max}, including the value recommended by Oster (2019) of 1.3 times the baseline R^2. The high absolute values of delta are reassuring: given the controls we have in our models, it seems unlikely that unobserved factors are 6 times more important than the observables included in our preferred specification, which makes it highly unlikely that our results can be explained by omitted variables bias.[42]

FIGURE 1.4: The Non-parametric Effect of 3G Rollout on the Desire to Emigrate

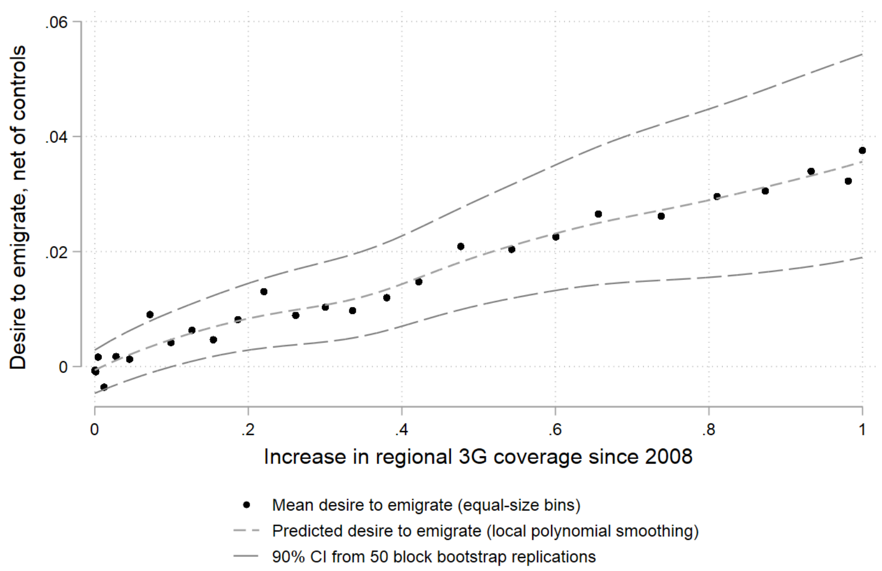

Notes: The figure shows a binned scatterplot of the individual-level desire to emigrate net of all controls by 40 bins of equal sample size, ordered by the relative increase in 3G coverage since 2008. 14 out of 40 bins are concentrated at a 3G coverage of 0, as these observations do not have coverage yet. The middle dashed line displays a local polynomial smoothing through the 40 points, and the outer dashed lines show a smoothened 90% block bootstrapped confidence interval using 50 bootstrap replications. Bootstrap samples are drawn in clusters of subnational districts and country-years, corresponding the the levels of clustering in our baseline specification.

Tables 1.1 and Figure 1.4 show that 3G internet coverage has a positive, sizable and statistically significant effect on the desire and plans to emigrate. Appendix Table 1.A.1 shows the effects of the rollout of 3G internet on a set of outcomes based on the questions on desires and plans to emigrate, as well as a question concerning the respondents self-assessed likelihood

[42]The rule of thumb to be able to argue that unobservables are unlikely to fully explain the treatment effect is for Oster's δ to be over the value of one (Oster, 2019).

he/she moves away from the area he/she currently lives in. The results in the last panel of Appendix Table 1.A.1 show that there is no statistically significant effect on the perceived likelihood of domestic migration, conditional on an individual not desiring to emigrate. This finding suggests that 3G expansion shapes emigration intentions and plans rather than domestic migration. This is intuitive as, even in the absence of internet connectivity, people are likely to be already well-informed about opportunities in their own country as opposed to opportunities in other countries.

1.6.2 de Chaisemartin and D'Haultfœuille Estimator and Testing for Pre-trends

In this section, we examine the validity of the pre-trends assumption and the properties of our TWFE regressions as the impact of 3G expansion is likely to vary across districts and over time. In particular, weight decompositions of group time-level treatment effects suggest that our results in Table 1.1 are susceptible to treatment effect heterogeneity.[43] To investigate whether our results are driven by this potential bias, we use a novel estimator by de Chaisemartin et al. (2024), which is valid even if the treatment effect is heterogeneous.

We proceed as follows to have a sufficient number of observations in every initial treatment group and a sufficient number of observations to include all baseline covariates.[44] We assign subnational districts with non-zero initial treatment in 2008 (that is, $ini > 0$) to a single bin. We also omit districts where treatment is not monotonically increasing. As many districts show minor decreases in coverage, we only omit districts where treatment decreases by more than 3 percentage points of population between any two subsequent years between 2008 and 2018.[45] To have a sufficient number of untreated and not-yet-treated observations in later time periods, we set the threshold for a first switch into treatment, Δ_{3G}, to 3 percentage points of population. We opted for a threshold of 3 percentage points as the largest proportion of minor increases is concentrated below 3 percentage points; results are qualitatively similar if using 2 or 5 percentage points as the threshold. A drawback of a higher threshold Δ_{3G} is that the never treated and not-yet-treated groups include districts that experienced small increases in treatment.

In Figure 1.5, we show the instantaneous and four dynamic estimators (referring to one, two, three, and four years after the expansion), DiD_l (where $l \geq 0$), and four placebo estimators (referring to two, three, four and five years before the expansion, with respect to the year), DiD_l^{pl} (where $l \geq 2$). We first show these estimators for our treatment variable $3G$ in Panel A (DiD_l^{3G}) and for the outcomes of interest Y in Panels B and C (DiD_l^{Y}).[46] The confidence interval of the placebo estimators should enclose 0 to support the parallel trends assumption.[47] Notably, the results reported in all panels of Figure 1.5 provide *no evidence* of pre-trends.

We turn next to the evolution of post-treatment effects. In Panel A, we find that 3G coverage increases steadily over time after the initial jump. In Panel B, we observe that the desire to emigrate increases immediately after an initial increase in 3G coverage and then remains stable.

[43] de Chaisemartin and d'Haultfoeuille (2020) developed a procedure (TWOWAYFEWEIGHTS) to calculate how many of the weights on the group time-level treatment effects are negative and what the sum of negative weights is (where all weights sum to unity). Using TWOWAYFEWEIGHTS while allowing for heterogeneous treatment effects, we find that the sum of negative weights for the TWFE regressions featured in Column (4) of the three panels from top to bottom in Table 1.1 are -0.77 and -0.44 (the total sum of weights is +1 by construction). As a substantial portion of the weights is negative, this suggests that our baseline results could be biased.

[44] As covariate adjustment is performed within every initial treatment group, increasing the number of initial treatment groups reduces the number of available observations for covariate adjustment. For a more detailed description of the covariate adjustment procedure, see Online Appendix 1.C.

[45] This happens in 234 out of 2,120 subnational districts.

[46] We use the user-written command DID_MULTIPLEGT in STATA 16. As two-way clustering of standard errors is not possible in this command, we cluster standard errors at the country level. Note that, in Table 1.E.14, we find that clustering at the country level gives somewhat bigger standard errors than our baseline estimates.

[47] To assess whether pre-trends between treatment and control are insignificant between the lth to the last period before treatment, we consider the null hypothesis that all of the placebo estimators are zero.

The instantaneous effect of an increase from no to full 3G coverage is 0.055 ($p = 0.04$), which exceeds our TWFE estimate. The average effect of all observations following a first increase in 3G expansion δ^5 is 0.057 ($p < 0.001$), which also exceeds our TWFE estimate. Two reasons can be underlying the difference between the TWFE and the de Chaisemartin-D'Haultfœuille estimator: (i) TWFE estimators are biased when treatment effects are heterogeneous and (ii) the average effects of the de Chaisemartin-D'Haultfœuille estimator are calculated over the first 5 periods after receiving treatment. As for the first point, the TWFE estimator is more likely to assign a negative weight to periods where a large fraction of groups are treated, and to groups treated for many periods (de Chaisemartin and d'Haultfoeuille, 2020). This may lead to a downward bias in the reported results if treatment effects are larger in the final years of the sample and if units treated for most periods in the sample (e.g. high-income countries) have higher treatment effects. As for the second point, we report the de Chaisemartin-D'Haultfœuille estimator for the first five periods after receiving treatment, but not the later time periods (as only few control groups are left in later time periods). As marginal treatment effects may be decreasing (for example, 3G expansions many periods after first expansion of 3G coverage in a district may draw few new mobile internet users), the de Chaisemartin-D'Haultfœuille estimator could capture the time period with the largest marginal effects. Panel C presents the results for plans to emigrate in the next 12 months. We observe that plans to emigrate increase gradually after receiving treatment. However, the instantaneous effect is not statistically significantly different from 0. The average effect of all observations following a first increase in 3G expansion δ^5 is 0.024 ($p < 0.001$).

FIGURE 1.5: De Chaisemartin-D'Haultfouille Estimates for the Effect of 3G Coverage on Desires and Plans to Emigrate

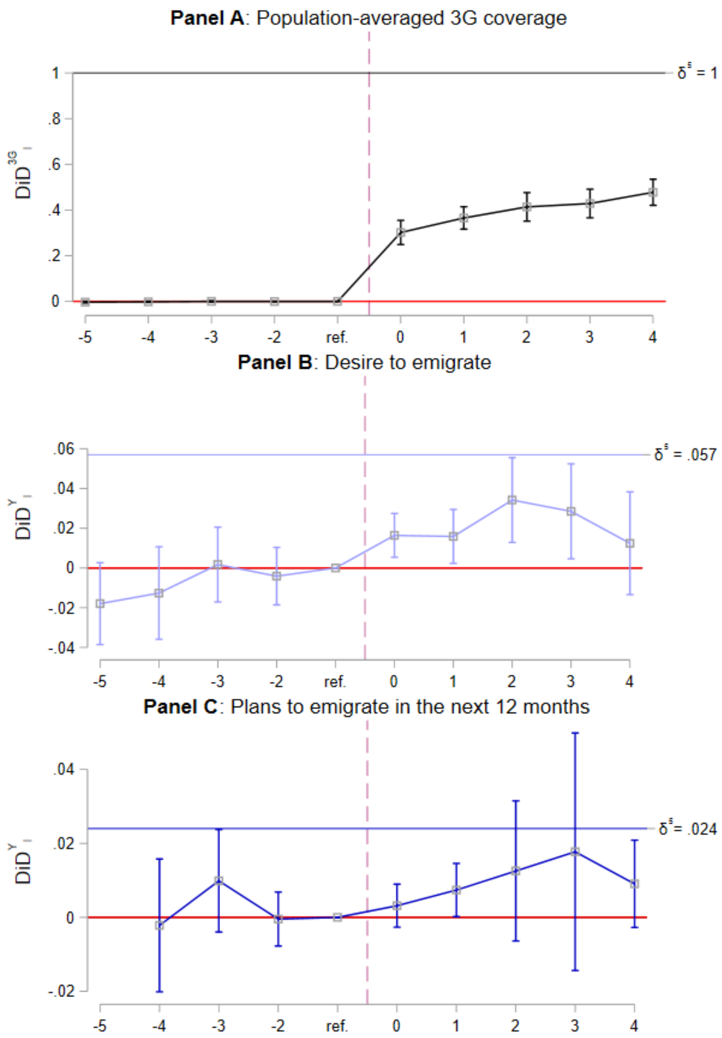

Notes: Event study plots based on the de Chaisemartin et al. (2024) DiD_l estimator. We show DiD_l around first switches in 3G coverage on the subsequent 3G coverage in Panel A and on the three main outcomes in Panels B, C and D. The *p*-value for jointly insignificant pre-trends equals 0.63 in Panel B, and 0.44 in Panel C. δ^5 denotes the estimated average effects of an increase from no to full 3G coverage using the instantaneous effects and the 4 dynamic effects. Per definition this is unity for panel A, as the outcome and the treatment variable are the same. The threshold for a switch into treatment is an increase of coverage of 3% of the population. Treatment and control groups are binned in two groups, those with initial treatment level $ini = 0$ or those with initial treatment level $ini > 0$ in 2008. Observations are weighted using the Gallup weights. After a district switches, an observation l periods after the switch can no longer be part of a control group and is only considered for the lth dynamic effect of its first switch. Note that the placebo estimators l are labelled by $-l$ on the x-axis. As the lth placebo estimator require $2l+1$ years of data (see Online Appendix 1.C), the $l = 5$ placebo estimator for panel C is infeasible. Standard errors are calculated using 50 bootstrap replications, clustered on the country level, 95% confidence intervals are shown.

Taken together, Tables 1 and 2 and Figure 5 show that 3G internet coverage has a positive, sizable and statistically significant effect on the desire and plans to emigrate, but no statistically significant effect on the perceived likelihood of domestic migration. This finding suggests that 3G expansion shapes emigration intentions and plans rather than domestic migration. This is intuitive as, even in the absence of internet connectivity, people are likely to be already well-informed about opportunities in their own country as opposed to opportunities in other countries.

1.6.3 Instrumental Variables Based on Incidence of Lightning Strikes

To alleviate concerns about the endogeneity of 3G network coverage, we instrument 3G expansion. Our instrumental variable is the logarithm of regional population-weighted lightning-strike frequency interacted with a linear time trend as outlined in section 1.5.3. Although our baseline includes district-level linear time trends, we omit those in the IV estimations as our instrument provides us with variation linear in time on the district level. Instead, we introduce a battery of controls related to geography and initial 3G coverage by district in 2008 interacted with a linear time trend.

Table 1.2 reports the two-stage least squares (2SLS) estimates at the subnational district by year level. Column (1) shows the district-year level equivalent of the baseline result from Column (4) of the upper panel of Table 1.1 for comparison, Column (2) reports the reduced form results omitting the district-level time trends in Column (1) but including additional controls, Column (3) reports the second-stage results, and Column (4) shows the first stage coefficients of the regression of 3G coverage on the instrument, separately for below- and above-median GDP per capita countries. The F-statistic is 11.04, which suggests a sufficiently strong first stage. The below-median income countries drive the first stage: districts in those countries with high frequencies of lightning strikes expand their 3G coverage more slowly. In these countries, a 10% higher lightning incidence implies that over 11 years the rollout (in terms of population covered) of 3G networks is 1.1 percentage points lower. For the above-median GDP per capita countries, the first-stage coefficients are statistically insignificant. To alleviate concerns related to the strength of the instrument, we report a 95% Anderson-Rubin (AR) confidence interval. We find that the confidence interval does not enclose zero. To further alleviate concerns about the validity of the exclusion restriction, we show in Appendix Table 1.E.20 that the two instruments do not have a direct effect on the desire to emigrate in the absence of 3G coverage exceeding 3 percent of the population.

In line with our baseline results and the results of the de Chaisemartin-D'Haultfœuille estimator, the IV estimates in Column (3) also indicate that 3G expansion leads to an increase in desire to emigrate. The IV estimate (0.156) is greater in magnitude than the average effects of the de Chaisemartin-D'Haultfœuille estimates (0.057) and the TWFE estimate (0.027). Guriev, Melnikov and Zhuravskaya (2021) use a similar estimation strategy with 3G network coverage as the main explanatory variable and also find IV estimates to be considerably larger than OLS estimates. This can be explained by the endogeneity of 3G network expansion due to expansion of 3G in better developing districts and measurement error in 3G coverage. It is plausible that 3G network is expanded earlier in districts that develop more positively, and where institutional quality and other types of infrastructure are better. Operators are likely to prioritize expanding 3G network in districts where economy is expected to grow faster, but expected faster growth also reduces push factors to emigrate, meaning that our baseline estimates may be downward-biased. Additionally, the TWFE and de Chaisemartin-D'Haultfœuille results may be subject to measurement error: not all mobile phone operators in a country provide network information, not all operators provide it timely or correctly, and 3G coverage may vary within years. Although measurement error may be non-classical, it is unlikely to be correlated with the district-year-level propensity to desire to emigrate.

1.6.4 Heterogeneity Analysis

We also look beyond average effects to understand how the causal effects vary with observable individual and country level characteristics. As we want to keep a flexible specification, allowing different trends for the different demographic groups, we run our fully saturated model as in

TABLE 1.2: Lightning-based IV Results

	(1)	(2)	(3)	(4)
Dependent variable:	Desire to emigrate			3G
Stage:	Baseline	Reduced	IV: second	IV: first
3G	0.028**	0.029***	0.156**	
	(0.012)	(0.009)	(0.078)	
Anderson-Rubin 95% Confidence Interval			*[0.031,0.390]*	
Below-median GDP per capita country $\times$ $\log(\mathcal{L}_d)$ $\times$ year				-0.011***
				(0.002)
Above-median GDP per capita country $\times$ $\log(\mathcal{L}_d)$ $\times$ year				-0.004
				(0.003)
First-stage F-statistic				11.04
Observations	12,929	12,929	12,929	12,929
R^2	0.797	0.732	0.720	0.877
Average dependent variable	0.230	0.230	0.230	0.392
District-level time trends	✓			
IV-related controls		✓	✓	✓
Baseline controls	✓	✓	✓	✓
District and time FEs	✓	✓	✓	✓

Notes: OLS regressions in Columns (1) and (2), 2SLS regression in Column (3) and (4). Standard errors in parentheses. See notes to Table 1.1 for details of the baseline control variables (which include demographic, life satisfaction and income-related controls). 3G expansion is instrumented by the log of lightning strikes per square kilometer on the subnational region level between 2005 and 2011 interacted with a yearly time trend for below- and above median income groups as used by Guriev, Melnikov and Zhuravskaya (2021). The unit of observation is the subnational region as defined in GWP. We aggregate all individual-level outcomes using the Gallup weights to the subnational district by year level. Column (1) shows our baseline estimate on the aggregated level with the same covariates, fixed effects and district-level time trend as Column (4) of Table 1.1. To include the instrument at the district *times* year level, Column (2) omits the district-level time trend but includes interactions of a linear time trend with the following district level variables: five bins of population density, maximum altitude of the district, the share of mountains, the share covered by deserts, the initial population-weighted 3G coverage, a dummy for 3G coverage being 0 in 2008, and a dummy for 3G coverage being 0 in 2008 on the country level. The 2SLS estimation reported in Column (3) and (4) uses the same controls, reporting the second-stage result in Column (3) and the first-stage with an F-statistic of 11.04 in Column (4). The standard errors are clustered two-way in all four columns: on the country-year and district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column (4) of Table 1.1 for separate subsamples.[48] Table 1.3 shows the split-sample results along seven important dimensions. When analyzing heterogeneity over education and employment, we study individuals aged 25–60 to focus on individuals most likely to have completed their education.

We find strong heterogeneity over education, age, employment, and per capita household income, but no statistically significant differences with respect to gender, the presence of children, and living in urban or rural location. The effect of 3G coverage is strongest among secondary educated, those between 25 and 45 years old, unemployed (or working only part-time even though they would like to find full-time job) or out of labor force, and those in the lowest two quintiles of the within-country-year income distribution. In terms of age, it is important to note that although the estimated effect of 3G coverage on desire to emigrate is strongest among those between 25 and 44 years, the mean desire to emigrate is highest among those younger than 25 years.

Why are the effects of 3G coverage non-monotonic in terms of education? Our explanation is that this is likely to reflect differences in potential to migrate and in the level of information available already without 3G mobile internet. Those with only primary education often lack language skills, which could be an important reason why they are both somewhat less likely to desire to emigrate than those with more education, and also respond less to 3G coverage as they are less able to search information on opportunities in foreign languages and even if they could, less likely to find suitable opportunities abroad. Those with tertiary education, instead, are more likely to have the information that the mobile internet offers already at their disposal. The finding that those who are unemployed or underemployed or out of labor force are most responsive to 3G coverage is in line with them having highest potential gains from migration, due to low returns to their skills at home.

Moreover, we explore heterogeneity of the most important individual-level variables over country groups in Table 1.E.21. We find that the effects of 3G coverage are strongest among those with secondary education in both low- and middle-income countries and high-income countries, as well as in non-democratic countries. In democratic countries, the estimated effects are quite similar among those with primary and among those with secondary education. We also find that the effects in countries with below-median unemployment rates are driven by those who are unemployed or involuntarily part-time employed. In countries with high unemployment, the effect is strongest for those out of labor force, which may reflect a significant fraction of those out of labor force having been effectively pushed out of labor force due to negligible prospects of finding employment. That the estimated effect among those who are unemployed in high-unemployment countries is not statistically significant is surprising at the first glance, but may result from this group having the highest average desire to emigrate in the first place, leaving not much scope for further growth/increase.

1.7 Mechanisms

In this section, we discuss potential mechanisms that can explain the relationship between mobile internet access and the desire to emigrate. First, we evaluate the role of the costs of acquiring information. We assess this by considering whether the effect is driven by those who do not have close personal networks abroad and how potential destinations change. Second, we consider whether mobile internet coverage affects perceptions of material well-being, trust in institutions and variables such as life satisfaction, optimism or sense of purpose in life.

[48] We perform split sample regressions over discrete variables, which is identical to properly saturated interaction models. For a discussion on split sample versus interacted heterogeneity analysis, see Feigenberg, Ost and Qureshi (2023).

TABLE 1.3: Heterogeneity of the Effect of 3G Coverage on Desire to Emigrate Based on Individual Level Characteristics

Panel A: Age and Gender					
	(1)	(2)	(3)	(4)	(5)
Sample:	Below 25	25 to 44	45 and above	Female	Male
3G	0.018	0.041***	0.020*	0.031***	0.024*
	(0.018)	(0.015)	(0.011)	(0.011)	(0.013)
Observations	201632	181955	233777	332127	285286
R^2	0.19	0.19	0.15	0.19	0.20
Mean dep. var.	0.328	0.229	0.120	0.204	0.240

Panel B: Education and Urban/Rural Status					
	(1)	(2)	(3)	(4)	(5)
Sample:	Primary	Secondary	Tertiary	Urban	Rural
3G	0.023	0.044***	0.007	0.025*	0.024*
	(0.016)	(0.014)	(0.019)	(0.014)	(0.013)
Observations	120448	205296	71692	245131	372236
R^2	0.22	0.18	0.20	0.20	0.19
Mean dep. var.	0.186	0.224	0.228	0.246	0.204

Panel C: Employment Status and Presence of Children					
	(1)	(2)	(3)	(4)	(5)
Sample:	Employed	Unemployed or Involuntarily Part Time Employed	Out of Labor Force	No Children	Children
3G	0.021	0.060**	0.059***	0.015	0.033***
	(0.014)	(0.025)	(0.020)	(0.013)	(0.012)
Observations	225417	56083	90525	276102	341305
R^2	0.18	0.26	0.22	0.19	0.20
Mean dep. var.	0.205	0.309	0.182	0.193	0.243

Panel D: Quintiles of Within-Country Per Capita Household Income					
	(1)	(2)	(3)	(4)	(5)
Sample:	Lowest Quintile	Second Lowest Quintile	Middle Quintile	Second Highest Quintile	Highest Quintile
3G	0.073***	0.049***	-0.001	-0.006	0.007
	(0.017)	(0.016)	(0.017)	(0.017)	(0.016)
Observations	130136	126366	123390	121307	116102
R^2	0.21	0.22	0.22	0.22	0.21
Mean dep. var.	0.215	0.211	0.217	0.222	0.238

Notes: OLS regressions. Standard errors, clustered by district and country-year, in parentheses. The specification used in Columns (1) to (5) of Panels A to D is the same as that of Column (4) of Table 1.1. Columns 1 to 3 of Panel C groups full-time employed, self-employed, and part-time employed who do not want to work full time in the first category, part-time employed who do want to work full-time and unemployed in the second category and all others in the "out of labor force category". The income quintiles in panel D are drawn per country per year, so that the measure accurately reflects the person's household income position. As employment status and education are poorly defined for young and old individuals we limit the scope of the regressions in Columns (1) to (3) of Panels B and C to those between 25 and 60. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.7.1 Reduced Costs of Information and Networks Abroad

Does internet access substitute for personal networks abroad?

To assess whether internet access decreases information costs, we consider whether the effect size is larger for individuals without first-hand access to information. As the GWP asked respondents whether they had someone abroad to rely on between 2008 and 2015, we can consider the differential effect on the group that has someone to rely on and the group that does not. These close prior networks have been shown to explain a substantial part of the variation in the desire to migrate (Manchin and Orazbayev, 2018).

Table 1.4 shows that the effect of 3G on the desire to emigrate is strong for the individuals without any close personal network abroad, and insignificant for the group with such a network abroad. This is striking when considering the lower baseline level of desire to emigrate for those without close personal network abroad. Only 16.1% of those desire to emigrate, whereas this is 29.8% for those with close personal networks abroad. For those without close prior networks abroad in Column (2), the relative effect size of a full rollout of 3G coverage is 27% of the average share of individuals desiring to emigrate. In contrast, this is only 14% for the full sample in Column (1). This suggests that internet access is likely to affect the desire to emigrate primarily through the cost of information acquisition, substituting for personal networks. If internet access would affect the desire to emigrate primarily by reducing migration costs or communication costs with those left in the home country after migration, then its effects should not depend on close personal networks abroad. Of course, absence of evidence is no evidence of absence: even though the estimated effect of 3G coverage is statistically insignificant among those with close personal network abroad, the point estimate is still positive, and we cannot rule out that internet access would also have an effect through migration costs.

TABLE 1.4: The Effect of 3G Coverage on the Desire to Emigrate According to Close Personal Network Abroad

Those with people to rely on abroad:	(1) All respondents	(2) No	(3) Yes
3G	0.030**	0.044***	0.016
	(0.015)	(0.016)	(0.025)
Demographic controls	✓	✓	✓
Amenities, satisfaction, and income controls	✓	✓	✓
Country-year-level controls	✓	✓	✓
Observations	388,368	252,172	136,130
R^2	0.19	0.18	0.21
Average dependent variable	0.209	0.161	0.298

Notes: OLS regressions. Standard errors, clustered by district and country-year, in parentheses. The specification of Columns (1) to (3) is identical to that of Column (4) of Table 1.1 for the subsample of (1) all respondents who answered the question whether one has someone to rely abroad (asked between 2008 and 2015), (2) only those who have no one to rely on abroad, and (3) only those who have someone to rely on abroad. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Does internet access change preferred destinations?

Because of the existence of prior networks, reductions in the cost of obtaining information are not equally shared across all origin-destination corridors. Origin-destination pairs with strong prior ties are less likely to be affected, as previous migrants from such districts can provide first-hand information to prospective migrants. Therefore, differential changes in information costs could divert migration flows from destination countries with strong prior networks to those without.

TABLE 1.5: The Effect of 3G Coverage and Pre-existing Migrant Networks on Preferred Destinations

	(1)	(2)	(3)	(4)	(5)	(6)
	Desired bilateral emigration					
$3G_{ot}$	0.291***	0.654***	0.065			
	(0.059)	(0.116)	(0.113)			
$3G_{ot} \times \ln(\text{Stock}_{od,2005}+1)$		-0.173***		-0.262***		-0.301***
(Standardized)		(0.061)		(0.059)		(0.072)
$3G_{ot} \times \ln(\text{GDPpc}_{dt})$			0.205**		0.010	0.218*
(Standardized)			(0.104)		(0.102)	(0.117)
$3G_{ot} \times$ Polity IV_{dt}						-0.006
						(0.013)
$3G_{ot} \times$ Common language_{od}						0.035
						(0.123)
$3G_{ot} \times \ln(\text{Distance}_{od})$						-0.004
(Standardized)						(0.053)
Observations	64,977	64,977	64,977	64,977	64,977	64,977
Origin-year-level controls	✓	✓	✓	-	-	-
Origin-destination FE	✓	✓	✓	✓	✓	✓
Destination-year FE	✓	✓	✓	✓	✓	✓
Origin-year FE				✓	✓	✓

Notes: Pseudo-Poisson Maximum Likelihood (PPML) regressions (see Silva and Tenreyro (2006)). Standard errors are clustered at origin-destination level. The dependent variable is the estimated (based on GWP responses) number of people desiring to migrate from a specific origin to a specific destination in a given year. $3G_{ot}$ is the population-averaged 3G coverage in country o in year t. For a discussion about the exact timing of 3G network coverage and the GWP, see section 3.2. $\ln(\text{Stock}_{od,2005}+1)$ is the log of the stock of migrants (plus one) in origin country o in destination d in 2005, $\ln(\text{GDPpc}_{dt})$ is the real GDP (PPP) per capita in destination d at time t, the Polity IV_{dt} integer score ranging from -10 (strongly autocratic) to +10 (strongly democratic), Common language_{od} is a binary indicator for whether the origin and destination countries share a language that is spoken by at least 9% of population in both countries, and $\ln(\text{Distance}_{od})$ denotes the natural log of the population-weighted distances between the origin and destination (the last three are obtained from the CEPII GeoDist database, for details see Mayer and Zignago (2011)). $\ln(\text{Stock}_{od,2005}+1)$, $\ln(\text{GDPpc}_{dt})$ and $\ln(\text{Distance}_{od})$ are standardized such that the means are 0 and standard deviations 1. The regressions in all six columns control for $\ln(\text{Stock}_{od,2005}+1)$, $\ln(\text{GDPpc}_{dt})$, Polity IV_{dt}, Common language_{od} and $\ln(\text{Distance}_{od})$. The additional origin-year-level controls include the unemployment rate, the total population and the polity score in the origin country. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Using the reported desired destination in Gallup, we calculate the number of people desiring to migrate from origin country o to destination country d in year t, as displayed in Figure 1.3. Table 1.5 reports estimates for the effect of origin country 3G coverage on constructed desired migrant flows from 2008 to 2018, where the unit of observation is the origin-destination-year. We estimate the following non-linear gravity equation using the Pseudo-Poisson Maximum Likelihood (PPML) estimator, where the parts in square brackets are included in some but not all columns of Table 1.5:

$$\begin{aligned} m_{odt} = \exp\big(&\beta 3G_{ot} + \boldsymbol{\alpha}'\mathbf{X_{odt}} + \gamma_{od} + \theta_{dt} + \\ &\big[3G_{ot} \times \ln(\text{Stock}_{\text{od},2005}+1) + 3G_{ot} \times \ln(\text{GDPpc}_{\text{dt}}) + 3G_{ot} \times \ln(\text{Distance}_{\text{od}}) + \\ &3G_{ot} \times \text{Polity IV}_{\text{dt}} + 3G_{ot} \times \text{Common language}_{\text{od}} + \phi_{ot}\big]\big) + \epsilon_{odt} \end{aligned} \quad (1.8)$$

where o indexes the country of origin, d indexes the preferred country of destination and time t the calendar year. $3G_{ot}$ is the population-average 3G mobile network coverage on the origin-year level.

Column (1) reports the effect of 3G access on the number of people desiring to migrate (from a specific origin to a specific destination). Moving from 0% to 100% 3G coverage increases desired bilateral migration flows by 29%, on average. In Column (2), we include an interaction between the log of the stock of migrants from origin o in destination d in 2005. We find that the effect of 3G on desired emigration is reduced for those dyads with a large prior stock of migrants. A one standard deviation larger log of dyadic migrant stock is associated with a 17% smaller effect of 3G on bilateral preferred flows. As an example, moving from the dyad Armenia-Spain (there were 7,200 Armenian migrants in Spain in 2005) to the dyad Argentina-Spain (there were 244,000 Argentinian migrants in Spain in 2005) corresponds to a one standard deviation difference in the log of prior stock of migrants.[49] In Column 3, we use a similar specification now interacting PPP GDP per capita on the destination country-year level with 3G coverage in the origin. We find that preferred flows are more sensitive to destination-country PPP GDP per capita when being covered by 3G networks, although the effect is just statistically significant at a 5% level. In Columns 4 and 5, we include origin-year fixed effects to control for unobserved time-varying country-level factors. These show similar results for the interaction with prior stock of migrants, but insignificant results for the interaction with destination country GDP per capita. As prior stocks of migrants may be correlated to other factors affecting migration aspirations, we include interactions with the Polity IV indicator (on the destination-year level), a dummy for sharing a common language (on the origin-destination level) and the log of weighted distance between the origin-destination pair (on the origin-destination level) in Column (6). We find that the interaction effect of the prior stock of migrants remains comparable and highly significant, whereas the interaction effect of GDP per capita is just insignificant at a 5% level. Column (6) indicates that a one standard deviation larger log of dyadic migrant stock is associated with a 30% smaller effect of 3G on bilateral preferred flows after controlling for other factors. Altogether, this suggests that internet access not only affects the extent to which people want to emigrate, but also the destination to which Gallup respondents desire to migrate. As destinations change towards countries with lower stocks of previous migrants from the same country, the reduction of costs associated with finding information about prospective destinations and actual emigration likely mediate this effect. Furthermore, the positive effects for the interaction of 3G and destination country-level GDP per capita in Columns (3) and (6) suggest that preferred flows redirect to

[49] As many dyads have small stocks of migrants in 2005, small absolute increases in preferred bilateral migration rates (based on only a few Gallup respondents) of low-stock dyads may inflate the percentage increase of these dyads a lot. To alleviate concerns that this drives the results found, we omit dyads with less than 1,000 migrants in 2005 in an additional sensitivity analysis (results available upon request). The estimate of $3G_{ot} \times \ln(\text{Stock}_{od,2005}+1)$ remains similar and statistically significant.

more prosperous destinations, which is consistent with an information-channel supplying information about earnings opportunities abroad.

If actual migration patterns change in line with desire to emigrate, immigrants' birthplace diversity increases in receiving countries, which may boost innovation (Alesina, Harnoss and Rapoport, 2016). Furthermore, increased immigration could influence the politics of receiving societies. For example, there is convincing evidence that ethnic diversity reduces support for income redistribution (Dahlberg, Edmark and Lundqvist, 2012; Alesina, Miano and Stantcheva, 2023). Both sending and receiving countries could benefit from additional networks boosting international trade (Gould, 1994; Parsons and Vézina, 2018). Sending countries could also benefit from additional knowledge flows (Kerr, 2008; Fackler, Giesing and Laurentsyeva, 2020).

1.7.2 Well-being and Satisfaction with Institutions

To further explore possible mechanisms, we consider the direct effect of 3G rollout on outcomes that may affect migration behavior.[50] We use various indices as constructed by Gallup, supplemented with reported log household income, a constructed aggregate index of material prospects, the first principal component of trust in the government as constructed by Guriev, Melnikov and Zhuravskaya (2021), and information on banking and remittances.

Does mobile internet access affect perceived material well-being?

The first mechanism is related to perceived material well-being. In particular, we test whether respondents' perceived economic and financial conditions change after obtaining mobile internet access. To do so, we consider four outcome variables in Panel A of Table 1.6. The outcomes across the columns in the top panel are as follows: "(log) household income" in Column (1); "material prospects index" in Column (2); "job climate index" in Column (3); and "financial well-being index" in Column (4).

In Column (1), we find no statistically significant relationship between our treatment variable and per capita household income (an objective measure of material well-being).[51] This also addresses the caveat in Hypotheses 1 and 2 that the results hold if the mobile internet access does not boost local wages substantially. Not only do we find no substantial boost, but the point estimate on the effect on log per capita household income is negative, although statistically insignificant. The results reported in Columns (2) to (4) indicate that access to the mobile internet leads to a fall in the material prospects index and job climate index (measures the attitudes about a community's efforts to provide economic opportunities). We also find that mobile internet access has a negative effect on the financial well-being index (measures respondents' subjective evaluations on their personal economic situations and the economic situation of the community in which they live).

Overall, these results suggest that individuals' perceived material well-being declines after mobile internet penetration, while there is no effect on their household income. Such increased dissatisfaction could be a push factor to emigrate.

Does mobile internet access affect views about life?

In Panel B of Table 1.6 we explore the impact of mobile internet access on views about life. In particular, we present evidence using four outcome variables. The outcome variables across the columns in the middle panel are as follows: "optimism index (measures respondents' positive attitudes about the future)" in Column (1); "daily experience index (a measure of respondents' experienced well-being on the day before the survey)" in Column (2); "life evaluation index

[50] Apart from log household income, the outcomes presented in this section all strongly correlate to the desire to emigrate, before 3G coverage arrives. Appendix Table 1.E.22 shows the results of similar regressions of the desire to emigrate on these outcomes.

[51] As household income per capita is an imperfect measure of wages, we perform the following two robustness tests: (i) we find that the effect remains insignificant if we focus on single person households and (ii) we find that 3G has no effect on whether the respondent is in employment or not.

TABLE 1.6: The Effect of 3G Coverage on Material Well-being and Satisfaction with Life and Institutions

Panel A: Material well-being	(1)	(2)	(3)	(4)
Dependent variable:	Log of household income (PPP) per capita	Material prospects first principal component	Job climate index	Financial well-being index
3G	-0.026	-0.030**	-0.036**	-0.114*
	(0.035)	(0.014)	(0.018)	(0.067)
Observations	617,402	569,708	614,435	172,653
R^2	0.71	0.24	0.19	0.23

Panel B: Life satisfaction and optimism	(1)	(2)	(3)	(4)
Dependent variable:	Optimism index	Daily experience index	Life evaluation index	Life purpose index
3G	-0.018	-0.005	0.030	-0.046
	(0.014)	(0.007)	(0.021)	(0.073)
Observations	617,220	615,880	580,644	172,467
R^2	0.22	0.12	0.21	0.20

Panel C: Institutional satisfaction	(1)	(2)	(3)	(4)
Dependent variable:	Law and order index	Corruption index	Community basics index	Trust in government first principal component
3G	0.015	-0.017	0.010	-0.037**
	(0.009)	(0.014)	(0.010)	(0.015)
Observations	616,783	588,979	617,402	486,283
R^2	0.19	0.22	0.25	0.23

Panel D: Mobile banking and remittances	(1)	(2)	(3)	(4)
Dependent variable:	Owns a bank account	Used cellphone to receive cash in last 12 months	Received money or goods from friend/ family from same country	Received money or goods from friend/ family from another country
3G	-0.020	0.003	-0.009	0.004
	(0.038)	(0.026)	(0.015)	(0.008)
Observations	169,581	161,081	566,956	566,956
R^2	0.40	0.21	0.12	0.10

Notes: OLS regressions. Standard errors, clustered by district and country-year, in parentheses. The specification used in Columns (1) to (4) of Panels A to D is similar to that of Column (4) of Table 1.1. We only exclude the control variables related to local amenities as some of these amenities are used in the construction of the GWP indices. The number of observations is varying by item because of imperfect variable coverage over time and subnational regions. Except for column 1 of Panel A and Panel B (2013-2015) and Columns (1) and (2) of Panel D (2011, 2014 and 2017), all outcomes are covered between 2008 and 2018. In Columns (1) and (2) of Panel D we omit the district-level time trends, as we only have 3 time periods available (2011, 2014 and 2017). All dependent variables in Panels A – C are GWP indices, except for "(log) household income" (which is the reported log of per capita household income), "material prospects" (a first principal component of the following questions (weights in parentheses): living comfortably on present income (0.69), now is a good time to find a job (0.34), and not having enough money to afford food (-0.65)), and "trust in goverment" (a first principal component of four questions related to trust in the government, as constructed by Guriev, Melnikov and Zhuravskaya (2021)). For all items in Panels A to C a higher value of the dependent variable implies a higher value of the item. For example, a higher value of "Material prospects first principal component" implies a better subjective evaluation of material well-being and a higher value of "Corruption index" implies a larger perception of corruption. For construction of the GWP indices, see https://www.oecd.org/sdd/43017172.pdf (Last accessed on 08-12-2021). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(respondents' perceptions of where they stand now and in the future)" in Column (3); and "life purpose index (measures whether one likes what she does daily and is motivated to achieve one's goals)" in Column (4). We find no effect on any of these outcomes.

Does mobile internet access affect satisfaction with institutions?

To investigate whether a fall in satisfaction with institutions can also increase desire to emigrate, we regress various outcomes on mobile internet access, the results of which are reported in Panel C of Table 1.6. The outcome variables across the columns in Panel C are as follows: "law and order index" in Column (1); "corruption index" in Column (2); "community basics index" in Column (3); and "trust in government" in Column (4).

The results in Columns (1) to (3) are based on indices constructed by Gallup and show that there is no effect on the law and order index (gauges respondents' sense of personal security), corruption index (measures perceptions in a community about the level of corruption in business and government) and community basics index (measures everyday life in a community, including environment, housing and infrastructure). In Column (4), in line with Guriev, Melnikov and Zhuravskaya (2021), we find that 3G mobile internet has a negative effect on trust in government.

Does mobile internet access affect access to financial services and remittances?

To investigate whether an increase in access to financial services can explain the results found, we consider whether mobile internet access has an effect on the use of financial services[52] and on domestic and international remittances directly. As mobile banking has the potential to reduce the costs of remittances, this could increase the benefits to migration for those staying behind, possibly fostering emigration. We report the results in Panel D of Table 1.6. The outcome variables across the columns in the last panel are as follows: "Owns a bank account" in Column (1); "Used cellphone to receive cash in the last 12 months" in Column (2); "Received money or goods from friend/ family from same country" in Column (3); "Received money or goods from friend/ family from another country" in Column (4). We find no evidence for mobile internet access having affected any of these outcomes.

Back-of-the-envelope estimates

Using the estimates of Tables 1.6 and 1.E.22, we can calculate a back-of-the-envelope estimate on how big part of the effect of mobile internet access on desire to emigrate could be driven by the proposed mechanisms. Altogether, four out of 16 proposed mechanisms in Table 1.6 are statistically significant. By multiplying the statistically significant coefficients in Tables 1.6 and 1.E.22, we get an idea on the relative strengths of the suggested mechanisms. The effect of moving from 0% to 100% 3G coverage on the desire to emigrate is 0.0030 through material prospects first principal component, 0.0026 through the job climate index, -0.0035 through the financial well-being index (a negative sign suggests a counteracting force), and 0.0046 through the trust in government. As some of these mechanism may be partly multicollinear, adding the mechanisms together is likely to give an overestimate on their joint effect. With this caveat in mind, the estimated combined effect would be 0.0067 if summing these. Along similar lines, we can estimate how much of the effect can be explained by access to information, using the estimates of Table 1.4. If we assume that the difference in effect size between those with and without prior networks is fully driven by a difference in available information related to close networks, other channels are able to explain only 0.016 pp, and the channel concerning information related to close networks an additional 0.028 pp. As the additional channel is available for the 65% of population without prior networks, this contributes 0.018 pp to the total effect size.

In summary, our results suggest that access to the mobile internet led to a decrease in perceived material well-being and trust in government, which could explain at most one third of the estimated relationship between mobile internet access and the desire to emigrate in Column

[52]For outcomes related to banking we make use of the FINDEX module, which is an add-on to Gallup conducted in 2011, 2014, and 2017, on the same individuals as in the GWP. Because we have only 3 time periods per subnational district available, we drop the district-level time trends.

(4) of Table 1. They also suggest that the availability of information previously only available through close networks could explain more than half of the found effect of mobile internet access on the desire to emigrate.

1.8 Does Mobile Internet Also Affect Real Emigration Behavior? The Case of Spain.

As few countries have reliable subnational emigration registries, estimating the effect of 3G coverage expansion on actual emigration on a large scale is infeasible. However, Spain has such data. The Spanish Statistical Office (INE) maintains a population registry where inflows and outflows are recorded by person based on municipal registrations, including supplementary information such as country of origin.[53] Data is published for all municipalities with more than 10,000 inhabitants. These municipalities contain 76% of the population of Spain (in 2008). We focus on emigration rates of individuals born in Spain, as we expect internet access to affect them most. For individuals not born in Spain emigration might mean simply returning to their country of origin and may in some cases be the end of a stay that was already initially planned to be temporary.

Using the Mobile Coverage Explorer and the Gridded Population of the World (GPW) population density, we calculated the share of population covered by 3G in these municipalities. The first time nonzero coverage is reported to the Mobile Coverage Explorer was in December 2008.[54] As population-averaged coverage was already 80% in all municipalities with more than 10,000 inhabitants in 2008, recorded variation in 3G coverage is limited over time and concentrated among smaller municipalities. In December 2008, the 50 province capitals of Spain already had a population-averaged reported 3G coverage of 87%, whereas the municipalities that are not province capitals had an average coverage of 71%. Between 2003 and 2015, migration from all municipalities in the sample increased gradually. In 2003, only 0.03% of the population emigrated, whereas, in 2015, 0.11% of the population emigrated.

To assess the question of whether 3G expansion has effects on actual emigration of Spanish-born individuals from Spain, we estimate the following linear continuous difference in differences model:

$$m_{dt} = \beta_1 3G_{d(t-1)} + \beta_2 u_{pt} + \phi_d + \theta_t + \epsilon_{dt} \tag{1.9}$$

where m_{dt} is the emigration rate of Spanish-born individuals from municipality d in year t. We control for the unemployment rate u_{pt} at the provincial level. Our sample contains 657 municipalities in 50 provinces, of which 29 have a population exceeding 200,000 in 2008. Our resulting sample covers the years 2010 to 2019, as prior years have no information on the first lag of 3G coverage.[55]

As variation in 3G coverage is larger within smaller municipalities, we report the results for small and large municipalities separately. Table 1.7 reports the estimation results of Equation 1.9, for all, the small and the large municipalities in Columns (1) to (3), respectively. We find that a 10 percentage point increase in 3G coverage on the municipality level increases emigration by 0.0016 percentage points for all municipalities, significant at a 5% level. For the smaller municipalities, we find that a 10 percentage point increase in 3G coverage on the municipality level

[53] This registry is called *Diseño de registro de la Estadística de Variaciones Residenciales (EVR)*. Data can be found here: https://www.ine.es/dyngs/INEbase/es/operacion.htm?c=Estadistica_C&cid=1254736177013&menu=resultados&idp=1254734710990#!tabs-1254736195469

[54] 3G networks were present in Spain prior to 2008. See, for example, https://www.elmundo.es/navegante/2004/10/26/empresas/1098805246.html (accessed on 21-10-2021)

[55] As in the empirical strategy in section 1.5, we use coverage reported in December as representative for coverage in the next year. Therefore, the first available year is 2010.

increases emigration by 0.0014 percentage points (significant at the 5% level) and for the larger municipalities with 0.0018 percentage points (insignificant). For the smaller municipalities, the average yearly emigration rate is about 0.09%, implying an increase in migration of around 1.5% due to a 10 percentage point increase in mobile internet. These results suggests that the rollout of mobile internet not only led to increases in stated aspirations and intentions, but also to more actual migration.

TABLE 1.7: The Effect of 3G Coverage on Emigration of Spanish-born Individuals from Spain

	(1)	(2)	(3)
Dependent variable:		Emigration rate ($\times$ 100)	
Population in 2008:	All	$\leq$200,000	$>$ 200,000
3G Coverage$_{t-1}$	0.016**	0.014***	0.018
	(0.007)	(0.005)	(0.028)
Observations	6,570	6,280	290
R^2	0.873	0.838	0.951
Average emigration rate ($\times$ 100)	0.094	0.093	0.105
Municipality and year FE	✓	✓	✓
Provincial unemployment	✓	✓	✓

Notes: OLS regressions. Standard errors in parentheses. The dependent variable is the average emigration rate of Spanish-born individuals from a municipality between 2010 and 2019, multiplied by 100. The unit of observation is the municipality. We control for yearly averaged unemployment rates on the provincial level. Column (1) includes all municipalities, Column (2) includes all municipalities with a population of less or equal than 200,000 in 2008, Column (3) includes all municipalities with a population exceeding 200,000 in 2008. Standard errors are clustered two-way: on the municipality and the province-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.9 Conclusion

In this article, we show that mobile internet access increases both desire and plans to emigrate. Our analysis combines Gallup World Polls data from more than 600,000 respondents living in 2,120 sub-national districts in 112 countries, collected between 2008 and 2018, and geo-coded data on worldwide 3G mobile internet rollout. The effects are sizeable: a 10 percentage point increase in 3G mobile coverage leads to a 0.27 percentage point increase in the desire to emigrate permanently, and a 0.09 percentage point increase in plans to emigrate permanently over the ensuing 12 months. As the average increase in 3G coverage between 2008 and 2018 was 36% across the 2,120 subnational districts, our estimates suggest that such an increase goes together with a 0.98 percentage points increase (95% confidence interval: 0.21 to 1.82) in the desire to emigrate permanently. These are conservative estimates, identified by exploiting within-district variation in 3G coverage that has been stripped of any influence of constant and linearly changing district-level characteristics.

That the desire to emigrate increases more than the emigration plans is in line with the idea captured in our model that only a fraction of those desiring to emigrate is actually able to do so. By using an alternative estimator suitable for settings with a staggered treatment, we show that the results are not spuriously driven by the use of a two-way fixed effects estimator and, if anything, are underestimated. Moreover, using an instrumental variable based on the incidence of lightning strikes, we provide supplementary evidence that the effects are causal. Furthermore, the effects estimated using instrumental variables are considerably larger, suggesting that the estimated effects without using instruments are likely to underestimate the true effect. These effects are likely to translate into subsequent actual migration behavior, as migration aspirations and intentions are strongly correlated with actual migrant flows.

We find substantial heterogeneity in the effects of mobile internet access. The effect on desire to emigrate is strongest for secondary-education individuals and in the age group 25 to 44, as well as among those who are currently unemployed, involuntarily part-time employed, or outside the labor force, and among those who have relatively low per capita household income. Those with secondary education and those in the lowest income tercile are most responsive to mobile internet access in both countries with low or middle income and in countries with high income.

Using data on actual emigration from Spanish municipalities confirms that increased mobile internet access goes together with increased emigration. Our estimates suggest that switching from no to full 3G coverage in Spain increased annual emigration by about 15 percent compared with emigration rates that could have been expected in the absence of 3G coverage.

Our theoretical model suggests that mobile internet access is likely to increase both desire and plans to emigrate by reducing the cost of information acquisition. In line with this prediction, our analysis reveals that mobile internet access has strongest effects on respondents who do not have personal networks abroad, which can be explained by internet access substituting for personal contacts. Furthermore, we find that increased mobile internet access reduces perceived material well-being and also erodes trust in own government. Such increased dissatisfaction could be an additional channel through which mobile internet access increases desire and plans to emigrate. Finally, our results suggest that mobile internet access may not only increase overall international migration but also redirect migration flows towards less popular destinations. This could have far-reaching implications on both origin and destination countries.

Appendix

1.A Additional Information on Outcome Variables

Construction of Outcome Variables from GWP

In addition to the questions on desire and plans to emigrate outlined in the main text, the Gallup World Polls contains two other questions on (international) mobility:

4. Self-assessed Likelihood to Migrate (2008 – 2018): *In the next 12 months, are you likely or unlikely to move away from the city or area where you live in?* **Likely / Unlikely / Don't know / Refused to answer**[56]

5. International Migration Preparations (2009 – 2015): *Have you done any preparation for this move?* (asked only of those who are planning to move to another country in the next 12 months) **Yes / No / Don't know / Refused to answer**

The first is answered positively by approximately 17% of individuals, whereas the latter is answered positively by approximately 1.7% of respondents. Appendix Table 1.E.1 provides all relevant questions as they were stated in the GWP, and provides information on how we combined the variables if any modification was needed. The leftmost column contains the numbers of the outcomes reported in the main text.

Question (1) refers to the desire to emigrate. Question (2) refers to emigration plans and comprises two questions that are slightly different. Individuals who did not name a country in WP3120 are not asked WP6880 and are thus flagged as not planning to emigrate. However, it is unlikely that a respondent planning to emigrate is unable to identify the intended destination country in the preceding question. A greater issue is posed by individuals planning to emigrate to a feasible destination country, instead of their preferred destination. They might have identified another country in WP3120, which they only desire to emigrate to. These individuals then would answer negatively to WP6880, as they do not plan to emigrate to the country mentioned in WP3120. Therefore, for some individuals we might underestimate their plans to emigrate when considering WP6880. However, within-country positive rates of WP10252 (which asked about plans to emigrate within 12 months) and WP6880 (which asked about plans to emigrate to the preferred destination country within 12 months) are comparable, suggesting that emigration plans usually refer to plans to emigrate to the preferred destination country, and combining the two questions is justified. By combining WP10252 and WP6880, we are able to obtain a measure of plans to emigrate between 2008 and 2015. Having a longer sample is especially valuable as the positive rate of Question (2) is low and thus expected effect sizes are small.

Descriptives

Appendix Table 1.E.2 presents an overview of the main variables, including the data source and the level of observation. Averaging across all country-years, 22% of respondents report that they would like to move permanently to another country, while only 3% report that they are planning to move permanently in the next 12 months. 17% report being likely to move away from the city or area in which they live in the next 12 months. 46% of survey respondents are men. The average age of respondents is 40, 15 (53)% have completed tertiary (secondary) education and 58% have a partner.

[56] This question relates to movements both within and across international borders with no constraint imposed on the distance of the move.

1.A.1 Additional outcome variables

We visually summarize our four main outcomes in a Venn diagram in Figure 1.A.1, which identifies six mutually-exclusive regions for migration aspirations and plans (ranging from planning likely emigrants to likely stayers). Region I is of particular interest as it combines plans to emigrate with a self-assessed likelihood of moving away within 12 months. Therefore, it captures more developed plans to emigrate (I), in comparison with general plans to emigrate (I+II). Among those *planning* to emigrate (2.8%), about two-thirds (1.8%) report that they are also likely to move within 12 months. Moreover, region V identifies those deeming migration likely, but do not desire to emigrate. Although not a perfect measure, V predominantly captures those who intend to migrate domestically.

FIGURE 1.A.1: Venn Diagram of the Three Migration-related GWP Questions

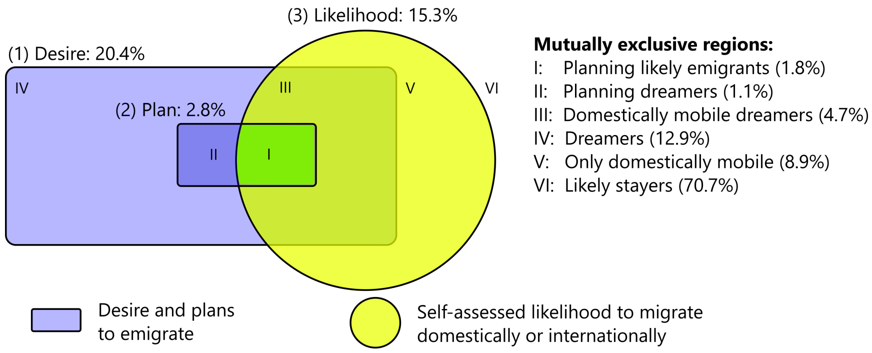

Notes: Venn diagram of the three migration-related outcomes (desire and plans to emigrate and self-assessed likelihood to migrate, including domestic migration) identifying six mutually exclusive regions. Note that the analysis is limited to the time period 2010 – 2015, as outside of this window not all underlying questions are asked in GWP. The Figure reports the unweighted proportion of respondents answering each question positively as two boxes (desire and plans to emigrate) and one circle (self-assessed likelihood to migrate). The share of respondents belonging to each of the mutually exclusive regions denoted in the Venn Diagram is reported on the right. The subsample where all three outcomes are available comprises $N = 342,328$ individuals.

1.A.2 Results

Table 1.A.1 reports estimates for four additional dependent variables. The dependent variables are a dummy indicating that the respondent “has any desire to emigrate or deems it likely to migrate in the upcoming 12 months” (first panel), that the respondent assess that he/she is "likely to migrate in the next 12 months" (second panel), that the respondent “likely plans to emigrate in the next 12 months” (third panel); and that the respondent "is only likely to migrate domestically in the next 12 months" (fourth and last panel). For each of the panels, we provide in Roman letters in brackets the mutually exclusive regions of Figure 1.A.1 the (constructed) outcome refers to. The first outcome captures a broad spectrum of migration aspirations: it measures whether respondents desire to emigrate or deem it likely to move away from their current residence. The second outcome captures the assessed likelihood of moving internally or internationally On the contrary, the third outcome captures a narrow intention to emigrate within the next 12 months, which is narrower than the outcome used in the second panel of Table 1.1 as it excludes those who plan to emigrate, but do not deem it likely they will actually do so in the next 12 months. The last solely captures domestic migration intentions in the next 12 months.

We find that 3G internet has a positive, sizable and statistically significant effect on all of these outcome variables with the exception of the category of likely internal migrants. This suggests that internal migration is not affected, which is in line with the hypothesis that internet access reduces migration costs more for lesser known (foreign) destinations than better known (domestic) destinations.

TABLE 1.A.1: The Effects of 3G Rollout on Alternative Outcome Variables

	(1)	(2)	(3)	(4)	(5)
Outcome:	Any desire or plans to migrate (I-V)				
3G	0.041***	0.040***	0.042***	0.041***	0.040***
	(0.014)	(0.014)	(0.013)	(0.013)	(0.013)
Observations	489,182	489,182	489,182	489,182	489,182
R^2	0.12	0.18	0.21	0.21	0.22
Average dependent variable	0.311	0.311	0.311	0.311	0.311
First year	2008	2008	2008	2008	2008
Last year	2018	2018	2018	2018	2018
Outcome:	Likely to migrate in the next 12 months (I+III+V)				
3G	0.027**	0.026**	0.027***	0.027***	0.027***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Observations	547,758	547,758	547,758	547,758	547,758
R^2	0.10	0.13	0.16	0.16	0.16
Average dependent variable	0.171	0.171	0.171	0.171	0.171
First year	2008	2008	2008	2008	2008
Last year	2018	2018	2018	2018	2018
Outcome:	Planning likely emigrant within 12 months (I)				
3G	0.010***	0.011***	0.011***	0.011***	0.011***
	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)
Observations	342,328	342,328	342,328	342,328	342,328
R^2	0.05	0.06	0.06	0.06	0.07
Average dependent variable	0.018	0.018	0.018	0.018	0.018
First year	2008	2008	2008	2008	2008
Last year	2015	2015	2015	2015	2015
Outcome:	Likely internal migrant within 12 months (V)				
3G	0.009	0.009	0.009	0.008	0.008
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Observations	489,182	489,182	489,182	489,182	489,182
R^2	0.04	0.05	0.06	0.06	0.06
Average dependent variable	0.092	0.092	0.092	0.092	0.092
First year	2008	2008	2008	2008	2008
Last year	2018	2018	2018	2018	2018
District and year fixed effects	✓	✓	✓	✓	✓
District-year trends	✓	✓	✓	✓	✓
Demographic controls		✓	✓	✓	✓
Life satisfaction-related controls			✓	✓	✓
Income controls			✓	✓	✓
Country-year-level controls				✓	✓
Country×income quintile fixed effects					✓
Country×education fixed effects					✓

Notes: OLS regressions. Standard errors in parentheses. For explanation of the controls and fixed effects, see notes to Table 1.1. The first measure is constructed using the union of positive answers to questions (1) and (3), the second measure using the union of positive answers to questions (2) and (3), and the third measure using the union of positive answers to question (3) and negative answers to question (1). For an overview of the potential intersections of questions (1), (2), and (3) and their interpretation, see Figure 1.A.1. The standard errors are clustered two-way on the country-year and district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

1.A.3 Power calculation for preparations to emigrate

The question on preparations to emigrate is only included in GWP between 2009 and 2015 and in less sub-national districts than the question on plans to emigrate, leaving limited variation in 3G network coverage. Furthermore, this question is on average answered positively by only 1.7% of respondents. If treatment effects are not very large compared to the share of respondents answering positively, the statistical power to uncover the effect may be low. Arnold et al. (2011) suggest to perform a simulation in order to find the power of a statistical test in settings with more complicated treatment assignment. We follow Arnold et al. (2011) and simulate the outcome variable (in this case the preparations to emigrate) to study the power of a statistical test on our estimate of the effect of 3G coverage on preparations to emigrate. Using the observed 3G coverage and covariates from the data, we simulated the data for preparations to emigrate as follows. We assume that a move from 0% to 100% 3G coverage could increase the probability to prepare to emigrate by 0.3 percentage points, an increase of around 20% compared to the baseline rate of 1.7%, giving the equation:

$$\Pr(Y_{idt} = 1) = 0.017 + 3G_{dt} \times 0.003$$

Based on this probability, we draw a binary outcome for all observations. Thereafter, we run our baseline specification (see Equation 1.5) using the generated binary outcome and observed 3G coverage and covariates on the sample of respondents who answered the GWP question on *preparations to emigrate within 12 months*. This sample contains 293,176 individuals. Using 500 realizations of the simulation described above, we reject the null hypothesis of no treatment effect at a 5% significance level only in 9% of the realizations. In comparison, the power using similar procedures (using the found effect sizes reported in Table 1.1 as the simulated treatment effect sizes) for desire to emigrate and plans to emigrate are 99% and 53%, respectively. This suggests that the sample size and/or the expected effect sizes for preparations to emigrate are unfortunately too small to study the effect of mobile internet coverage, in contrast to the other outcomes with better sample coverage and higher prevalence.

1.B Robustness Checks

In this section we report further analyses establishing the robustness of our findings.

Event Study Approach and Assessment of Pre-Trends

Our preferred alternative to TWFE is the de Chaisemartin-D'Haultfœuille estimator for continuous treatments. Another alternative is event study approach, focusing on districts experiencing an increase of at least 50 percentage points in 3G coverage between two subsequent years. An advantage of focusing on sharp increases is that instantaneous and dynamic effects can be distinguished, as there is no significant increase in 3G coverage before or after the event. However, in our case the event study approach comes at the cost of observations. In the event study, we focus on the 25% of the sample that lives in a district with a sharp expansion of 3G coverage. Nonetheless, an event study approach provides a valuable complementary perspective on the effects of mobile internet access, which is why we present also results using it.

In Figure 1.5, we find no significant pre-trend in the desire to emigrate using the de Chaisemartin-D'Haultfœuille estimator for continuous treatments. In the estimates in Figure 1.5, the control group also contains units with only small year-on-year increases in 3G coverage, whereas the event study contains only (treatment and control) districts that receive a sharp increase in 3G coverage between 2008 and 2018. A test on the presence of pre-trends using the de Chaisemartin-D'Haultfœuille estimator that includes also small treatments and a test using the event study approach are complementary for two reasons. First, treated districts receiving a large treatment may differ from those units that receive only a small first increase in treatment. Second, control districts never receiving a sharp increase in 3G coverage may be on a different trend in the desire to emigrate compared with units that receive a sharp increase.

In the event study, we focus on districts that experienced an increase of at least 50 percentage points in 3G coverage between two subsequent years, and analyze how the desire to emigrate develops with regard to this event, net of all baseline controls. As such an event design is subject to the same issues as TWFE estimators (see e.g. Sun and Abraham (2021)), we estimate the model using the de Chaisemartin-D'Haultfœuille estimator for binary treatments.

Appendix Figure 1.D.1 shows the results from the event study. We show the first three placebo, the instantaneous and the first four dynamic estimators. Earlier placebo estimators and later dynamic estimators are omitted as they are based on less than 1,000 observations each. These estimators are derived using the de Chaisemartin-D'Haultfœuille procedure used in Section 1.6.2. In contrast to the use case of the de Chaisemartin-D'Haultfœuille estimators in Section 1.6.2, here we have a binary treatment indicator (which is 1 in a subnational district in all periods after receiving an increase of at least 50 percentage points in 3G coverage between two subsequent years, and 0 otherwise). Therefore, we have only one initial treatment group (pooling those with no 3G coverage and some 3G coverage before treatment), which is the case discussed in the first paragraph of Online Appendix 1.C.

The black line shows the event study estimates of 3G expansion around a sharp increase and confirms that the treatment is sharp, showing little change in 3G coverage before and after the steep increase. Prior to the event, 3G coverage increases on average by 4 percentage points in the three periods before treatment, as districts could have nonzero treatment prior to a treatment corresponding to an increase of at least 50 percentage points. The 3G coverage rises, on average, by around 75 percentage points during the sharp increase in 3G coverage. After the first period, treatment only rises by 6 percentage points in the subsequent 4 periods. This justifies the use of an event study. The light blue line displays the event study estimates of the desire to emigrate. The instantaneous coefficient is around 0.02 and statistically significant. The first and second dynamic effects are just insignificant, but in magnitude comparable to the instantaneous estimator. In contrast, the third dynamic effect is almost double as large, whereas the fourth dynamic effect is close to 0 and insignificant. The fourth dynamic effect is less precisely estimated than the third as it is based on less observations (3,774 compared to 5,556) than the third dynamic effect. As the treatment is very sharp, we can interpret the dynamic estimators as actual dynamic effects, compared to combinations of instantaneous and dynamic results as in section 1.6.2. These results thus further suggest that the first three dynamic effects after receiving 3G coverage are positive as well, implying that the found effect is not just transitory in nature. None of the pre-event estimates are significantly different from 0. The p-value of a joint test for significance of any of the pre-event estimators is 0.22.

Controlling for Nighttime Light Density as an Alternative Measure of Regional Development

To alleviate concerns that 3G expansion and regional development coincide and that the coefficient on 3G coverage is biased because it captures regional development, we control for the mean of subnational-district-year level of per capita income in the household. However, as this is a self-reported measure of income and a mean of a relatively small number of observations at the subnational-district-year level, we show that using other measures of regional development does not alter the main result. More specifically, we use the nighttime light density as an alternative measure of regional development in Column (1) of Appendix Table 1.E.6 and the median, instead of the mean, of district-year personal income in Column (2). Our results remain similar.

Robustness to Excluding Potentially Bad Controls

One might worry that some of the individual characteristics (life satisfaction and local amenities) are themselves affected by the 3G rollout. Therefore, we omit sets of controls in Columns (3) to (5) of Appendix Table 1.E.6. Excluding life satisfaction and living standard-related controls in Column (3), satisfaction with amenities in Column (4) and whether someone can count on friends in Column (5) separately hardly alters the coefficient on 3G coverage.

Robustness to Including an Extensive Set of Additional Controls

As many questions in GWP are only covered for a part of the sample, we omitted some potentially relevant controls. However, adding controls for employment status in Column (6) of Appendix Table 1.E.6, financial support from home country or abroad in Column (7) of Table 1.E.6 and the aforementioned extra controls and various other controls (related to views about hard work, life satisfaction in five years, whether the current region is good for immigrants and whether the respondent has health problems) in Column (8) of Table 1.E.6 do barely change the estimated effect of 3G coverage.

Falsification Exercise: Using 2G Expansion as a Treatment

The main function of 2G technology is the transmission of information via voice signals while that of 3G technologies is internet browsing, data transfer, and downloading. At the same time, the expansion of cellular 2G and 3G networks is strongly correlated because of the technologies' shared infrastructure. This raises as a potential concern that the estimated effect of the 3G expansion could have arisen, at least partially, already because of improved communication allowed by the coinciding expansion of 2G networks. However, in Column (1) of Appendix Table 1.E.7, we find that 2G coverage has no statistically significant effect on the desire to emigrate, which is consistent with the idea that 3G affects the desire to emigrate through improved internet access and is not driven by an improved ability for mobile bilateral communication. In Column (2), we find that inclusion of 2G in the main specification does not alter the point estimate of 3G coverage.

Falsification Exercise: Using Leads as Treatments

By regressing the desire to emigrate on leads in 3G coverage, we can assess whether future increases in 3G coverage predict previous changes in desire to emigrate. If this is the case, the parallel trends assumption may be violated or treatment may be anticipated.

Appendix Table 1.E.7 shows that the instantaneous value of 3G coverage (Column 5) has an effect on the desire to emigrate while lags (Columns 3 and 4) and leads of 3G (Column 6 and 7) have no effect on the desire to emigrate.[57] The insignificance of the leads alleviates the concern that both 3G coverage and the desire to emigrate may be related to a (slowly moving) omitted variable and therefore display non-parallel trends. If the main result would be driven by different longer-run pre-trends for treated and untreated units, we would expect the leads to have a significant effect on the outcome. Therefore, the insignificance of the first lead of 3G coverage renders it implausible that non-parallel pre-trends in desire to emigrate are present.

Ruling Out Influential Observations

We rule out the importance of influential observations by showing the coefficients of our preferred specifications by omitting one year at a time. Appendix Table 1.E.8 shows that our coefficient estimates are quite stable even as a specific survey year is excluded from our main sample in each iteration.

We repeat a similar analysis in Appendix Table 1.E.9, in which we exclude one global region at a time in each estimation and again find that our estimates are not driven by a single global region. We classify OECD countries outside Europe and Latin America as one group.

Robustness to Excluding Top 10 Refugee-origin Countries and Countries with High or Low Desire to Emigrate

In order to alleviate concerns that the found results are driven by a few countries in distress, we omit the 10 countries of origin with the most refugees.[58] Additionally, we omit countries where a large ($\geq$40%) proportion of GWP respondents desires to emigrate and those where a small ($\leq$10%) proportion desires to emigrate. Appendix Table 1.E.10 reports the baseline results for these three omissions. The coefficient on 3G is robust to omission of these country groups.

Measurement Error in Mobile Coverage Data

As the data on mobile network coverage is based on reports of mobile network operators, it may be susceptible to various kinds of measurement error. First of all, reporting may be delayed. Second, coverage is not necessarily reported by all network operators, possibly underestimating the 3G coverage. As both of those sources of measurement error may be related to mobile network operator, industry structure, as well as country- or district-level characteristics, these may potentially bias the results we reported. To alleviate concerns about such measurement errors affecting our estimates, we omit groups of countries in Appendix Table 1.E.11 based on several criteria, which are:

- Districts that report report sharp decreases (defined as a drop of 10 percentage points or more) in 3G coverage. It is unlikely that coverage drops sharply within one year. This may be the caused by a reporting error.[59]

- Countries with large initially reported 3G coverage:

[57]Please note that using the n^{th} lag (lead) disregards the observations in the n earliest (last) years of the sample.

[58]We consider the 10 countries with the largest number of refugees under the UN High Commissioner for Refugees mandate in 2015. These include Syria, **Afghanistan**, Somalia, South Sudan, **Sudan**, **Democratic Republic of the Congo**, Central African Republic, Myanmar, Eritrea and **Colombia**. The countries in bold are part of out baseline sample. For the raw data, see: `https://www.unhcr.org/refugee-statistics/download/?url=738dpE`

[59]This happens in 109 districts in the baseline sample, most of which are located in Europe (31 in six countries) or in the former Soviet Union (36 in five countries). A striking example is Finland, where six districts reported decreases greater than 50 percentage points in 2016, to (more than) fully recover in 2017.

We omit countries for which the first year of nonzero 3G coverage is 2009 or later, and more than 20% of population is covered already in the first year.[60] In this case, we deem it plausible that, prior to that year, the country already had nonzero 3G coverage.[61]

- Countries with much lower 3G coverage than mobile broadband subscriptions in 2015:

  Countries that have at least four times as many mobile broadband subscriptions per capita than population-averaged 3G coverage in 2015. In this case, it is plausible that 3G coverage is under-reported.[62]

Excluding these country groups individually in Columns (1) to (3) of Appendix Table 1.E.11, and all of them simultaneously in Column (4), does not change our results qualitatively.

Balancing Test

3G expansion depends on the choices by network operators and authorities giving permissions to network expansion. If these choices are correlated with the characteristics of local population, our econometric analysis risks associating parts of the estimated effects of endogenous network expansion to control variables with which it is correlated. To address this concern, we ran a balancing test to check whether our treatment variable is correlated with respondents' observable demographic and socio-economic characteristics, with results shown in Appendix Table 1.E.12. In line with our identification assumption, none of the estimates is statistically significant at a 5% level. Furthermore, the p-value on the joint insignificance of all covariates equals 0.11.

Multiple Hypothesis Testing

We also conducted multiple hypothesis testing based on a randomization inference technique, as recently suggested by Young (2019). This helps to establish the robustness of our results, both for individual treatment coefficients in separate estimations and for the null hypothesis that our treatment does not have any effect across any of the outcome variables (i.e., treatment is irrelevant). The method builds on estimating the distribution of treatment effects by randomizing the treatment assignment under the null hypothesis that the treatment effect is 0 for all observations, and comparing the pool of randomized estimates to the estimates derived in the baseline specification. Using 500 iterations, the results presented in Appendix Table 1.E.13 show that our three findings in Column (4) of Table 1.1 remain robust. The null hypothesis of the Westfall-Young test for irrelevance of the 3G treatment in all three regressions is also rejected, with a p-value of 0.014.

Robustness to Alternative Levels of Clustering

In our main specification, we cluster the standard errors in two ways: at the district level (2120 groups) and at country-year level (791 groups). We show that our results are robust to using alternative assumptions about the variance-covariance matrix: the results remain significant when clustering at gender-education-country level (assuming that residuals move collectively within these units) as well as clustering at country-level (see Appendix Table 1.E.14).

Are the Results Driven by Non-comparable Samples?

Not all countries and districts are consistently included in GWP between 2008 and 2018, especially in earlier years in our sample. Thus, the results could conceivably be biased by heterogeneous, non-comparable samples. We therefore consider the baseline result on the sample

[60] We do not omit countries that show such increases before 2009, as it does not affect our sample period.

[61] This is the case in Armenia, Burkina Faso, Cameroon, Dominican Republic, Ecuador, Ghana, India, Kuwait, Malta, Mauritius, Montenegro, Qatar and Tunisia.

[62] We calculate country-level averages of population-weighted 3G coverage and we compare this to the number of mobile broadband subscriptions in 2015 as indicated by the International Telecommunication Union (ITU) `https://tcdata360.worldbank.org/indicators/h1e032144`. This is the case in the following countries: Belize, Bhutan, Colombia, Costa Rica, El Salvador, India, Kyrgyzstan, Mozambique, Namibia, Nepal, Nigeria, Oman, Senegal, Thailand, Trinidad and Tobago and Venezuela.

of countries and districts that are included in all years. The results reported in Appendix Table 1.E.15 confirm that our findings are robust across balanced samples.

Robustness to Using Population Weights and Using No Weights

We weight our observations in the baseline using the within-country weights based on the inverse probability of being included in the Gallup surveys. These weights are based on the demographic characteristics of the respondent and of the country of residence.[63]

We show that found results are robust to the choice of weights in Table 1.E.16. Column (1) reports the results for the unweighted baseline regression, whereas Column (2) reports Gallup weights only (our baseline). We find that the effect size is largest when using individual-level population weights (Column 3). Although the estimate using population-weighted observations provides truly global evidence, we have chosen as our baseline the more conservative Gallup weights only, due to a concern that a few large countries could drive the found effect when using population weights. That the qualitative effects are similar is an important robustness test, as the preferred population and Gallup weights vary significantly between countries and, to a lesser extent, between individuals.

Robustness to Alternative District-specific Trends

In our baseline regressions, we use district-specific time trends to alleviate concerns about spurious correlations between district-level 3G coverage and desire to emigrate driven by unobserved drifts on the district level. However, to show that our results do not critically depend on the inclusion of these linear time trends, we consider alternative specifications in Appendix Table 1.E.17. Omitting the time trend reduces the effect size found by around one standard deviation (Column 2), whereas adding a quadratic time trend does not alter the results by much (Column 3).

Robustness to Omission of Phone Interviews

The Gallup World Polls are conducted in-person, except when countries have a phone penetration exceeding 80%. As reaching this threshold may be correlated with mobile network roll-out, differences in answers between in-person and phone interviews could drive our effects. Therefore, we check whether our results are robust to the omission of phone interviews. Table 1.E.18 shows that omitting phone interviews or omitting countries with at least one phone interview completely does not alter our results.

Robustness across Sub-Periods

As the internet and its contents developed rapidly in the period of study, we are interested in knowing whether the effect of 3G coverage is driven by either early or late time periods. Table 1.E.19 shows that this is not the case: the effect size before 2014 and from 2014 onwards are statistically significant at a 5% and 10% level respectively and the magnitude of the point estimates are very close to our baseline estimate.

[63] GWP supplies a within-country weight variable based on unequal inverse probability of selection, calculated from (among others) national demographics, number of phone connections per household and number of household members. This allows the calculation of average statistics on the national level and to weight regressions accordingly. We refer to those weights as Gallup weights. Moreover, GWP aims to cover each country with at least 1,000 interviews per country-year. This implies that small countries are oversampled in GWP with regard to their populations. One can calculate population-adjusted country weights by using the Gallup weights w_i^{Gallup}, country-level population data obtained from the World Bank in 2015, N_c, and the total number of respondents between 2008 and 2018 in GWP per country, N_c^{Gallup}:

$$w_{ic}^{pop} = w_i^{Gallup} \cdot \frac{N_c}{N_c^{Gallup}} \tag{1.10}$$

We refer to w_{ic}^{pop} as the individual-level population weights.

1.C Implementation of the de Chaisemartin-D'Haultfœuille Estimator

In this section we discuss the use of the estimator proposed by de Chaisemartin et al. (2024) as an alternative to a TWFE regression using a continuous treatment variable. First, we introduce the estimator in the case of a binary staggered treatment. Thereafter, we discuss the adaptation of this estimator to the case where treatment is continuous and can change more than once over time. For a full discussion of this estimator and further extensions, we refer the reader to de Chaisemartin et al. (2024).

dCDH Estimator for a Binary Treatment

In the staggered adoption case with binary treatment, DiD_l is an estimator comprising a weighted average over groups g of $DiD_{g,l}$. This elementary building block is the difference (between first-treated units and a weighted average of suitable not-yet-treated units) in differences (over the length of l periods after being treated) of those units first treated at time F_g and being untreated prior to that. g indexes the unit receiving treatment, in our case a subnational district. As this estimator computes $DiD_{g,l}$ at group g level, all variables are aggregated on the group-time (indexed by group g and time t) level prior to estimation.[64] The weights on $DiD_{g,l}$ are proportional to the number of observations in group g.

As it uses only clean control units (meaning that they have never been treated yet at t), this estimator is robust to treatment effect heterogeneity and dynamic effects.[65] Although this estimator is robust to those, for identification of a causal effect we still have to rely on a common trends assumption, which can be assessed using the placebo estimators.[66]

One can modify the estimator to allow for the inclusion of relevant covariates.[67] Including covariates allows for a weaker common trends assumption: common trends of treatment and control groups only needs to hold after conditioning on covariates.

---

[64] Symbolically, we can write this as: $DiD_{g,l} = Y_{g,F_g+l} - Y_{g,F_g-1} - \sum_{g':D_{g',t=1}=0,F_{g'}>F_g+l} \frac{N_{g',F_g+l}}{N_{F_g+l}} \left(Y_{g',F_{g'}+l} - Y_{g',F_{g'}-1}\right)$. $Y_{g,t}$ is the (weighted) group-time level average outcome of the individual outcome $Y_{i,g,t}$. Groups $g^{'}$ are suitable control groups if they are untreated at period 1 ($D_{g',t=1}$) and remain untreated until at least l periods after g receives treatment for the first time ($F_{g'} > F_g + l$). N_{g',F_g+l} are the number of observations in the suitable treatment group g'. N_{F_g+l} is the sum of the number of suitable groups g' ($\sum_{g':D_{g',t=1}} N_{g',F_g+l}$), such that the weights on the outcome differences of the control groups sum to 1.

[65] Importantly, to calculate the DiD_l using all available groups, one needs a treatment variable that is balanced on the group level, as knowledge of a group's past treatment status is essential for determining if it is a clean control group and whether the unit switches into treatment for the first time. Although we do not observe every district every year in the GWP, we do observe the value of 3G coverage in the gaps of the GWP sample. We leverage this information, which is not used in a TWFE setting, to identify the exact timing of switching into treatment.

[66] The placebo estimators DiD_l^{pl} calculate the difference-in-differences between the treatment and control units between l periods before and 1 period before the treated unit is treated for the first time. To ensure that we calculate the placebo estimators on (a subset of) the same observations as we calculate dynamic effects, we restrict the sample for the lth placebo estimators to the groups that are used for calculation of the lth dynamic effect. If we would not restrict the placebo estimators accordingly, the earlier (larger l) placebo estimators would predominantly cover later treated (larger F_g) units that can not be used for the later (larger l) dynamic estimators, which rely of earlier treated (smaller F_g) units. This problem arises due to the finite panel length: we do not observe outcomes (i) many periods after treatment for groups treated late in the panel and many periods before treatment of groups treated early. Therefore, these estimators are important assessments of differential pre-trends between treatment and control units prior to the first treatment.

[67] Covariate adjustment of the elementary building blocks $DiD_{g,l}$ is performed in two steps. First, we run an OLS regression of the first differences in outcome on the first differences in covariates and time fixed effects on the sample of all never treated and treated groups prior to first treatment. Secondly, we residualize the l^{th} temporal difference in outcomes using the coefficients of the first step multiplied by the l^{th} temporal difference in covariates. The covariate-adjusted $DiD_{g,l}$ are then the differences between treatment and control in the difference over time relative to first treatment l unexplained by the covariates. Covariate adjustment has implications for the feasibility of the estimator as there may be fewer observations in the regression than there are covariates in the first step.

Extending to the Case of Non-Binary Treatments

In our main empirical strategy, we use the population-averaged 3G coverage for every sub-national district. This is a non-binary treatment that gradually increases over time.[68] Nevertheless, we can still apply the principle of units switching into treatment for the first time to identify difference-in-differences between treatment and clean control groups. Some units receiving treatment for the first time during our sample period may already have a stable level of nonzero treatment for several periods (in our case, at least since the beginning of the sample period in 2008). We refer to the initial level of treatment in 2008 as *ini*. The elementary building block $DiD^{ini}_{g,l}$ is now differentiated over initial treatment status *ini* and we calculate the $DiD^{ini}_{g,l}$ using treatment and control groups with the same *ini*. As 3G coverage is continuous, it is necessary to bin the initial treatments *ini*, as otherwise all districts are in different groups and we are unable to find a control group for a group that switches to a higher treatment.[69]

If those bins become too wide, treatment and control groups with fairly different initial levels of treatment are compared. In order to estimate the DiD^{ini}_{l} unbiasedly, we have to assume that the treatment effects between the binned treatments do not vary over time.[70] As the (adoption of) internet and the activity of users changed considerably between 2008 and 2018, it is likely that treatment effects are heterogeneous over time. Any binning of initial treatment groups thus requires justification. As 3G coverage for many groups increases at least somewhat in most years between 2008 and 2018, it is helpful to define a stable treatment as an increase exceeding some threshold (which we call Δ_{3G}) for an increase in 3G coverage between two subsequent years. Without this adjustment, for some initial treatment levels *ini*, it is impossible to find control groups, as most of the sub-national districts have changing levels of 3G coverage during the time span studied. As such a threshold biases the control group somewhat towards the treatment group, this is a conservative adjustment. However, if Δ_{3G} is too large, some levels of *ini* may not have a single group switching into treatment and DiD^{ini}_{l} is not defined.

Figure 1.C.1 diagrammatically presents examples of time series of 3G coverage in the case of two initial treatment levels $ini = 0$ and $ini > 0$ and to which group they belong. Units that are never treated are indicated with C and units that are treated are indicated by T. Treated units have a subscript τ to indicate the time period in which they switch into treatment. For $t < \tau$ ever treated units are not yet treated and thus also valid control groups. An illustration of an increase smaller than the threshold Δ_{3G} is given in the time series for $C^{ini>0}$, as well as increases larger than the threshold in the time series for $T^{ini=0}$ and $T^{ini>0}$.

As with the binary staggered adoption design, we calculate the dynamic effects DiD_l where $l > 0$ are the cumulative effects of receiving treatment l periods ago. The interpretation of the

[68] It is important to note that our treatment 3G is not exactly monotonically increasing, as the level of 3G coverage is allowed to decrease between two periods. In only 234 out of 2,120 districts in the main sample the coverage decreases by more than 3% of population between any two years. As we omit these districts from our econometric analysis due to concerns about data reliability, we do not discuss here designs in which treatment can decrease. For a discussion of estimators when such declines are prevalent, see de Chaisemartin et al. (2024).

[69] Except for those districts with $ini = 0$, which constitute approximately 40% of our sample.

[70] If this is not the case, the counterfactual of remaining in treatment *ini* is not exactly the counterfactual treatment of staying in a slightly different initial treatment status $ini' \neq ini$ and the elementary building block $DiD^{ini}_{g,l}$ is biased through the differences in outcomes for control units (in symbols for all l: $Y^{ini}_{g,F_g+l} - Y^{ini}_{g,F_g-1} = Y^{ini'}_{g,F_g+l} - Y^{ini'}_{g,F_g-1}$ only holds if $TE^{ini\to ini'}_{g,F_g+l} = Y^{ini}_{g,F_g+l} - Y^{ini'}_{g,F_g+l} = Y^{ini}_{g,F_g-1} - Y^{ini'}_{g,F_g-1} = TE^{ini\to ini'}_{g,F_g-1}$). This bias is greater for (1) larger l, as treatment effects likely vary slowly as well as for (2) larger bins (implying larger $|ini - ini'|$), such that the treatment effect $TE^{ini\to ini'}$ is larger. This issue is mitigated if there is a balance in the various binned levels and their period of first treatment, as the biases cancel each other. In the simplified case of binning observations into two distinct treatment levels *ini* and *ini*′, we use both groups with *ini* as control groups for first switches from *ini*′ as well as groups with *ini*′ as control units for first switches from *ini*. In this case, the two contributions counteract, and the estimator binning observations with *ini* and *ini*′ together ($DiD^{ini\cup ini'}_{l}$) is less biased.

FIGURE 1.C.1: Examples of Relevant Treatment and Control Groups for the de Chaisemartin and D'Haultfœuille Estimator

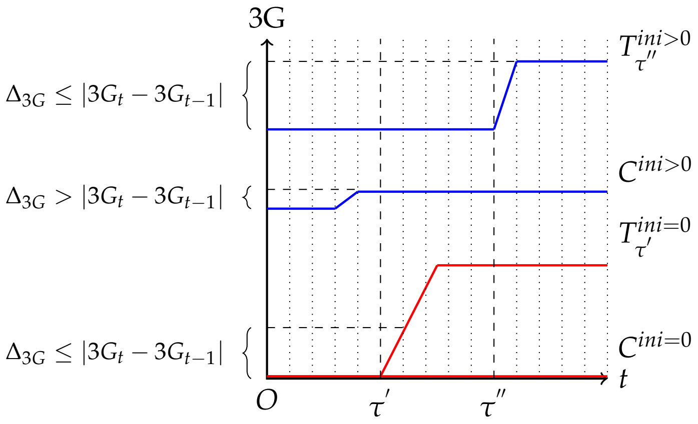

Notes: This figure shows an example of a treatment (T) and control (C) unit for both $ini = 0$ (in red) and $ini > 0$ (in blue). Note that $C^{ini=0}$ overlaps with the horizontal axis and $T^{ini=0}$ overlaps with horizontal axis until τ'. The control unit for $ini > 0$ is treated with less than the threshold Δ_{3G}, so we consider it as a control unit also after the marginal treatment. The treated group for $ini = 0$ receives a treatment exceeding the threshold between τ' and $\tau' + 1$ and is considered a treated group from $\tau' + 1$ onwards. The treated group for $ini > 0$ receives treatment exceeding the threshold between τ'' and $\tau'' + 1$ and is considered a treated group from $\tau'' + 1$ onwards.

DiD_l for the case of a monotonically increasing non-binary treatment is different from the binary staggered case. In the staggered case when $l \geq 1$, one can interpret DiD_l as the cumulative effect of being treated for l periods. However, as treatment may have increased further since the first time the district receives treatment (the 'first switch'), DiD_l is a weighted average of the instantaneous effect of increased coverage in period l and the dynamic effects of the first switch and the earlier period increases, respectively. Using the DiD_l, we can calculate the following quantity (de Chaisemartin et al., 2024):

$$\hat{\delta}^L = \frac{\sum_{l=0}^{L} w_l DiD_l^Y}{\sum_{l=0}^{L} w_l DiD_l^{3G}} \tag{1.11}$$

$\hat{\delta}^L$ is the treatment effect per unit of treatment which is calculated using the ratio of the DiD_l on the outcome of interest Y and the DiD_l on the treatment (3G), weighted by the share of observations in the lth effect. de Chaisemartin et al. (2024) shows that this is equivalent in interpretation to an IV estimator as the numerator in Equation 1.11 is the average treatment effect of a first switch, whereas the denominator is the average treatment following a first switch. Only if there would be no dynamic effects and treatment would be staggered, $\hat{\delta}^L$ denotes the ATT. Nevertheless, the estimator allows us to study the average treatment effect using an estimator robust to heterogeneous and dynamic treatment effects. Therefore, $\hat{\delta}^L$ identifies a convex combination of (heterogeneous) instantaneous and dynamic treatment effects.

1.D Additional Figures

FIGURE 1.D.1: Event Study around Sharp Increases in 3G Coverage

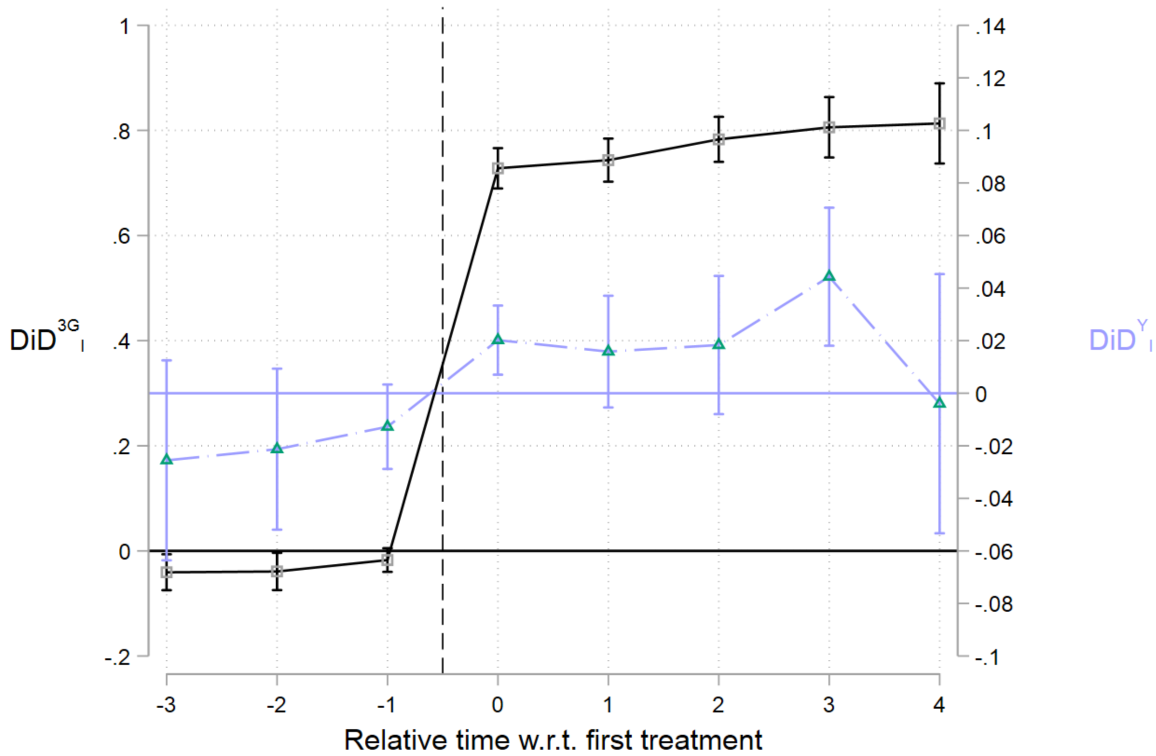

Notes: Event study estimates around treatment of 50 percentage point increase in 3G in a year with a 95% confidence interval using the estimator proposed by de Chaisemartin et al. (2024). The black (blue) line depicts the event study estimates with 3G coverage (desire to emigrate) as dependent variable. All units that experience a decrease of more than 10 percentage points in 3G coverage between any two subsequent years are omitted, to exclude districts with possibly poor data quality. The sample covers 116,413 respondents in 380 districts with a sharp increase in 3G coverage. A test of joint insignificance of the pre-treatment period (placebo) estimators for the desire to emigrate gives a p-value of 0.22.

FIGURE 1.D.2: Oster's δ for Increasing Values of Maximally Admissible R^2_{max}

.

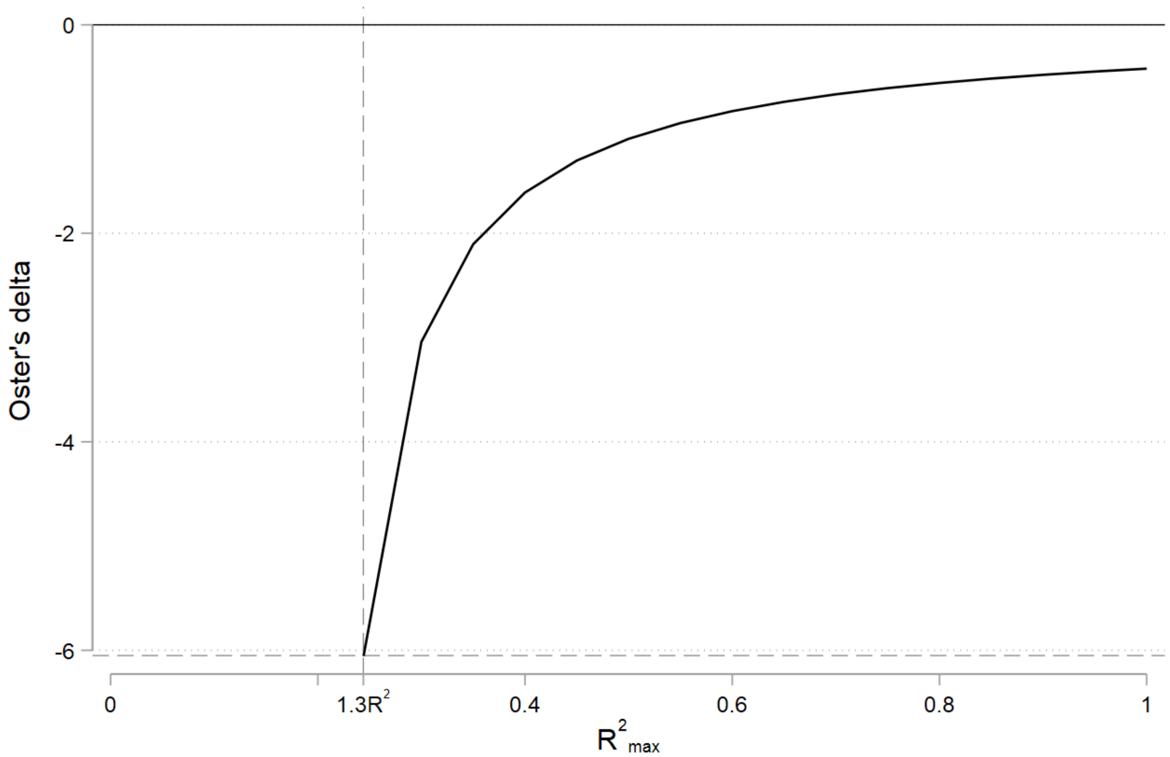

Notes: This Figure shows Oster's Delta (Oster, 2019) for different values of the maximum allowed variation in outcome that covariates can explain. Oster's Delta indicates how much stronger (and with what sign) the selection on unobservables should be compared to selection on observables to fully explain the found effect. The analysis here is based on our main specification, as found in Column (4) of Table 1.1. Oster's δ is equal to -6.05 for the value recommended by Oster (2019) of $R^2_{max} = 1.3R^2$ and Oster's δ is still equal to -0.4 when we allow unobservables to explain all remaining variation.

1.E Additional Tables

TABLE 1.E.1: Questions in GWP relating to Respondents' Aspirations and Intentions to Migrate

Variable	GWP ID	Question / construction	Coverage
Panel A			
(1): Desire to emigrate	WP1325	Ideally, if you had the opportunity, would you like to move permanently to another country, or would you prefer to continue living in this country?	(2008 – 2018)
(1*C*)	WP3120	To which country would you like to move? (Asked only of those who would like to move to another country (WP1325))	(2008 – 2018)
Panel B			
*Mig*10252	WP10252	Are you planning to move permanently to another country in the next 12 months, or not? (Asked only of those who would like to move to another country - WP1325)	(2010 – 2015)
*Mig*6880	WP6880	Are you planning to move permanently to that country in the next 12 months, or not? (Asked only of those who specified a country to which they would like to move. - WP3120)	(Mostly 2008/09)
(2): Plan to emigrate	WP10252& WP6880	*Mig*10252, *Mig*6880 if *Mig*10252 unavailable	(2008 – 2015)
(2*C*)	WP3120& WP10253	WP3120 if question (2) answered positively (2008 – 2009) and WP10253 (2010 – 2015)	(2008 – 2018)
Panel C			
(3): Preparation to emigrate	WP9455	Have you done any preparation for this move (asked only of those who are planning to move to another country in the next 12 months)	(2009 – 2015)
(3*C*)	WP10253	WP10253 if *Preparation to emigrate* answered positively	(2009 – 2015)
Panel D			
(4): Likely to move	WP85	In the next 12 months, are you likely or unlikely to move away from the city or area where you live?	(2008 – 2018)

TABLE 1.E.2: Summary Statistics and the Data Sources

Panel A: Baseline	Mean	S.D.	Observations	Source	Level
Desire to emigrate	0.22	0.42	617,402	GWP	Individual
Plan to emigrate	0.03	0.16	376,801	GWP	individual
Likely to move	0.17	0.37	544,022	GWP	Individual
District-level 3G coverage	0.37	0.39	617,402	Collins Bartholomew	District-Year
District-level 2G coverage	0.77	0.30	617,402	Collins Bartholomew	District-Year
Male	0.46	0.50	617,402	GWP	Individual
Age	40.10	17.02	617,402	GWP	Individual
Urban	0.39	0.49	617,402	GWP	Individual
With partner	0.58	0.49	617,402	GWP	Individual
Separated/divorced	0.06	0.24	617,402	GWP	Individual
Presence of children	0.56	0.50	617,402	GWP	Individual
Secondary education	0.53	0.50	617,402	GWP	Individual
Tertiary education	0.15	0.36	617,402	GWP	Individual
Born in country of interview	0.96	0.19	617,402	GWP	Individual
Log of HH per capita income	7.74	1.51	617,402	GWP	Individual
Log of district per capita income	8.15	1.15	617,402	GWP	District-Year
Life satisfaction	0.46	0.50	617,402	GWP	Individual
Can count on friends/relatives	0.82	0.39	617,402	GWP	Individual
Satisfied with living standard	0.62	0.48	617,402	GWP	Individual
Living standard is getting better	0.46	0.50	617,402	GWP	Individual
Lack of money for food	0.35	0.48	617,402	GWP	Individual
Lack of money for shelter	0.25	0.43	617,402	GWP	Individual
Satisfied with the city	0.78	0.41	617,402	GWP	Individual
Satisfied with public transport	0.62	0.49	617,402	GWP	Individual
Satisfied with roads	0.55	0.50	617,402	GWP	Individual
Satisfied with education	0.68	0.47	617,402	GWP	Individual
Satisfied with healthcare	0.58	0.49	617,402	GWP	Individual
Satisfied with housing	0.52	0.50	617,402	GWP	Individual
Had money or property stolen	0.16	0.37	617,402	GWP	Individual
Log of GDP per capita	8.44	1.40	617,402	World Bank	Country-Year
Polity 2	5.44	5.01	617,402	Center for Systemic Peace	Country-Year
Share of respondents below 30	0.32	0.13	617,402	GWP	Country-Year

Notes: All individual-level variables are binary apart from log of HH per capita income, with 1 denoting yes and 0 denoting no. All income-related variables are measures in USD PPP terms.

TABLE 1.E.3: The Effects of 3G Coverage on Access to the Internet

	(1)	(2)
Outcome:	Internet Access	
3G	0.049***	0.051***
	(0.015)	(0.014)
Broadband subscription rate		✓
Observations	614,945	606,541
R^2	0.52	0.52
Average dependent variable	0.435	0.435

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on the specification. Standard errors are clustered two-way: at the district and country-year level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.4: The Effects of 3G Coverage on Access to the Internet *at Home*

	(1)	(2)
Outcome:	Internet Access	
3G	0.012	0.014
	(0.016)	(0.015)
Broadband subscription rate		✓
Observations	409,966	405,351
R^2	0.53	0.53
Average dependent variable	0.363	0.363

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. Sample is restricted to 2008-2015, as the question on internet access explicitly referred to access at home up to and including 2015. See notes to Table 1.1 for details on the specification. Standard errors are clustered two-way: at the district and country-year level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.5: The Effects of 3G Expansion and Internet Access to the Internet *at Home* on the Desire to Emigrate

	(1)	(2)	(3)	(4)
Outcome:		Desire to Emigrate		
3G	0.033**		0.033**	0.039***
	(0.014)		(0.014)	(0.014)
Internet access at home		0.037***	0.037***	0.042***
		(0.003)	(0.003)	(0.004)
3G $\times$ Internet access at home				-0.013
				(0.008)
Observations	409,960	409,960	409,960	409,960
R^2	0.18	0.18	0.18	0.18
Average dependent variable	0.224	0.224	0.224	0.224

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. Sample is restricted to 2008-2015, as the question on internet access explicitly referred to access at home up to and including 2015. Standard errors are clustered two-way: at the district and country-year level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.6: Robustness to Inclusion of an Extensive Set of Additional Controls and Omission of Selected Baseline Controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome:	Desire to emigrate							
3G	0.029***	0.027**	0.028***	0.024**	0.027**	0.029***	0.030***	0.026**
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.012)	(0.013)
Nighttime light density	-0.000							
	(0.001)							
Log of district-year median per capita HH income		0.003						
		(0.005)						
Log of district-year mean per capita HH income			0.003	0.002	0.003	0.004	0.004	0.004
			(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Demographic and country-level controls	✓	✓	✓	✓	✓	✓	✓	✓
District-level trend and district and year FEs	✓	✓	✓	✓	✓	✓	✓	✓
Can count on friends/relatives	✓	✓	✓	✓		✓	✓	✓
Satisfaction with local amenities	✓	✓	✓		✓	✓	✓	✓
Satisfaction with life situation	✓	✓		✓	✓	✓	✓	✓
Employment status						✓		✓
Received money/goods (from home country and abroad)							✓	✓
Additional controls								✓
Observations	606,712	617,402	617,402	617,402	617,402	579,507	566,873	471,622
R^2	0.19	0.19	0.18	0.17	0.19	0.19	0.19	0.20

Notes: OLS regressions. Standard errors in parentheses. Columns (1) and (2) include the baseline controls, except for the log of average per capita income in the household on the district-year level. Column (1) includes the nighttime light density, whereas Column (2) includes the log of median per capita income in the household on the district-year level. Columns (3), (4) and (5) include the baseline controls, except for life satisfaction, satisfaction with living standards, whether the respondent believes to be financially better off in five years, whether the respondent has sufficient means for food, for shelter, and whether the respondent had something stolen in the past year in Column (3), satisfaction with housing, healthcare, education, roads, transportation and the city in Column (4), and whether the respondent can count on family or friends in Column (5). Columns (6), (7) and (8) include the baseline controls and additionally include a dummy for unemployment, involuntarily part-time employment and being out of the workforce in Column (6), whether the respondent received money or goods from abroad and whether the respondent received money or goods domestically in Column (7), and whether the respondent believes people can get ahead in life by working hard, expect to have higher life satisfaction in five years, whether the respondent believes his or her current living area to be good for immigrants, and whether the respondent has health problems in Column (8). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.7: Effect of 2G Coverage and Lags/Leads of 3G Coverage on the Desire to Emigrate

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Outcome:				Desire to emigrate			
$2G_t$	0.019	0.018					
	(0.014)	(0.013)					
$3G_t$		0.027**			0.027**		
		(0.011)			(0.011)		
$3G_{t-2}$			0.017				
			(0.012)				
$3G_{t-1}$				0.001			
				(0.012)			
$3G_{t+1}$						0.010	
						(0.013)	
$3G_{t+2}$							-0.001
							(0.015)
Observations	617,402	617,402	551,021	581,401	617,402	548,152	473,783
R^2	0.19	0.19	0.19	0.19	0.19	0.19	0.18
Average dependent variable	0.214	0.214	0.214	0.214	0.214	0.214	0.216

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1, apart from the reported coefficients. See notes to Table 1.1 for details on the specification. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 1.E.8: Robustness to Omission of Single Years from Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Outcome:	Desire to emigrate										
Omitted year:	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018
3G	0.028**	0.024**	0.018	0.040***	0.025**	0.030***	0.031***	0.031***	0.028**	0.031**	0.016
	(0.011)	(0.012)	(0.012)	(0.011)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.011)
Observations	590,636	586,273	565,156	551,182	558,148	565,846	562,549	547,900	546,699	541,586	558,045
R^2	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19
Average dependent variable	0.222	0.221	0.222	0.221	0.223	0.221	0.220	0.221	0.218	0.218	0.215

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 1.E.9: Robustness to Omission of Global Regions from Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Outcome:	Desire to emigrate						
Global region omitted:	Europe	Former USSR	AUS+CAN+ISR+JPN+ KOR+NZL+TUR+USA	Middle East	Rest of Asia	Americas without CAN+USA	Africa
3G	0.022*	0.027**	0.024**	0.032***	0.036***	0.022**	0.027**
	(0.011)	(0.012)	(0.012)	(0.012)	(0.013)	(0.011)	(0.012)
Observations	509,276	539,645	575,734	523,460	498,829	609,016	448,452
R^2	0.19	0.19	0.19	0.19	0.19	0.19	0.17
Average dependent variable	0.229	0.224	0.225	0.211	0.245	0.220	0.183

Notes: OLS regressions. Standard errors in parentheses. Results omitting a mutually exclusive global region at a time. Column (1) omits European countries (including the Baltic countries), Column (2) omits former USSR countries (excluding Baltic countries), Column (3) omits a group of developed non-European countries: Australia, Canada, Israel, Japan, New Zealand, South Korea, Turkey, and the United States, Column (4) omits the Middle East, Column (5) omits the remaining Asian countries (except the Middle East, former USSR, Israel, Japan, South Korea and Turkey), Column (6) omits the Americas (excluding USA and Canada), and Column (7) omits Africa. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 1.E.10: Robustness to Excluding Countries with Many Refugees and High or Low Share of Respondents Desiring to Emigrate

	(1)	(2)	(3)
Outcome:	Desire to emigrate		
Excluding countries:	Top 10 refugee	≥40% desire to emigrate	≤10% desire to emigrate
3G	0.026**	0.023**	0.036***
	(0.011)	(0.011)	(0.013)
Observations	599,017	565,042	515,940
R^2	0.19	0.16	0.17
Average dependent variable	0.216	0.194	0.251

Notes: OLS regressions. The specification follows the specification used in Column (4) of Table 1.1. Standard errors in parentheses. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. Column (1) omits respondents in Afghanistan, Sudan, Democratic Republic Congo and Venezuela. Column (2) omits countries where, on average, more than 40% of GWP respondents desires to migrate. Column (3) omits countries where, on average, less than 10% of respondents desire to migrate. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 1.E.11: Robustness to Dropping Observations with Potentially Poor-quality 3G Coverage Data

	(1)	(2)	(3)	(4)
Outcome:	Desire to emigrate			
Omits:	Districts with a more than 10 p.p. drop in 3G coverage between 2008 and 2018	Countries where first-reported 3G coverage exceeds 20%	Countries where 3G coverage is less than one-quarter of the number of mobile broadband subscriptions in 2015	All aforementioned
3G	0.031**	0.031**	0.032***	0.037**
	(0.012)	(0.012)	(0.012)	(0.014)
Observations	580,253	522,958	501,979	427,062
R^2	0.19	0.18	0.18	0.18
Average dependent variable	0.224	0.221	0.231	0.219

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column 4 of Table 1.1. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. Column (1) omits districts that experience a sharp drop of more than 10 percentage points in 3G coverage anytime between 2008 and 2018, Column (2) omits districts in countries that report a country-average population coverage exceeding 20% in the first year of nonzero reported coverage, Column (3) omits districts with a population-averaged 3G coverage lower than one-quarter of the number of mobile broadband subscriptions in 2015, as reported by ITU. Column (4) omits all units omitted in Columns (1-3) compared to the baseline displayed in Column (4) of Table 1.1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 1.E.12: Balancing Test of 3G Coverage on Baseline Demographic Covariates

Outcome:	3G × 100
Male	0.008
	(0.032)
Age	-0.001
	(0.006)
Age-squared	0.000
	(0.000)
Urban	0.028
	(0.147)
With partner	-0.102*
	(0.053)
Separated/divorced/widowed	-0.170*
	(0.099)
Presence of children	0.100
	(0.064)
Secondary education	-0.032
	(0.087)
Tertiary education	-0.101
	(0.121)
Not born in country of interview	-0.015
	(0.142)
Log of personal income	-0.009
	(0.050)
Log of district-year mean per capita HH income	-0.063
	(0.555)
N	617,402
R2	0.932

Notes: OLS regression. p-value of the F-test of joint insignificance: 0.1154. Standard errors are clustered two-way: on the district and the country-year level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.13: Robustness to Randomization Inference and Multiple Hypothesis Testing

Outcome:	(1) Desire to emigrate	(2) Plans to emigrate	(3) Likelihood to migrate	(4) Joint test of irrelevance
3G	0.027^{**}	0.009^{**}	0.027^{***}	
Young(2019) Randomized p-value	(0.020)	(0.023)	(0.004)	(0.014)

Notes: OLS regressions. Young (2019) randomization inference p-values in parentheses, based on 500 bootstrap replications. Column (1) to (3) denote the point estimates of Table 1.1 and the standard errors corrected for multiple hypothesis testing. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.14: Robustness to Alternative Variance-Covariance Matrix Structure

	(1)	(2)
Outcome:	Desire to emigrate	
3G	0.027***	0.027*
	(0.009)	(0.014)
Observations	617,402	617,402
R^2	0.19	0.19
Level of clustering	Country-Education-Gender	Country
Number of clusters	658	112

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.15: Robustness to Omission of Non-balanced Countries and Districts

	(1)	(2)
Outcome:	Desire to emigrate	
3G	0.047***	0.059***
	(0.018)	(0.020)
Observations	202,378	179,138
R^2	0.16	0.15
Average dependent variable	0.156	0.164
Level of balancing	Country	District

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.19: Interaction of 3G Coverage with Time Period Dummy

Outcome:	Desire to emigrate
3G × I(Year<2014)	0.029**
	(0.012)
3G × I(Year≥2014)	0.024*
	(0.012)
Observations	617,402
R2	0.19

Notes: OLS regression. Standard errors are clustered two-way: on the district and the country-year level. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.16: Robustness to Alternative Choices of Weighting Observations

	(1)	(2)	(3)
Outcome:		Desire to emigrate	
3G	0.033***	0.027**	0.039***
	(0.010)	(0.011)	(0.013)
Observations	617,402	617,402	617,402
R^2	0.19	0.19	0.22
Average dependent variable	0.222	0.222	0.222
Weights	Unweighted	Gallup only (baseline)	Population and Gallup

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.17: Robustness to Different Specifications of District-specific Time Trends

	(1)	(2)	(3)
Outcome:		Desire to emigrate	
3G	0.027**	0.018*	0.032***
	(0.011)	(0.009)	(0.012)
Observations	617,402	617,402	617,402
R^2	0.19	0.18	0.20
Average dependent variable	0.222	0.222	0.222
District-level trend	Linear	-	Linear + Quadratic

Notes: OLS regressions. Standard errors in parentheses. Column (1) presents the baseline result. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.18: Robustness to Omission of Telephone Interviews

	(1)	(2)	(3)
Outcome:		Desire to emigrate	
3G	0.027**	0.029**	0.028**
	(0.011)	(0.012)	(0.013)
Observations	617,402	514,637	506,326
R^2	0.19	0.19	0.19
Average dependent variable	0.231	0.231	0.231
Telephone Interviews	All	No	No (country)

Notes: OLS regressions. Standard errors in parentheses. The specification follows the specification used in Column (4) of Table 1.1. See notes to Table 1.1 for details on control variables. Column (2) omits all phone interviews, whereas Column (3) omits all countries with at least 1 phone interview in the sample. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.20: The Effect of the Lightning-based Instrument on Desire to Emigrate, prior to 3G rollout

	(1) Desire to Emigrate
Below-median GDP per capita countries $\times \log(\mathcal{L}_d) \times$ year	-0.001
	(0.002)
Above-median GDP per capita countries $\times \log(\mathcal{L}_d) \times$ year	0.003
	(0.002)
Observations	4,151
R-squared	0.769
Mean dep. var	0.239

Notes: OLS regression. Standard errors in parentheses. See notes to Table 1.1 for details on control variables. Standard errors are clustered two-way: on the district and the country-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.21: Heterogeneity of the Effect of 3G Coverage on Desire to Emigrate Based on Individual- and Country Level Characteristics

Panel A: Educational Attainment and Country Income Groups

	(1) Primary Educated in Low- or Middle-income Countries	(2) Secondary Educated in Low- or Middle-income Countries	(3) Tertiary Educated in Low- or Middle-income Countries	(4) Primary Educated in High-income Countries	(5) Secondary Educated in High-income Countries	(6) Tertiary Educated in High-income Countries
3G	0.019	0.042**	0.003	0.060	0.054**	0.014
	(0.017)	(0.018)	(0.028)	(0.055)	(0.023)	(0.025)
Observations	112520	141721	38361	7928	63575	33331
R^2	0.22	0.20	0.22	0.24	0.14	0.17
Mean outcome	0.187	0.242	0.267	0.163	0.183	0.183

Panel B: Educational Attainment and Polity Democracy Status

	(1) Primary Educated in Non-democratic Countries	(2) Secondary Educated in Non-democratic Countries	(3) Tertiary Educated in Non-democratic Countries	(4) Primary Educated in Democratic Countries	(5) Secondary Educated in Democratic Countries	(6) Tertiary Educated in Democratic Countries
3G	0.008	0.054**	0.034	0.042*	0.032*	-0.011
	(0.021)	(0.022)	(0.032)	(0.024)	(0.017)	(0.023)
Observations	75408	99697	28997	45040	105599	42695
R^2	0.21	0.21	0.23	0.22	0.16	0.19
Mean outcome	0.215	0.252	0.258	0.137	0.197	0.208

Panel C: Employment Status and Country-level Unemployment Rates

	(1) Employed in Below Median Unemployment Rate Countries	(2) Unemployed or Involuntarily Part Time Employed in Below Median Unemployment Rate Countries	(3) Out of Labor Force in Below Median Unemployment Rate Countries	(4) Employed in Above Median Unemployment Rate Countries	(5) Unemployed or Involuntarily Part Time Employed in Above Median Unemployment Rate Countries	(6) Out of Labor Force in Above Median Unemployment Rate Country
3G	-0.003	0.088**	-0.016	0.037**	0.045	0.107***
	(0.021)	(0.039)	(0.024)	(0.019)	(0.033)	(0.028)
Observations	110436	26277	44951	109665	28662	44365
R^2	0.20	0.28	0.23	0.15	0.24	0.20
Mean outcome	0.179	0.275	0.160	0.229	0.333	0.201

Panel D: Within-country Income Terciles and Country Income Groups

	(1) Lowest Income Tercile in Low- or Middle-income Countries	(2) Middle Income Tercile in Low- or Middle-income Countries	(3) Highest Income Tercile in Low- or Middle-income Countries	(4) Lowest Income Tercile in High-income Countries	(5) Middle Income Tercile in High-income Countries	(6) Highest Income Tercile in High-income Countries
3G	0.075***	0.017	-0.011	0.069***	0.011	0.007
	(0.017)	(0.019)	(0.016)	(0.023)	(0.023)	(0.023)
Observations	116687	113624	110656	49583	49110	47978
R^2	0.21	0.22	0.21	0.16	0.15	0.15
Mean outcome	0.187	0.201	0.223	0.166	0.135	0.144

Notes: OLS regressions. Standard errors, clustered by district and country-year, in parentheses. The specification used in Columns (1) to (4) of Panels A to D is similar to that of Column (4) of Table 1.1. In Panel A and D, lower- and middle income groups include low-income, lower-middle-income and higher-income groups according to the World Bank in 2018. In Panel B, non-democratic countries are those with an average polity score lower than 6 between 2008 and 2018, whereas democratic countries have an average polity score of 6 or higher. In Panel C, the median unemployment rate across the country-averages between 2008 and 2018 is 7%. See notes to Table 1.3 concerning the definitions for income in Panel A (here with terciles rather than quintiles) and employment status in Panel B. For the same reasons as in Table 1.3, we restrict the analysis in Panels B, C and D to those between 25 and 60. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 1.E.22: The Effect of Material Well-being and Satisfaction with Life and Institutions on the Desire to Emigrate, prior to 3G Coverage

Panel A: Material well-being				
	(1)	(2)	(3)	(4)
X-variable:	Household income (log)	Material prospects first principal component	Job climate index	Financial well-being index
Outcome: Desire to migrate	0.000	-0.100***	-0.076***	0.031***
	(0.002)	(0.010)	(0.007)	(0.005)
Observations	181,505	166,227	179,298	49,864
R^2	0.19	0.19	0.19	0.21
Panel B: Life satisfaction and optimism				
	(1)	(2)	(3)	(4)
X-variable:	Optimism index	Daily experience index	Life evaluation index	Life purpose index
Outcome: Desire to migrate	-0.078***	-0.105***	-0.017***	0.025***
	(0.009)	(0.011)	(0.004)	(0.004)
Observations	181,341	180,996	170,071	49,807
R^2	0.19	0.19	0.17	0.21
Panel C: Institutional satisfaction				
	(1)	(2)	(3)	(4)
X-variable:	Law and order index	Corruption index	Community basics index	Trust in government first principal component
Outcome: Desire to migrate	-0.113***	0.060***	-0.117***	-0.125***
	(0.009)	(0.006)	(0.010)	(0.009)
Observations	181,505	174,984	181,505	146,794
R^2	0.19	0.19	0.19	0.18
Panel D: Mobile Banking and Remittances				
	(1)	(2)	(3)	(4)
Independent variable:	Owns a bank account	Used cellphone to receive cash in the last 12 months	Received money or goods from friend/ family from same country	Received money or goods from friend/ family from another country
Outcome: Desire to migrate	0.017***	0.053***	0.010***	0.068***
	(0.006)	(0.014)	(0.004)	(0.007)
Observations	53,658	53,268	162,009	162,009
R^2	0.18	0.18	0.19	0.19

Notes: OLS regressions. Standard errors, clustered by district and country-year, in parentheses. The specification used in Columns (1) to (4) of Panels A to D is similar to that of Column (4) of Table 1.1. In Columns 1 and 2 of Panel D we omit the district-level time trends, as we only have 3 time periods available. We only exclude the control variables related to local amenities as some of these amenities are used in the construction of the GWP indices. All independent variables of interest in Panels A to C are GWP indices, except for "(log) household income" (which is the reported log of per capita household income), "material prospects" (a first principle component of the following questions (weights in parentheses): living comfortably on present income (0.69), now is a good time to find a job (0.34), and not having enough money to afford food (-0.65)), and "trust in goverment" (a first principle component of four questions related to trust in the government, as constructed by Guriev, Melnikov and Zhuravskaya (2021).). For all items in Panel A to C a higher value of the independent variable of interest implies a higher value of the item. For example, a higher value of "Material prospects first principal component" implies a better subjective evaluation of material well-being and a higher value of "Corruption index" implies a larger perception of corruption. For construction of the GWP indices, see `https://www.oecd.org/sdd/43017172.pdf` (Last accessed on 08-12-2021). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Chapter 2

Low-cost Language Learning: a Boost to Move? Evidence from Duolingo

2.1 Introduction

The rapid rollout of mobile internet and the subsequent adoption of smartphones has transformed international migration. Not only does this enable more than 4.5 billion people (GSMA, 2023) to search for information about opportunities abroad and stay in contact with family and friends, it also unlocks access to modern educational technologies.[1] One of the most popular educational technologies is language learning platforms, which strongly reduced the financial and convenience cost of foreign language learning. As foreign language skills are crucial to integrating into a host society (Adserà and Pytliková, 2016), this technology is of special interest to migrants. It is especially relevant as many (prospective) immigrants lack good language skills, which has generated interest in policies that try to remedy that (Foged, Hasager and Peri, 2024). Improving their language skills improves welfare of immigrants and natives through better integration and enables to reap the large gains from immigration from less to more productivity countries (Clemens, 2011).

A decrease of the cost of foreign language learning is expected to increase learning and improve skills, and increases the attractiveness of migration destinations with attainable languages. However, it is an empirical question to what extent it does. Moreover, it is a priori unclear how the introduction of a low-cost technology impacts the selection of migrants in terms of general and language skills and whether it improves the preparedness of migrants for host-county labor markets. Hence, examining how the introduction of *low-cost language learning* platforms has affected migration patterns and migrant integration provide valuable information to academics understanding human capital decisions of cross-border workers and policy makers to design sound immigration and integration policies.

In this paper, I study the effect of online low-cost language learning on migration and integration patterns using the staggered rollout of language courses on the platform Duolingo. Duolingo provides freely available language courses, includes gamified elements[2] and targets learners at low initial levels of proficiency as it offers learning a target language using instructions in a particular source language. Each course consists of a series of topical lessons lasting a few minutes consisting of several items which consist of matching through listening, translating, speaking, multiple choice items and stories for some courses. Importantly, Duolingo was found to be an effective way for English speakers to study Spanish at a beginner level, indistinguishable from in-class instructions (Rachels and Rockinson-Szapkiw, 2018; Ersoy, 2021). Partially due to these features, Duolingo soon became the market leader on the language learning app market and has 74 million monthly active users in 2023.[3] Duolingo first rolled out courses in 2012, starting with three courses in 2012, and gradually rolled out more courses in subsequnt years. As of 1st of January 2023, 84 courses to living languages are have been introduced.

To study the impact of low-cost language learning, I exploit the feature that Duolingo courses are *dyadic* in nature by enabling learning from source to target language. Combining this with information on languages spoken by country gives rise to rich variation in low-cost language learning availability on the country-language pair level and country pair level. This is best illustrated with an example. After introduction of a course between French and English, learning English

[1] In accordance, the share of population engaging in online courses has risen rapidly in recent years. Across the EU, the percentage of population doing any online course has risen from 5% in 2015 to 15% in 2023. Figure 2.B.1 shows the share of population doing an online course across the EU over time, showing considerable divergence across countries.

[2] To motivate and engage learners, Duolingo provides several *gamified* elements, such as competing on a leaderbord, virtual badges for additional content and regular reminders to self-set target. A course typically provides approximately several thousands of lessons, including 10s of thousands unique sentences and 1,000s of unique words, and is under constant development.

[3] See `https://www.businessofapps.com/data/language-learning-app-market/`. The market share of Duolingo in 2023 is 60%.

for those in French-speaking countries becomes easier. I study (among others) whether interest in English increased in French-speaking countries, whether English test scores improved among native French speakers, whether migration intentions and flows increased between French-speaking and English speaking countries, and whether migrants from French-speaking countries hosted have better language skills upon arrival and integrate faster economically. Beyond this stylized example, reality is less dichotomized as many languages are spoken by a part of population in a given country. I proxy the intensity of exposure to Duolingo between an origin country and target language by the probability that a given course enables a randomly picked person in the origin to be able to newly learn the target language. Likewise, I calculate the probability a course between two languages enables two randomly picked individuals in two countries to newly communicate.[4] Combining these communication probabilities with course rollout dates, I construct a measure of exposure to low-cost language learning.

All analyses in this paper compare individuals or geographic units more exposed to those less exposed to Duolingo before and after exposure, which requires the assumption of counterfactual parallel trends. This is jeopardized if courses are rolled out in anticipation of stronger future demand for language learning (which could be related to trends in language skills and migrant integration). I argue that this is unlikely the case as Duolingo was supply-constraint for most of its history. As Duolingo was established in 2011, it had to rapidly keep up with demand and rapidly developed courses for the largest languages. Hence, Duolingo likely anticipated *levels* in demand for language learning across two languages rather than *trends*. In addition, several courses were developed by the user-community, who had no other motive than making their language available for others to learn. When it comes to migration-related outcomes, it is particularly unlikely that future trends are anticipated. Duolingo courses are not primarily targeted towards migrants; most users learn for work- or education-related reasons. Across applications, I nevertheless test for the presence of pre-trends. I proceed in three steps. First, I study how course availability has shaped language-related outcomes, such as search interest on Google Trends, test scores in the TOEFL and GRE scores and the impact on traditional language learning. Second, I estimate its effect on migration aspirations in the Gallup World Polls (GWP) and flows through a theory-guided gravity model. Third, I study its effect on the language and general skills of migrants upon arrival on subsequent integration in the European Union and the United states.

First, I study how the rollout of Duolingo courses affected online search behavior for Duolingo and available languages. I find that the first relevant course rollout in a country strongly increases search for Duolingo courses on Google Trends. Moreover, I find that online search interest in available languages increases as well, which suggest that the availability of a language course also spurs off-platform learning. To study the impact of low-cost language learning on language skills, I turn to the English-languaged TOEFL and GRE tests. These tests are typically taken by young individuals, among many whom take the test to apply for an educational program of job abroad. Using the TOEFL test, I find that exposure to Duolingo increases passive elements of language skills (reading and listen) but not active elements of language skills (writing and speaking), which is in line with the skills online platforms such as Duolingo predominantly practice. Using the GRE, which test both language and quantitative skills, I find that Duolingo exposure increases language skills, but does not change the number of test takers or changes the scores on quantitative scores, suggesting that the improved English skills are not driven by selection, but by learning.

To understand how the availability of new language learning technologies interacts with traditional language learning, I study how the participation in adult and school-based foreign language learning develops when Duolingo becomes available. Using information on German learning in *Goethe* language institutes across the world, I find that the availability of Duolingo

[4] In absence of information of the joint distribution of language skills within countries, these require the simplyifying assumption that language knowledge is randomly among a country's population.

courses to German seem to reduce the number of course participants, but not the number of exams taken. Using information on foreign languages taught in schools across the EU, I find that full Duolingo exposure increases the share of pupils learning available languages by about two percentage points.

Second, I study the effect on migration intentions and flows. For the main analysis, I use the Gallup World Polls (GWP), which is a representative survey in more than 150 countries globally and includes a question on migration aspirations and where people would like to migrate. A large advantage of the migration intention data over the flow data is that information is available for all destination countries globally. Reassuringly, migration aspirations have been shown to be predictive of subsequent migration flows (Tjaden, Auer and Laczko, 2019) and are thus a meaningful income. Motivated by a random utility model of migration, I estimate a gravity model of migration including exposure to Duolingo. I can control for pair, origin-year and destination-year fixed effects, which can control for multilateral resistance terms as well as general push and pull shocks. I find that upon large increases of Duolingo exposure, migration intentions increase gradually in the first three years. On average, I find that migration intentions increase by 45% if a course newly enabling communication between origin and destination country becomes available. I find that less than half of this effect is driven by individuals who otherwise would not have desired to emigrate, and the rest by diversion of migration intentions towards destinations with learnable languages. To examine whether this has also led to larger migration flows, I turn to the OECD yearly bilateral migration migration flow data, which includes all origin countries, but only 37 OECD countries as destinations. I find no conclusive evidence that migration flows have increased, but estimates are relatively imprecise. As an alternative, I use a recent database of scholarly migration flows developed by (Akbaritabar, Theile and Zagheni, 2024) between all countries in the world. I find that exposure to Duolingo increases these flows by 4 percentage points, which is driven by courses from English to other languages.

Third and last, I study the integration of migrants across the EU using the EU Labor Force Survey (EU LFS) from 2007–2021. Exploiting variation in timing of the rollout of relevant Duolingo courses across country pairs, relative to the arrival cohort of immigrants and survey interview timing, the effects of pre-arrival and post-arrival exposure to Duolingo can be disentangled. An advantage of this data over data from a single country is that it includes many destination countries, enabling me to control for origin-year fixed effects which capture time-varying selection into migration. Hence, I exploit variation in the availability of low-cost language learning across origin group-destination pairs over time. I find that the probability that migrants speak a language at least at beginner level increases by 20 percentage points. These effects are driven by migrants who found a job already before arrival, family migrants and those coming for education. Moreover, I find that the availability of Duolingo increases the share of migrants coming with a job upon arrival and the share of male migrants. The initial gain of employment upon arrival of about 10 percentage points decreases over time in the destination. Nevertheless, exposure to Duolingo after arrival also has a strong positive effect on employment rates.

I complement this with a study of migrant integration in the US, using the American Community Survey (ACS). Although I can not control for origin-year fixed effects, there are no trends in migrant characteristics before strong increases in Duolingo exposure. Contrary to the EU, language skills within the first year after arrival do not improve. Exposure to Duolingo after arrival does increase language skills and increased employment rates by four percentage points. In both the EU and US results seem to suggest that the availability of Duolingo before arrival shifts the educational composition of migrants towards low-skilled individuals. In addition, the English language intensity of immigrants' occupations in the US decreases. This suggests that the basic skills that Duolingo enables users to attain could be sufficient for low-skilled workers to find suitable employment that requires some language skills but less than average, which reduces the average language requirements of immigrant jobs.

The remainder of the paper is structured as follows. Section 2.2 reviews related literature and discusses this paper its contributions to it. Section 2.3 introduces the language learning app Duolingo and describes the roll-out over time and languages. Section 2.4 sets up a simple model of investments in language learning to understand how a decrease in the cost of language learning can impact migration patterns, and introduces and discusses the main empirical strategy. Section 2.5 encompasses an analysis how Duolingo course availability has impacted language learning and proficiency, section 2.6 studies its effects on migration aspirations and flows to OECD countries and among scholars and section 2.7 studies the effects on migrants' language skills and economic integration in the European Union and the United States. Section 2.8 concludes and discusses implications of this study.

2.2 Literature

This paper relates to three strands of literature: the literature on the economic, cultural and linguistic determinants of international migration, the literature on the role of new technologies in international migration, and – most closely related to this work – the literature on the role of language learning in international migration.

The first strand of literature is that aims to understand the drivers of (international) migration. Many authors have assessed economic determinants of international migration, focusing on the extent of migration flows as well as the selection and sorting of migrants Borjas (1987); Grogger and Hanson (2011). Furthermore, the role of language and culture in migrants' earnings has been thoroughly studies. Micro-level evidence has shown that relevant language skills contribute to higher labor market earnings (Chiswick and Miller, 1995; Dustmann and Fabbri, 2003). Apart from financial costs, it may also reduce the burden of applying for visas in a foreign language and many other important frictions (Jaschke and Keita, 2021). Linguistic distance between languages is a key determinant, as it enables individuals attain languages easier. Isphording and Otten (2011) show that language attainment among migrants in Germany strongly correlated to the distance between languages based on a measure of lexical distance between languages. Adserà and Pytliková (2016) survey the literature, finding that the host-country language premium ranges between 5 and 35%.[5] As these studies show that labor markets reward foreign language skills, countries sharing a language or speaking a similar language should be positively related to the size of bilateral migration flows. Belot and Ederveen (2012) have shown that between OECD countries cultural and linguistic distance is associated with lower migration flows. White and Buehler (2018) have shown that differences in individualism, uncertainty avoidance and perceived gender roles are the most important cultural impediments to international migration. Adserà and Pytliková (2015) have advanced the study of language by showing that lower linguistic distance is associated with larger international migration flows, after controlling for sharing a language. A fundamental limitation of such studies is that one can not control for all pair-level unobserved heterogeneity, such as unobserved cultural factors correlated to language. I contribute to this literature by showing that not only sharing a language or linguistic proximity between languages affects bilateral migration, but also the ease of learning a destination country language through available technology. Contrary to studies using time-independent measures of language and culture, I can do so controlling for unobserved dyadic factors.

The second strand of literature concerns (digital) technologies that affect international migration. The internet changed the velocity and way information spreads across the globe, which is likely to have large consequences for (prospective) international migrants. Chapter 1 of this

[5] These large differences supported motivated policy makers to introduce (obligatory) language courses for some immigrant groups. In a recent survey of the literature, Foged, Hasager and Peri (2024) conclude that language learning policies for refugees have positive effects on earnings in the short run.

thesis show that the worldwide rollout of 3G mobile technologies increased the desire and intentions to emigrate, using data from 120 countries of origin. In addition, their analysis suggests that preferred destinations change: as the internet lowers the cost of acquiring information about previously lesser known destinations, preferred destinations become more diverse. Böhme, Gröger and Stöhr (2020) show that online search behavior predicts migration flows, suggesting online search is important to finding information about potential destinations. Diving into one important element of the modern-day world wide web, Dekker and Engbersen (2014) conceptualize how social media has transformed international migration by interviewing individuals from three origin countries in the Netherlands, showing that social media reduced perceived distance to family and friends at home and enabling migrants to leverage weak social networks to organize migration and integration, thereby facilitating migration. I contribute to this literature by showing how a very specific type of internet technology, the availability of language learning apps, shape migration aspirations.

The third and final strand of literature is that of the role of language learning in international migration. The seminal work of Bleakley and Chin (2004, 2010) has documented that immigrants age at arrival is crucial for attaining host-country language skills. Using this relationship, they found that immigrants' language skills increases earnings and intermarriage with natives. Dustmann (1999) has shown that temporary migrants with a longer horizon in Germany have a larger incentive to attain the host country language and have better languages skills. Along similar lines, Wong (2023) exploits random allocation of refugees across the linguistic regions of Switzerland, finding that linguistic proximity is related to better labor market outcomes through quicker language learning. Adserà and Ferrer (2021) find that linguistic distance to English reduced earnings upon arrival more for college educated than for non-college educated migrants in Canada. Furthermore, they find that labor market earnings of men from linguistically distant countries increased substantially over time. These two studies are suggestive of the fact that labor market potential increases when language learning is easier. In line with this, policy makers have introduced (mandatory) language courses for some immigrant groups, particularly refugees. Using quasi-experimental variation in the availability such integration policies, their efficacy has been studied in several countries. In a recent survey of the literature, Foged, Hasager and Peri (2024) conclude that integration policies for refugees, including language training, have small positive effects on earnings in the short run. Only few studies have examined the role of language learning in isolation. Foged and Van der Werf (2023) study the availability of language learning in Denmark for refugees through variation in the proximity to language training centers, finding strong impacts to stay in that locality and some positive results on labor market outcomes.[6]

Contrary to language learning after arrival, less is known about language learning prior to migrating. Nocito (2021) show that English as a language of instruction in Master's degree strongly increase graduate migration from Italy. Huber and Uebelmesser (2023) and Jaschke and Keita (2021) study the closing and opening of German language institutes (Goethe Institutes; GI) where up to 100,000 individuals study German each year. Huber and Uebelmesser (2023) show that six years after opening a GI international migration flows from the country where the institute located sends more migrants to Germany. Jaschke and Keita (2021) find in the same setting that GIs affect the self-selection of migrants: upon arrival they have better language skills and are higher educated. Nevertheless, course participants still pay a considerable fee for a language course in the Goethe institutes. Freely accessible online language courses provide an opportunity to many people across the world, including to those who would not be able to afford institutional language courses. Furthermore, compared to the yearly attendance of the Goethe Institutes, the number of Monthly Active Users on Duolingo is order of magnitudes larger. Therefore, studying

[6] Di Paolo and Mallén (2023) use a similar design in Barcelona, finding that distance to a language center improves Catalan language skills, but not labor market outcomes.

how low cost language learning affects international migration is complementary to studying the availability of certified language courses and provides a setting for studying the availability of language learning not in an isolated setting of migration to Germany, but practically the whole world.

2.3 Duolingo: an Educational Technology

Duolingo is a freely available online language learning platform with gamified features, consisting of bilateral and directional courses: a course enables one to learn a specific target using a source language. Hence, it particularly targets language learning at low initial levels of target-language proficiency. Although Duolingo was not the first online language learning platform, it gained considerably more traction than its competitors for two reasons: Duolingo is available for free and it contains many features that increase user engagement (Shortt et al., 2023). Users can set targets and Duolingo reminds users to meet the target and users can compete against each other on a leaderboard. This allowed Duolingo to amass a market share of 64% in 2022, more than six times that of its closest competitor.[7] Duolingo's entry threshold is very low. First of all, most of Duolingo's content is available for free[8] Secondly, it allows learning from scratch, as the language of instruction is the source language the user has command over. This feature gives rise to rich variation in availability of low-cost language learning across speakers of different languages.

Users can access the platform through a desktop browser or a mobile application. Figure 2.B.2 shows a series of typical tasks of a course: it includes translation, sentence completion, written conversations and dictation. These elements are mostly very helpful to attain passive skills such as reading and listening, but potentially lack active elements of languages, such as writing and speaking. Moreover, Duolingo's learning philosophy is based on learning by doing, and does not include explicit grammar exercises (Freeman et al., 2023). By forcing users to set learning targets and reminding users of these regularly, Duolingo keeps users engaged, which could aid in overcoming commitment problems.[9] The rightmost screenshot of Figure 2.B.2 shows how Duolingo encourages users to fulfull their targets. In addition, Duolingo fosters engagement through several gamification elements, which allow users to collect points through learning and to compete against others on the platform.

As Duolingo provides language learning of a target language by instruction in a source language, it naturally targets language learning at low levels of proficiency. To also target more advanced users, a placement test is offered, so that users can start at an appropriate level. Many courses are extensive: they comprise 1,000s of practice lessons; several courses reach up to and including CEFR level B2.[10] Nevertheless, courses were gradually extended over time, and not all courses include lessons up to B2 to date.[11] Duolingo was found to be an effective way for English speakers to study Spanish at a beginner level, indistinguishable from in-class instructions (Vesselinov and Grego, 2012; Rachels and Rockinson-Szapkiw, 2018; Ersoy, 2021). Moreover, Duolingo's research department has extensively studied the efficacy of its platform on reading and listening skills, finding outcomes on par with several semesters of university courses (Jiang et al., 2021*a*,*b*). Nevertheless, there is less independent evidence on its efficacy at more advanced stages of language learning and on speaking and writing skills. However, Duolingo should not

[7] https://seekingalpha.com/article/4570169-duolingo-stock-gamified-learning-great-growth-potential

[8] Duolingo is free to use with advertisements. An ad-free premium version is available too. In 2024, Duolingo had more than 5 million paying subscribers for the premium version.

[9] The recent work of Brade et al. (2024) shows that overcoming commitment problems improves academic performance.

[10] A B2 user can understand main ideas from complex text, interact with native speakers without strain and produce detailed test on a wide range of subjects.

[11] For an overview of the extent of specific courses, see https://duolingodata.com/.

be seen as a pure substitute to traditional language learning. After obtaining a basic level of a language on Duolingo, an individual can continue learning a language at an institution that offers certification. As the individual possesses some initial language knowledge, the individual can start institutional language learning at a higher level, reducing the total cost spent on language certificates.

The first courses where made available in 2012, with English to Spanish, Spanish to English and English to German.[12] An important element of course development is that subsequent courses were built using strong support of volunteers during most of the history of Duolingo, who suggested courses in the *Duolingo Incubator*.[13] Nevertheless, many Duolingo language courses are introduced with commercial motives to attract more users to the platform and increase engagement. I discuss the determinants of rollout of Duolingo courses in section 2.4.2, and the uptake of courses in section 2.5.1.

I obtained information on the roll-out dates from an online source.[14] Courses go through three stages of development. Although courses may be available to a smaller audience before the final phase, I identify the rollout date as the day the course entered the final phase.[15] To validate the relevance of the roll-out dates, I study the impact of on online search behavior for Duolingo and available target languages in section 2.5.

I discard Esperanto, Klingon, High Valyrian and Latin, as they are not widely used and therefore irrelevant to international migrants. Until 2022, 110 courses involving two existing languages have been developed, of which 87 have reached the final phase. I do not include Irish, Hawaiian and Scottish Gaelic, because of the low number and lack of reliable information on the current day speakers, leaving 84 courses. Figure 2.3.1 shows the number of courses rolled out by year of introduction, as well as the total number of monthly active users across all courses. The first four courses were rolled out in 2012 and the most recent introduction took place in 2021. Figure 2.B.3 shows all available courses in 2021 in a Sankey diagram by connecting source and target languages. In total, there are 23 unique source- and 30 unique target languages. The diagram highlights that English is the most prevalent source (27 courses) and target language (22 courses), but that there is considerable variation across other languages.

[12] More than half a million people signed up for the *beta* versions in June 2012, mobile applications on iOS and Android were released in November 2012 and June 2013, respectively.

[13] This option ended in 2021, see `https://blog.duolingo.com/ending-honoring-our-volunteer-contributor-program-2/` and `https://duolingo.fandom.com/wiki/Incubator`. As an example, the English to Russian course was fully developed by volunteers.

[14] These dates can be found on `https://duolingo.fandom.com/wiki/Course_list`. I verified these dates using the available languages on Duolingo through the Wayback Machine.

[15] In the first phase, courses are being developed, but can't be used by the general public. In the second (beta) phase the course is in testing and can be accessed by users, although it is typically not widely used. In the third and final phase the course is operational and widely used. As the number of users in the beta phase is small, I use the final phase as the rollout date. As it takes time for a course to achieve a wide audience, especially when no other course is available from the same source language or in early years when Duolingo had limited visibility, the introduction date of a course may be an imperfect proxy of when a course becomes widely adopted. Section 2.5 shows that online search interest in Duolingo is elevated but still very low in the quarters before a country experiences the first rollout of a relevant language course, increasing rapidly in the first 1.5 years after introduction.

FIGURE 2.3.1: Introduction of Duolingo Courses and the User Base over Time

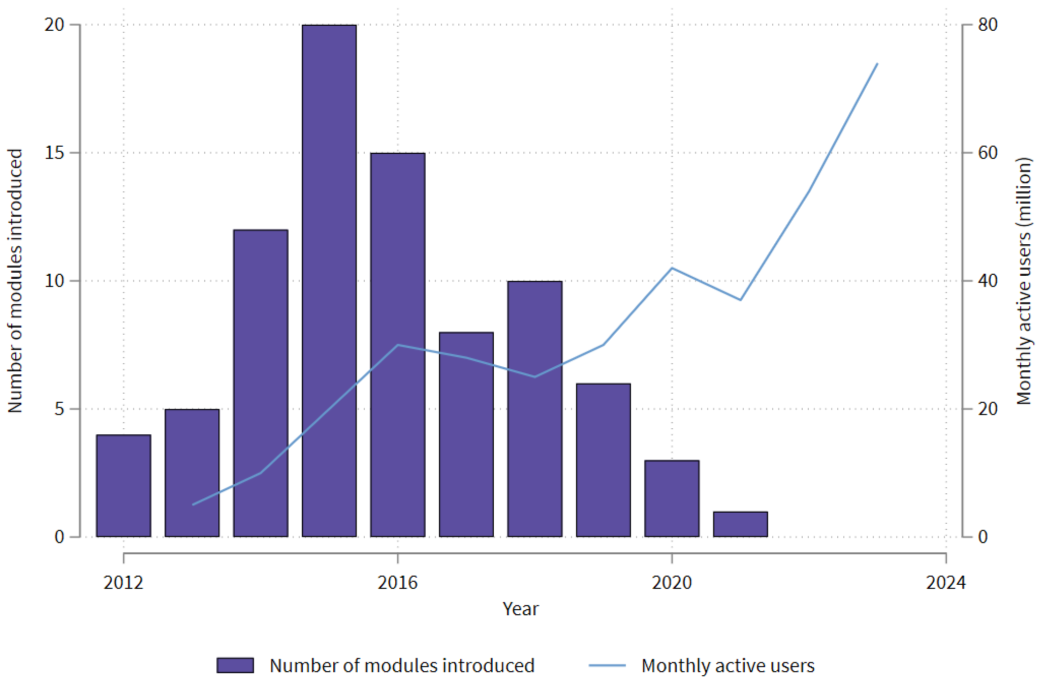

Notes: Total number of courses by year of introduction and the number of Monthly Active Users (MAUs). A course is counted within a year if it reached the final stage of development within that year. MAUs are defined as unique users who engage with our Duolingo App or the learning section of the Duolingo website each month.

In 2013, Duolingo had 5 million monthly active users, which gradually rose to 74.1 million in 2023, of whom 21.4 million are active on a daily basis. In 2023, 1.5 billion hours were spent on learning, suggesting an average of about 20 hours per user. Duolingo's user base consist for 50% of females, and is relatively young: 29% of users is aged between 18 and 24, 26% is between 25 and 34, 18% between 35 and 44, 13% between 45 and 54, 9% between 55 and 64, and just 6% over 65.[16] Users' main reason for learning on Duolingo vary across country-language pairs and age categories. English learners in the U.S. are most likely to use the platform for work-related reasons, whereas those in non-English speaking countries and younger generations are more likely to use it for school.[17] Although I use the rollout of Duolingo on the language level, it is still insightful to see where Duolingo's users are located. Table 2.B.2 shows internet traffic data to Duolingo by global region. Although 27% of traffic comes from North America, considerable traffic originates from other regions across the world.

2.4 Model and Empirical Strategy

2.4.1 A Model of Language Learning and Migration

Duolingo is available to anyone in possession of a desktop or smartphone with internet connection. The use of the ad-based version of Duolingo is free of cost regardless of the number of courses and intensity of learning. This stands in stark contrast to traditional institutional language courses. For example, in 2024, a German course at CEFR A2 level costs 340 euro in

[16] Data obtained from `https://www.similarweb.com/website/duolingo.com/#demographics`, which only lists information on users 18 and over. Data on the numbers of users in the US suggests that about 20% of users are below 18.

[17] See `https://blog.duolingo.com/english-learner-motivations/`.

Colombia and 180 euro in Bangladesh (Goethe Institute, 2024*a*,*b*). Motivated by this, I model the availability of a relevant Duolingo course as a decrease in the cost of language learning in a modified random utility model (RUM) of migration, in which agents choose how much to invest in language learning before observing their utility from each destination.

I model costly foreign language acquisition and migration from origin country o speaking language S to a destination country d where language T is spoken as a two-step process. I focus on one language skill as a Duolingo course enables to learn one specific *target* language, but discuss the extension to multiple language skills below. I assume that the discount factor is one. In the first step, individuals choose the optimal level of a foreign language skill $l_{oS}^T \in (0,1)$. As individuals are ex-ante indistinguishable, optimal language skill decisions are homogeneous and l_{oS}^T is not indexed by i. The cost of acquiring a target language skill T is convex in l_{oT}:

$$c_{oST}(l_{oS}^T)^2 = \frac{\kappa_{oST}}{2(1+\eta_{oST}Duolingo_{ST})}(l_{oS}^T)^2 \tag{2.1}$$

Here, c_{oST} depends positively on parameter κ_{oST} and negatively on the availability of a Duolingo course with effectiveness η_{oST}. In the second step, an individual observes the idiosyncratic benefits of migration to all alternatives d, ϵ_{iod}, and migrates to the destination which offers highest utility. Following the literature, I model the idiosyncratic term as an i.i.d. EVT-1 shock (Beine, Bertoli and Fernández-Huertas Moraga, 2016), which gives a convenient closed-form solution for destination choice probabilities. The utility for individual i with language skills l_{oS}^T from country o when moving to d is:

$$U_{ioSd} = \ln w_{od} + \epsilon_{iod} = \mu_{od} + l_{oS}^T b_{oSd}^T + \epsilon_{iod} \tag{2.2}$$

Here, μ_{od} are earnings in absence of relevant language skills in country d, l_{oS}^T are language skills of individuals from o in language T and b_{oSd}^T denotes the return to language skill for individuals from o speaking language S in country d, which I assume to be always finite, positive, and strictly positive for at least one d. For sake of simplicity, I do not explicitly include migration costs, but these can be thought of as absorbed in μ_{od} and in b_{oSd}^T, if language skills reduce migration costs. Using the properties of the EVT-1 shock, utility maximization gives the following migration probabilities:

$$\mathbb{P}_{od} = \frac{e^{\mu_{od}+l_{oS}^T b_{oSd}^T}}{\sum_{d'} e^{\mu_{od'}+l_{oS}^T b_{od'T}}} \tag{2.3}$$

The denominator is the sum of the deterministic part of utility in all potential destinations d', which also includes the origin. The number of migrants from o to d is given by $M_{od} = \mathbb{P}_{od}P_o$, where P_o is the initial population of origin o. Dividing the share of individuals migration $\mathbb{P}_{od}$ by the share staying $\mathbb{P}_{oo}$ and taking the natural logarithm gives a convenient expression for the log odds of migrating to d over staying in o, which is independent of the utility of alternative destinations (Bertoli and Fernández-Huertas Moraga, 2013; Beine, Bertoli and Fernández-Huertas Moraga, 2016):

$$\ln\left(\frac{\mathbb{P}_{od}}{\mathbb{P}_{oo}}\right) = \mu_{od} - \mu_o + (b_{oSd}^T - b_{oSo}^T)l_{oS}^T \tag{2.4}$$

The log odds of migrating increase in l_{oS}^T if returns to the language skill are higher in the destination country than at home ($b_{oSd}^T > b_{oSo}^T$). However, this result does *not* imply that total migration to d is increasing in l_{oS}^T. Section 2.A.1 shows that this is only the case when b_{oSd}^T exceeds the migration probability-weighted in all alternative destinations (including the origin). equation 2.1 shows that total emigration from o increases only when the weighted foreign return to language skills exceed domestic returns to the language skill.

Returning to the first step, individuals from o decide how much to invest in the language skill, given their expected utility from language skills. The expected indirect utility for someone from country o in period one is given by the expected utility in period two minus the cost of language learning:

$$U^*_{oS} = \mathbb{E}\left(U_{oSd}\right) = \sum_d U_{oSd}\mathbb{P}_{od} - c_{oST}(l^T_{oS})^2 \tag{2.5}$$

Language skills affect indirect utility in two ways. First, language skills increase the utility of destinations where the skill is valuable (keeping migration probabilities constant). Second, language skills increase migration probabilities towards destinations where the skill is valuable (keeping the destination-specific utility constant). As the indirect utility function is optimized w.r.t. migration probabilities, I can use the Envelope Theorem: $\frac{\partial U^*_{oS}}{\partial \mathbb{P}_{od}} = 0$. Hence, the net change in indirect utility from the second channel is zero. I obtain the following first order condition w.r.t. l:

$$2c_{oST}l^T_{oS} \approx \sum_d \mathbb{P}_{od}b^T_{oSd} \tag{2.6}$$

The left hand side represents the marginal cost of one unit of language skills, whereas the right hand side represents the expected marginal benefit. Importantly, the migration probabilities $\mathbb{P}_{od}$ depend on l^T_{oS}. An equilibrium pinning down l^*_{oT} exists and is unique.[18] For most countries of origin, migration probabilities are small compared to the probability of staying. In this low migration case, the right hand side only weakly depends on l^T_{oS} and an expression for optimal language skill acquisition l^{T*}_{oS} can be derived:

$$l^{T*}_{oS} \approx \left(\mathbb{P}_{oo}b^T_{oSo} + \sum_{d\neq o}\mathbb{P}_{od}(0)b^T_{oSd}\right)\frac{1+\eta_{oST}Duolingo_{ST}}{\kappa_{oST}} \tag{2.7}$$

Here, $\mathbb{P}_{od}(0)$ denotes the migration probability in absence of the language skill. Two motives for language learning can be isolated. First, higher earnings on the domestic labor market increase incentives for language skill acquisition. Second, as language skills increase earnings abroad, the expected benefit is a migration-probability weighted average of returns to skill abroad. Plugging l^*_{oT} from equation 2.7 into equation 2.4 yields the following expression for the log odds of migration:

$$\ln\left(\frac{\mathbb{P}_{od}}{\mathbb{P}_{oo}}\right) = \mu_{od} - \mu_o + (b^T_{oSd} - b^T_{oSo})\left(\mathbb{P}_{oo}(0)b^T_{oSo} + \sum_{d\neq o}\mathbb{P}_{od}(0)b^T_{oSd}\right)\frac{1+\eta_{oST}Duolingo_{ST}}{\kappa_{oST}} \tag{2.8}$$

The log odds of migration depend on the distribution of returns to language skills across countries. Importantly, it depends linearly on the *difference* in earnings abroad and at home. Hence, both the size and magnitude of the effect of low-cost language learning (an increase in $Duolingo_{ST}$) depends on domestic (in o) and foreign (in d) returns to language skills. Moreover, the term in large brackets represents the degree to which language skills respond to changes in costs: it depends on the migration probability-weighted return across all countries. As most people never migrate, migration probabilities are small. Thus, this term is particularly large for foreign languages which are rewarded on domestic labor markets, such as English. Although equation 2.8 pertains to the low migration limit, l^* is an increasing function of $Duolingo_{ST}$ in

[18] To see this, note that the right hand side is already strictly positive when l^T_{oS} is 0 ($\sum_d \mathbb{P}_{od}(l=0)b^T_{oSd} > 0$), finite when l^T_{oS} is large and everyone migrates to the country with largest returns to skill ($\lim_{l\to\infty}\sum_d \mathbb{P}_{od}(s)b^T_{oSd} = max_d b^T_{oSd}$), and that its derivative w.r.t. l^T_{oS} is always positive, as with increasing l^T_{oS} the migration probability weights on destinations with larger b^T_{oSd} become larger.

more general cases. However, the effect is generally not linear and its precise functional form depends on the distribution of mean wages and return to language skills across countries.

Limitations. The model makes several simplifying assumptions. First of all, the model assumes that the EVT-1 shock realizes only after decisions on language skills, which renders language learning origin country-specific but not individual-specific. However, in reality individuals have heterogeneous skills and migration preferences and are more likely to invest in language skills relevant to preferred destination countries. Hence, effects are plausibly heterogeneous across individuals in the origin country. Nevertheless, the main mechanism remains the same: given expectations about migration probabilities, decreasing the cost of language learning increases language learning, which increases foreign earnings conditional on migrating and increases migration probabilities. Moreover, I assume that the EVT-1 shock is i.i.d. distributed, which is a strong assumption. In reality, preferences on the individual level are driven by preferences for (unobserved) country characteristics, which are likely to be correlated across countries. Relaxing this assumption introduces a complex error structure which depends on the characteristics of alternative destinations (Beine, Bertoli and Fernández-Huertas Moraga, 2016). I come back to the consequences of potential violations of the i.i.d. assumption in section 2.4.2. Furthermore, the model excludes non-earnings related explanations for language learning. As an example, language learning may have a consumption value as well (Huber and Uebelmesser, 2023). In this case, akin to the love-of-variety argument underpinning many trade models (Krugman, 1980), language learning may actually be stronger among more (linguistically) distant languages. This could explain the popularity of the Duolingo courses for Japanese, despite the limited returns to Japanese on domestic labor markets and limited number of migrants in Japan. Furthermore, it also excludes mechanisms affecting migration (intentions) through other mechanisms than language learning. Duolingo may spark interest may make target-language speaking countries more salient and induce interest in a country's culture. Such channels may ultimately affect migration.

The model does not incorporate heterogeneity across the foreign and general (i.e. education levels) skill distribution, which could have profound effects on migrant selection. First, as Duolingo targets language learning at low levels of proficiency, it could be particularly beneficial to those without or with limited initial proficiency. This could exert a downward pressure on the average foreign language skills of migrants. Second, the effects on selection into migration in terms of general skills is ex-ante ambiguous. On the one hand, as Duolingo is freely available, it could enable low-skilled individuals to whom language learning was previously too costly to acquire foreign language skills. On the other hand, foreign language learning could be complementary to general skills. As higher-skilled individuals have higher propensities to emigrate, they face stronger incentives to learn foreign languages, which in turn could improve their language skills upon migrating. Moreover, high-skilled individuals could use foreign languages as source languages (e.g. English) to learn other foreign languages. I empirically study the net effects of these forces by considering the impact of Duolingo availability on migrants' characteristics and language skills upon arrival in section 2.7.4. Secondly, the model only considers pre-migration learning, but not post-migration language learning. As Duolingo courses enable language learning both before and after migration, this poses a dynamic decision problem of prospective migrants deciding on the relative timing of language learning and migration. For example, migrants could delay language learning, knowing that they can utilize low-cost language learning technologies after arrival. Moreover, unanticipated access to language learning after arrival could further improve migrants' language skills. I empirically study the role of post-migration in section 2.7.4.

2.4.2 From Model to Empirical Strategy

To bring equation 2.8 to the data, one needs to proxy for returns to foreign language skills specific to the directed country pair od. In the following, I discuss how to approximate the returns using information on the distribution of languages spoken across countries.

Proxying returns to language skills. Returns to language skills play a crucial role in the model outlined above. However, these returns are not extensively measured across countries and languages. Nevertheless, the literature has provided estimates in several salient settings. Adserà and Pytliková (2016) survey the literature on immigrants' returns to destination country language skills, finding returns between 5 and 35% across contexts. Additional language skills can also increase earnings on domestic labor markets.[19] Returns to English has been found to be related to 10-50% higher earnings across countries.[20] This also extends to other widely spoken languages, such as French, German, Russian and Spanish, although the size of the estimated returns have been found to be smaller than for English.[21] An additional caveat to many of these studies of returns to skills is that they are correlational or identify local treatment effects among a specific sub-population.

In absence of such estimates of returns across dyads and languages, I proxy the foreign return of a Duolingo course between S and T by approximating the probability the course enables communication between two randomly drawn individuals in o and d, $\mathbb{P}(comm_{od}|DL_{S\to T})$. Using the distribution of speakers of languages across countries and the (strong) assumption that languages within countries are independently distributed, I can calculate the likelihood that a random individual speaking S in o can communicate with a random person in d, and how much this probability increases when one would also speak T.[22] This approach can be implemented both for spoken and official languages. In the special case that o and d do not share any languages, this measure is simply equal to product of the share of source- and target language speakers in the origin and destination, respectively: $\alpha_{oS}\alpha_{dT}$. The higher the share of source language speakers in the origin and the share of target language speakers in the destination, the larger the potential gains. However, the larger the existing overlap of languages between o and d, the lower the potential gains from a given Duolingo course. In section 2.A.3 I discuss how to calculate this object in the general case. However, this approach does not work for domestic returns to foreign languages. As most countries have one dominant language, the probability to newly communicate with someone in your own country from a language course is negligible. Instead, I assume that the domestic return to learning T from S is simply the product between the share of speakers of S and T, $\alpha_{oS}\alpha_{oT}$. The reasoning behind this is that non-native languages T that are valued on domestic labor markets amass a considerable number of speakers.

I obtain the share of speakers across countries and languages, α_{cL} and official languages by country from Melitz and Toubal (2014) and Ginsburgh, Melitz and Toubal (2017), who collected information about all languages spoken (natively) by at least 4% of population in most countries worldwide, as well as all official languages. For several missing observations, I complete the data using the most recent CIA World Factbook. I implement this approach for spoken source

[19] These include both foreign and non-foreign languages. In multilingual countries, which have multiple official or widely spoken languages, returns to non-foreign languages which are not one's native language may yield considerable returns. This concerns for example German for native-French Swiss and English in India.

[20] English has large returns in India 34% (Azam, Chin and Prakash, 2013), Turkey 40% (Di Paolo and Tansel, 2015), Poland 50-60% (Adamchik et al., 2019), China 10% (Wang, Smyth and Cheng, 2017), Germany 13% (Hahm and Gazzola, 2022; Stöhr, 2015), Spain (Isphording, 2013) and across Europe (Ginsburgh and Prieto-Rodriguez, 2011).

[21] In the US foreign language skills yield small positive returns (Saiz and Zoido, 2005), in Turkey, Russian, French and German (Di Paolo and Tansel, 2015), in Poland, French, German and Spanish (Liwiński, 2019), and across Europe for French, German and Spanish (Ginsburgh and Prieto-Rodriguez, 2011).

[22] This approach is borrowed from the trade literature: Melitz and Toubal (2014) study the effect of common languages on trade flows between countries. They use the probability two individuals speak the same spoken, native and official languages to explain the role of language in international trade.

languages, as command of a non-native source language may facilitate the learning of other target languages. For target languages, I calculate it using spoken languages as well, but also for official languages as a robustness test. The reason for the latter is that official languages may better reflect returns to language skills b^T_{oSd}. For example, spoken languages include minority languages and foreign languages, which have limited use on the country's labor markets.

Constructing foreign and domestic exposures. To construct a time-varying measure of foreign and domestic Duolingo exposure, I interact the proxy for returns with a binary indicator that takes value one the first year a Duolingo course between S and T has been fully available.[23] In practice, multiple courses ST may "bridge" a country pair od. For sake of simplicity, I take the largest exposure at any point in time. This implies that the treatment can increase more than once. For example, exposure from Netherlands to France increased for the first time when the English to French module became available, for the second time when German to French became available and for the third time when Dutch to English became available (see Figure 2.B.12). By construction, these measures are bounded between 0 and 1.

$$
\begin{aligned}
DL_{odt} &= \max_{S,T} \mathbb{P}(comm_{od}|DL_{S\to T})Duolingo_{STt} \\
DL_{oot} &= \max_{S,T} \alpha_{oS}\alpha_{oT}Duolingo_{STt}
\end{aligned}
$$

Figure 2.B.10 shows a heatmap of the intensity of exposure DL_{odt} across country dyads, showing considerably variation within and across dyads over time. Although by construction the largest exposure is one, plenty of dyads have treatment intensities less than one because a course only newly enables communication between part of the origin and destination country. Figure 2.B.11 shows the average exposure by origin countries show that there is plenty of variation across origin countries over time, except for several linguistically isolated countries, or countries that only speak languages not covered by Duolingo. Figure 2.B.13 shows the domestic exposure across countries, showing that about half of the countries are unexposed.

Throughout the remainder of this paper, I also estimate treatment effects of Duolingo exposure on outcomes varying at different levels than the country pair level. In such cases, I construct the exposure analogously. For example, for analyses on the country-language level I construct the exposure as the probability a Duolingo course enables communication to a person speaking target language T, which is an increasing function of α_{cS} and a decreasing function of α_{cT}.

Estimation equation. I estimate equation 2.8 by stacking all origin countries o, exponentiating both sides, plugging in the exposure measures and adding a well-defined error term with mean 1 and adding a time component t. This boils down to the following gravity model of migration with two continuously varying treatments:

$$
\frac{M_{odt}}{M_{oot}} = \exp\left[\beta_1 DL_{odt} + \beta_2 DL_{oot} + \gamma^{'}\mathbf{X_{odt}} + (\phi_{ot}) + \theta_{dt} + \psi_{od}\right]\eta_{odt} \tag{2.9}
$$

$\frac{M_{odt}}{M_{oot}}$ are the migration odds, comprised of M_{odt}, the number of (aspiring) migrants from country o to country d at time t, and M_{oot}, the number of (aspiring) stayers from country o. The regression coefficients β_1 identifies the effect of the availability of a Duolingo course enabling communication between the full populace of o in d on the value of living in d. Importantly, this is not the same as the effect of a typical Duolingo course on migration odds, due to the multilateral resistance exercised because alternative same-language destinations receive treatment at the

[23] For example, this switches from zero in 2012 to one in 2013 for the course English to Spanish, which has been introduced during 2012.

same time, which could be close substitutes to the focal destination. One can think of β_1 of the counterfactual effect on migration odds if a single destination would have received treatment. β_2 identifies the effect of being able to learn another domestically spoken language.

The covariate vector $\mathbf{X}_{odt}$ includes a dummy for joint EU membership, WTO trade agreements and bilateral trade flows. ψ_{od} captures unobserved pair-level unobserved factors. ϕ_{ot} and θ_{dt} indicate a set of origin-year and destination-year fixed effects that capture unobserved heterogeneity at those levels. Without origin-year level fixed effects and destination-year level fixed effects, language area-specific shocks temporally coinciding with the rollout of Duolingo courses could generate spurious effects[24] and to account for bias due to the inward and outward multilateral resistances.[25] As some of the terms following from the model only vary at the origin-time level, I estimate models with and without ϕ_{ot}. η_{odt} is an error term with unit mean.

2.4.3 Identification

The identification strategy underlying equation 2.9 is a generalized differences-differences strategy. Hence, to interpret the estimates of β's in prior sections as Average Treatment effects on the Treated (ATT), the following identifying assumptions need to be satisfied. First, there should be no anticipatory effects of Duolingo availability. It seems plausible that language learning and migration intentions are not affected by the future availability of Duolingo courses. Second, trends between exposed and unexposed units in absence of treatment should be common, for all levels of treatment intensity. For the multiplicative model of 2.9, this requires migration odds of origin-destination pairs to follow parallel trends in proportions in absence of treatment, conditional on the (origin-year) and destination-year fixed effects. This implies that growth rates in migration odds would have developed similarly in units with more or less Duolingo exposure, if the relevant Duolingo courses would not have been made available.

The parallel trends assumption could be violated if the rollout of courses anticipates trends in demand for language learning that are correlated to trends in migration outcomes. However, there are two reasons why this is likely not the case for Duolingo courses: courses did not target migrants and course rollout was plausibly supply-constrained and did not depend on trends in demand for learning, but rather on levels. The primary motive of language learning on Duolingo is often not preparation for migration. Duolingo asks users for the main motivation for studying foreign languages. In 2020, 33.8% indicated to learn English for school, 15.8% for work, 13.2% for brain training, 9% for family and 7.3% for cultural reasons. Only 12.6% learns English because of travel (which may be partially for tourism-related reasons) and 8.4% because of other reasons (which could include migration-related reasons). In fact, Duolingo's founder, Luis von Ahn, was motivated to increase domestic wages of Guatemalans on their home-country labor market. In addition, as Duolingo was released in 2012, the rollout of courses was initially supply-constrained as courses had to be developed from scratch. The development stage of the average course took several 100s of days. Courses are thus plausibly developed based on the current demand for language learning across a language pair, rather than on the growth rate. If Duolingo nevertheless aimed to anticipate future trends in the demand for language learning driven by migration trends, it is unlikely to happen for migration aspirations, which is the first step in the migration process.

To alleviate remaining concerns about parallel trends violations, I perform several diagnostic and robustness exercises. First, I study pre-trends in migration and language learning outcomes between strongly exposed and other country dyads before the strongly exposed received treatment. Using event study estimators that are robust to the issue of negative weights in staggered effects regressions de Chaisemartin and d'Haultfoeuille (2020); Nagengast and Yotov (2023) (as

[24]For example, this could happen when Hispanophone countries experience higher unemployment rates, or Anglophone countries introducing stricter immigration laws

[25]Section 2.4.4 discusses the challenges multilateral resistance poses to my estimates and how I deal with it.

further discussed in section 2.4.4), I find no evidence of differential pre-trends in migration intentions and flows. Moreover, there are no pre-trends in levels of the following language-related outcomes: online search interest in languages, language instruction among language skills among English test takers, and language skills conditional on migration in the EU. Second, courses may be developed to tailor to the origin and destination countries with the most speakers by language. To exclude that my main results are driven by these countries, I set the contribution for language-country pairs to zero if the country is the one where the language is spoken by the most people globally. For example, if a course is developed between French and English, it could be driven by anticipation of trends for individuals in France. However, also all other French-speaking countries to English speaking are affected by the Duolingo course. This exercises removes the former variation, but keeps the latter. I perform this exercise at the source language-origin and the target language-destination pair, as well as both at the same time. Third and last, I construct two Instrumental Variables approaches, one based on the notion that Duolingo courses are rolled out between languages with many speakers and the notion that course development is easier if there are more courses using the same source- or target languages. I re-estimate the model of equation 2.9 using a control function approach (Wooldridge, 2015). The results are robust to these two exercises.

What predicts course development?

To study what determines course rollout and hence Duolingo exposure, I analyze the timing of rollout on the language-pair and country-pair level. For the language-level analysis, I regress (1) a binary indicator of whether a course has been rolled out by the end of my sample period, and conditional on rollout, (2) the year of rollout on (bilateral) characteristics of languages. Second, I perform a similar analysis using the country pair-level foreign Duolingo exposure country pair level and the year the Duolingo exceeds 0.5. Table 2.4.1 shows the main results. I find unsurprisingly that courses are rolled out between larger languages, and courses to languages with less target speakers are rolled out later. Moreover, the probability a module is rolled out increases strongly if both languages have many speakers. Turning to the country-level analysis, I find that the Duolingo exposure in 2023 is almost 20 percentage points smaller for countries sharing a language and increasing in GDP per capita of both the origin and target languages, suggesting that courses are rolled out to languages spoken by rich countries. After inclusion of origin and destination fixed effects, distance between countries is related to a lower Duolingo exposure. Importantly, the log of the bilateral stock of migrants does not explain Duolingo exposure in 2023, nor does it explain rollout timing once unobservable country-level characteristics are accounted for. However, this analysis does not exclude dynamic selection into treatment. I assess this by estimating pre-trends in various outcomes in subsequent sections.

TABLE 2.4.1: The Determinants of the Rollout of Duolingo Courses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$Duolingo^{ST}_{2022}$		Year of rollout		$DL_{od,2023}$		Year of large rollout	
Log source speakers	0.004**		-0.051					
	(0.002)		(0.265)					
Log target speakers	0.003**		-0.447**					
	(0.002)		(0.175)					
Log source $\times$ Log target speakers		0.002**		0.460				
		(0.001)		(.)				
Sharing an official language					-0.191***	-0.202***	1.848***	0.410
					(0.027)	(0.028)	(0.419)	(0.330)
Log population-weighted distance					0.054***	-0.021***	-0.031	0.048
					(0.009)	(0.006)	(0.215)	(0.046)
Log GDP pc PPP in origin					0.036***		0.239**	
					(0.008)		(0.095)	
Log GDP pc PPP in destination					0.031***		0.373*	
					(0.009)		(0.195)	
Log bilateral migrant stock + 1 (2005)					0.010***	-0.000	-0.213***	-0.004
					(0.002)	(0.002)	(0.048)	(0.012)
Observations	13225	13225	84	52	22217	22217	4098	4098
Source and Target FE		✓		✓				
Origin and Destination FE						✓		✓

Notes: OLS regressions of language-pair and country-pair level exposure and year of rollout on language- and country characteristics. I include all N source and target languages Column 1 and 2 regress a binary module for the presence of a module (by the end of 2023) on the log of source- and target language speakers and its interaction, where column 2 adds fixed effects. Columns 3 and 4 regress the year of rollout on the same characteristics, where column 4 add fixed effects, dropping all source and target languages appearing only once. Similarly, column 5 and 6 regress the measure of Duolingo exposure in 2023 on country and dyadic characteristics, and column 6 adds origin- and destination fixed effects. Column 7 regress the year in which exposure to Duolingo first exceeds 0.5 on country and dyadic characteristics, and column 8 adds origin- and destination fixed effects. Hence, columns 7 and 8 only include dyads for which exposure exceeds 0.5 by the end of 2023. Columns 5-8 also control for log of origin country population, log of destination country population, a dummy for countries sharing a border, linguistic proximity between both country's main languages (Adserà and Pytliková, 2015), log of trade value in 2005, origin country in EU, destination country in EU and both countries in the EU. Data on speakers by language is obtained using Ginsburgh, Melitz and Toubal (2017) and World Bank population data, data on country characteristics, except linguistic proximity, is obtained from CEPII. Linguistic proximity is obtained from Adserà and Pytliková (2015). Standard errors reported in parentheses are two-way clustered: on the source language and target language level (1-4) or on the country of origin and country of destination level (5-8). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.4.4 Estimation

To estimate treatment effects from multiplicative models of dyadic trade or migration flows, scholars typically estimate gravity models by Pseudo-Poisson Maximum Likelihood (PPML) rather than log-like transformed outcomes for two main reasons.[26] First, the presence of zero flows make the estimates effect size dependent on the unit in which the outcome is stated when there is an effect on the extensive margin (Chen and Roth, 2024). Second, heteroskedasticity can introduce bias in OLS estimates of non-linear models (Silva and Tenreyro, 2006).[27] Moreover, non-linear models are biased due to the incidental parameter problem (IPP). The (panel) Poisson

[26] I borrow the term log-like from Chen and Roth (2024), who define it as "functions $m(y)$ that are well-defined at zero but behave like $log(y)$ for large values of y, in the sense that $m(y)/log(y) \to 1$". This includes the often-used $log(y+1)$ and inverse hyperbolic sine $log(y+\sqrt{1+y^2})$ transformations.

[27] An often overlooked but important difference between the Poisson and OLS of a log-transformed outcome model is that the estimand it targets is different, as noted by Tyazhelnikov and Zhou (2021) and Chen and Roth (2024). Chen and Roth (2024) calls $e^{\hat{\beta}} - 1$ the average proportional treatment effect on the treated. The Poisson estimator targets the relative average effect, whereas the log-like estimator targets the average relative effect. Hence, a given treatment effect has the same effect on the Poisson estimate when it happens to units with low- or high levels of untreated potential outcomes.

estimator is the only non-linear estimator that does not face this issue. However, this is not the case for three-way fixed effects Poisson estimation. Weidner and Zylkin (2021) develop a correction to the prevailing IPP bias. Although this bias is limited in most cases, I report bias-corrected estimates.

Negative weights. A wealth of recent literature has shown that two-way fixed effects regressions of staggered treatments do not identify an estimand of positively-weighted treatment effects (Goodman-Bacon, 2021; de Chaisemartin and d'Haultfoeuille, 2020; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Borusyak, Jaravel and Spiess, 2024; Callaway and Sant'Anna, 2021; Wooldridge, 2021). In extreme cases, the negative weights on some treatment effects may lead the researcher to find results that take opposite signs. Alternative estimators have been proposed for the staggered setting with binary absorbing treatment, as well as with staggered adoption of multi-valued and continuous treatments Callaway, Goodman-Bacon and Sant'Anna (2021) as well as fully continuously distributed treatments (de Chaisemartin et al., 2024). Wooldridge (2023) pointed out that the same issue arises in non-linear models and Strezhnev (2023) and Nagengast and Yotov (2023) have shown that this problem naturally extends to the three-way fixed effects setting. In this paper, I estimate event study specifications using the regression-based estimator developed by Wooldridge (2023) and implemented by (Yotov, Nagengast and Rios-Avila, 2024). This estimator prevents "forbidden comparisons" by estimating treatment effects heterogeneity by separately estimating treatment effects by cohort and time period. These treatment effects can be aggregated by event time, enabling the study of pre-trends and dynamic effects. A major advantage of this estimator is that it can be implemented for non-linear models and that it flexibly allows for the inclusion of additional fixed effects. A limitation of this approach is that its usage is limited to binary and absorbing treatments. For implementation, I study large increases in Duolingo exposure. As a large share of variation is driven by a single course rollout per dyad, this allows me to estimate event studies using most identifying variation.

Multilateral resistance to migration. An additional estimation concern for studying migration-related outcomes arises from the strong assumptions underlying the random utility of migration, including that the discrete choice problem fulfills the Independence of Irrelevant Alternatives (IIA) assumption. However, in reality, individual-level preference shocks for different destinations are not independent. Bertoli and Fernández-Huertas Moraga (2013) show that this generates additional terms in the error term, giving rise to an endogeneity problem in equation 2.9 and could bias estimates if the independent variable of interest is correlated across destinations. In limiting cases, this term does not vary by dyad over time. One such case, as Ortega and Peri (2013) and Bertoli and Fernández-Huertas Moraga (2013) discuss, assumes that correlations of the EVT-1 shock only happens within two nests: one for the origin country, and one for all foreign destinations. The resulting multilateral resistance term only varies on the origin country by year, which can be accounted for by the inclusion of origin-time fixed effects.

More generally, for a given origin country, the correlation structure of the EVT-shock across destinations may be complicated, giving rise to multilateral resistance terms varying at the origin-destination-year level. Relevant to my setting, destination-specific shocks may be correlated across destinations sharing languages. To see how this can effect my estimates, consider the following example. Duolingo availability from a given language to Spanish increases the attractiveness of all Spanish speaking destinations. However, the Spanish speaking countries are also close substitutes and preference draws for these destinations are correlated. Hence, the effect size is underestimated because the competing effect all these destinations exert on each other makes the migration rate increase less than if only one destination would have been treated in isolation. This would generate a bias towards zero in my estimates. As origin-time fixed effects may not account for all resistance terms, an additional remedy is to include origin-time-destination nest fixed effects, where nests are chosen based on sharing relevant observable characteristics of destinations that determine the attractiveness to migrants (Bertoli and Fernández-Huertas Moraga, 2013; Beine, Bierlaire and Docquier, 2021). I address this in section 2.6.2.

Inference. Standard heteroskedasticity-robust standard errors may overestimate the precision of regression estimates because of correlation of the error term across observations. Conventional knowledge is to cluster at the level of treatment assignment (Abadie et al., 2023). In dyadic data, however, all flows from o to d can be correlated to all dyads where either o or d are represented. In practice, clustering two-way at both the sending and receiving unit gives approximately correct standard errors Cameron and Miller (2014). However, geographically, economically or culturally close countries often speak the same language or languages from the same language family. Due to this, observations across origin and destination countries sharing languages may not be fully independent. As the treatment is correlated across origins sharing languages to a given destination and vice versa, this could lead to an underestimate of the variance of the estimates. To account for this, I consider alternative two-way clustered standard errors: on the main spoken language in the origin country and on the main spoken language in the destination country.

2.5 Language Learning

In this section, I study the impact of Duolingo course rollout on language learning and learning outcomes. First, I study the determinants of take-up of Duolingo courses and its effect on traditional language learning. Second, I examine to what extent the introduction of a course induces online search behavior in Duolingo and towards the target language of the rolled out course by origin country. Third, I study whether the language skills of English test takers is impacted by Duolingo modules to English. Fourth and last, I study how adult and in-school language learning intensity is affected by Duolingo course availability.

2.5.1 Language learning

To study the determinants of take-up, I rely on the number of learners by language course, which I obtained from the Duolingo website in October 2022. Because of users quitting Duolingo, the total number of ever learners likely exceeds the number of learners at a given point in time. Figure 2.B.8 and 2.B.9 show the total number of learners by language and as a share of the number of speakers by language, aggregated by source and target language. Figure 2.B.8 shows that English is by far the most used source language as well as the most learnt target language. The former suggests that English is also used by many non-native English speakers to learn third languages. The latter reflects that returns to English proficiency are high due to its status as a *lingua franca*. Even among other widely spoken languages, such as Spanish and French, the number of learners exceeds 10% of the number of speakers. Moreover, several more widely spoken languages have garnered a relatively large number of speakers likely for tourism and cultural reasons, such as Greek, Italian, Japanese, and Korean. Table 2.B.1 shows the results from a regression of the number of users by course on the number of speakers of the source and target languages. and its interaction. The results show that the number of users increases by about 8 for every 1000 source language speakers as well as for every 1000 target language speakers.[28] This shows that not only the total pool of potential learners is relevant in the decision to learn a language, but also the extent of applicability of the target language. The positive interaction effect

[28] This is the slope of users to speakers. The percentage of learners of total source language speakers is about 2%. As individual can speak multiple languages, this number is lower than the total share of world population that have used Duolingo.

between the number of source and target language speakers in column 2 and 3 also suggest that demand for courses is particularly high between languages with many speakers, such as English and Spanish.

2.5.2 Interest in Duolingo and available languages

Online search behavior is a useful proxy of the interest in a topic. To validate that the course rollout dates and to study how courses affect demand for language-related information, I turn to online search interest for Duolingo and available languages. I collect information about relative search intensity on the widely used search engine *Google* through the Google Trends API. This API enables one to query time series of normalized search intensity of a search term or topic over time, referred to as the Google Trends Index (GTI). To ensure that the GTI represent the same relative search intensities across countries, I always scale the GTI across regions. For a detailed description about how consistent time series are constructed, Google Trends is queried for search terms and topics, see 2.C.

Interest in Duolingo. As Duolingo was founded in 2011, global search interest went from practically zero to large values (see Figure 2.C.1a). To study the effect of course rollout on the search interest for Duolingo, I study the scaled GTI four years before and after first increases of 50 percentage points or larger in foreign Duolingo exposure on the origin country-level.

Figure 2.5.1a shows the development of interest in Duolingo four years before and after the introduction of the first relevant Duolingo course on the country of origin level, controlling for country of origin and year fixed effects. After course introduction the country-specific interest starts increasing gradually, which is in line with the increasing popularity of Duolingo over time documented in Figure 2.3.1. As in some cases Duolingo was available as a beta version before the course was released, there was some interest several quarters before courses became available. This could also be driven by foreign language speakers (migrants or visitors) that search for Duolingo as a course may have already been available to them.

FIGURE 2.5.1: The effect of influential course introductions on interest in Duolingo and Languages

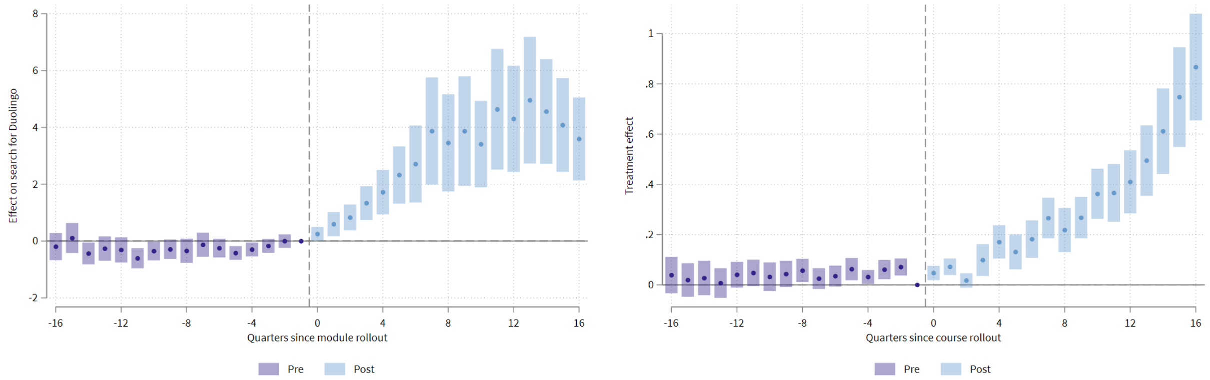

(A) Duolingo

(B) Languages

Notes: Results from Wooldridge (2023) event study estimators around increases in Duolingo exposure. (A) shows estimates of the effect of increases in Duolingo exposure on the relative search interest in Duolingo and its transliterations across countries proxied be the scaled Google Trends Index. A country is counted as treated if the share of inhabitants speaking a source language with a Duolingo course is at least 50 percentage points. The two-way fixed effect counterpart of the Nagengast and Yotov (2023) event study corresponds to a regression of (a) the GTI in Duolingo on origin country and quarter fixed effects and an aggregation of the Duolingo exposure: $GTI_{ot}^{Duolingo} = \beta \mathbb{1}\left(\left(\max_{S,T} \alpha_o^S DL_t^{ST}\right) > 0.5\right) + \psi_o + \phi_t + \epsilon_{ot}$. N = 8,058 from 158 countries, of which 112 treated. Shaded blue bars indicate 95% confidence intervals based on cluster-robust standard errors at the country level. (B) shows estimates of the effect of increases in Duolingo exposure on relative search interest in available target languages proxied be the Google Trends Index scaled across countries but not languages. The regression counterpart is a three-way fixed of the GTI on origin-language, origin-quarter and target-quarter fixed effects: $GTI_{ot}^{\tilde{T}} = \beta \mathbb{1}\left(\left(\max_{S} \alpha_o^S DL_t^{ST}\right) > 0.5\right) + \psi_{oT} + \phi_{tot} + \theta_{Tt} + \epsilon_{oTt}$. In the latter case, I restrict the time period to 2011-2022, as Google Trends Indices for less frequently spoken languages are noisy due to the limited search interest. N = 545,220, 15,145 pairs of which 1,279 treated from 233 origins and 65 queried target languages. Shaded blue bars indicate 95% confidence intervals based on two-way cluster-robust standard errors at the country and language level. Data obtained by repeatedly querying Google Trends. For a further discussion on the construction of the Google Trends Indices, see section 2.C.

Interest in target languages. As Duolingo courses enable to learn a particular language, it could also spur the interest in the specific language or increase the intensity of learning. Relative search interest in languages was stable since 2009, but started increasing in 2016 (see Figure 2.C.1b). As for Duolingo, I study the interest in languages four years before and after course roll-out. Contrary to the previous section, as I obtain dyadic search interest from origin countries to target languages, I can partial out all origin-time and source language-time variation with fixed effects. I study the scaled GTI around first increases of 0.5 or larger in dyadic Duolingo exposure on the origin country-source language level.

Figure 2.5.1b shows the event study results around the introduction of a salient course on the *bilateral* interest between the origin country and the target language. I find that prior to course roll-out, interest in languages is not trending before course rollout and that interest in the language slowly starts increasing after the introduction of the course and increases steadily thereafter. The continuing gradual increase over time reflects that Duolingo has become considerable more popular over time, but also reflect that course availability increases online search to further study languages by e.g. searching for translations or further study material. These two exercises also validate that the rollout dates, based on the date courses enter the final phase, reasonably capture the relevant timing of course introduction.

2.5.3 Language skills

The previous sections have shown that many people took up Duolingo courses, and that course roll-out spurred interest in Duolingo and available target languages. A pressing question is whether access to low-cost language learning also improved language skills among the general population, as predicted by the model. However, internationally comparable data on foreign language skills among the general population is scarce. As an alternative to representative data, I consider origin-specific test scores from widely-used language proficiency tests.

TOEFL. The Test of English as a Foreign Language (TOEFL) test organized by ETS, which is an English language test taken annually by more than 2 million individuals across the world, mostly for (foreign) university enrollment. Hence, the base of test takers likely includes many prospective migrants. The TOEFL test is scored on a scale of 0 to 120 points, where each of the four sections (Reading, Listening, Speaking, and Writing) receives a score from 0 to 30. However, test participants may be differentially selected after introduction of low cost language technologies. For example, individuals without prior access to language learning may take up low-cost language learning, but still have worse skills at the time of the test than the rest of the population. Moreover, individuals could defer the test to improve their test scores as they can continue learning. The latter concern is somewhat alleviated as many test takers take it at the end of high school or a degree program, which leaves little margin to defer the test. Furthermore, differential selection is unlikely to affect the relative performance across elements. In particular, as Duolingo predominantly enables the learning of passive language skills, I can test whether scores on passive (reading and listening) rather than active elements (writing and speaking), providing a partial test of the influence of low-cost language learning on language skills.

ETS compiles yearly reports of average section scores by test takers' native language. Figure 2.5.2 shows the results of event studies of test scores around the introduction of 22 Duolingo courses to English. In the first years before course rollout there are no discernible pretrends. In the first five years after course rollout, scores have increased by 0.12 s.d. for reading and 0.08 s.d. for listening, but no effects for writing and speaking. Altogether, the results suggests that passive skills have improved considerably, without an effect on active skills. Based on the nature of Duolingo vis-a-vis in-person language courses with more active components, this is not surprising. This exercise does not identify the effect of Duolingo on the average population for two important reasons. First, TOEFL takers are much more likely to have used Duolingo. Second, Duolingo availability may affect selection into taking the test. As Duolingo enables the study of English language at beginner and intermediate levels, it seems unlikely only initally more proficient individuals take the test. Nevertheless, test scores have become more positive on average, suggesting this effect is not dominating. As Duolingo may be more useful to those with low levels of initial skills, I study whether the effects are stronger by native languages with low pre-Duolingo test scores. Figure 2.D.1 shows effects separately for the subsample of languages with below-median scores in 2010. Effects are considerable larger, and the results also hint at small positive effects for active language skills.

FIGURE 2.5.2: The Effect of Duolingo Rollout on Component Scores of English language (TOEFL) test (2007-2022)

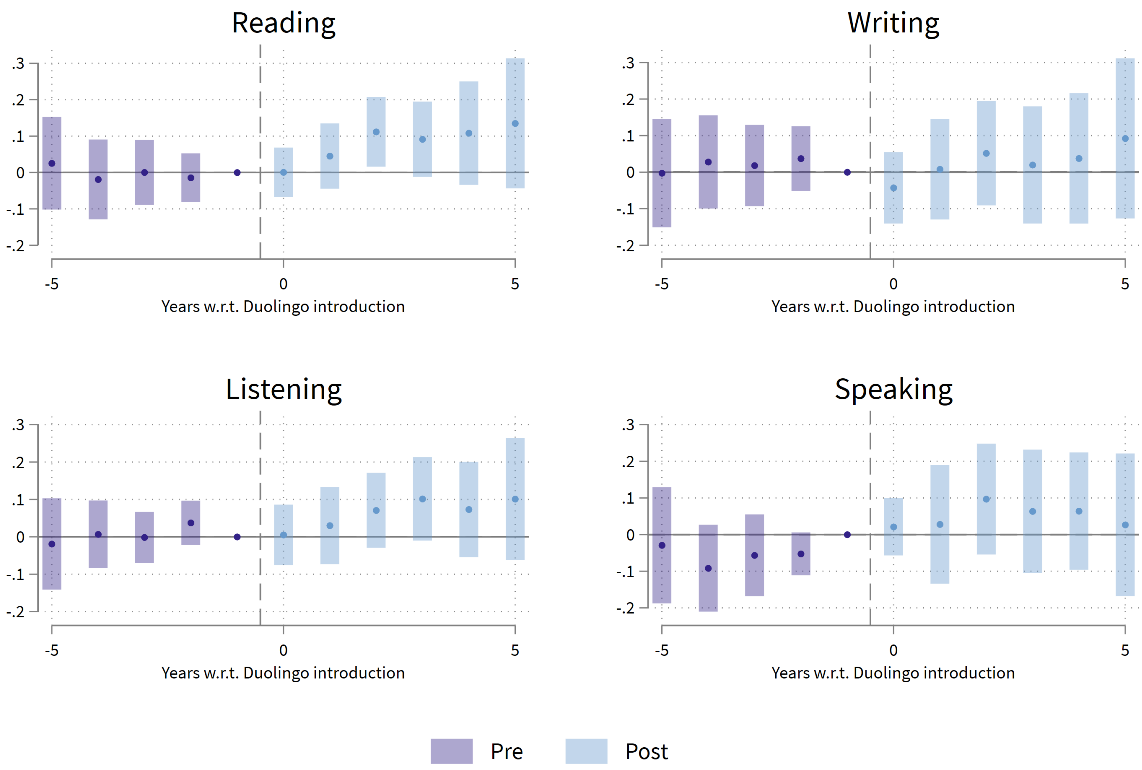

Notes: Results from linear Wooldridge (2023) event study estimators around introduction of a Duolingo course to English. The panels report results for four different outcomes: standardized TOEFL scores reading (upper left), writing (upper right), listening (lower left) and speaking (lower right) by native language of the test takers. As the unit of observation is the native language level, the Duolingo exposure is binary. N = 1,808 from 121 languages, of which 22 are treated. Shaded blue bars indicate 95% confidence intervals based on standard errors clustered at native language level. Data is obtained from the yearly TOEFL iBT Test and Score Data Summary between 2007 and 2022.

GRE. To further address the potential risk posed by differential selection into test taking, I turn to the GRE test. The GRE is a English-languages general ability test taken by approximately 300,000 individuals every year. Most test takers use the GRE test for university admission in a foreign country (not necessarily English-speaking). Contrary to the TOEFL test, for the GRE the number of GRE test takers are reported, which enables testing whether Duolingo availability has changed the size of the pool of test takers. Moreover, the GRE includes separate linguistic and quantitative elements, which allows us to test whether selection in terms of general skills has changed. As Duolingo exposure should not have a causal effect on quantitative skills, this provides a test of changes in selection of test takers. As the GRE scores are published by origin country, I construct a measure of exposure on the origin country-level for Duolingo courses to English. Figure 2.D.2 reports results from an event study specification using the Wooldridge (2023) estimator. The number of test takers has not changed in a statistically significant way. Verbal and Analytical writing skills, however, have increased by up to 0.2 S.D. after three years. Reassuringly, I find null effects on quantitative scores, which improve confidence that the availability of Duolingo did not affect selection into the GRE based on other skills.

2.5.4 Traditional Language Learning

It is a priori unclear how low-cost language learning interacts with traditional (in-class) language learning. On the one hand, potential learners may use Duolingo instead of traditional in-class language courses. On the other hand, Duolingo may spur language learning at basic levels and generate interest in destination language countries and culture, which could increase in-class learning, particularly at higher proficiency levels. This could increase the total number of learners, and increase average proficiency, as suggested by the model of section 2.A. I test which effect prevails for two different types of language learning: adult language learning in German language learning institutes and in-school foreign language instruction across the EU.

Goethe institutes offer German language learning in more than 90 countries worldwide. These institutes both offer languages courses, as well as exams providing generally accepted certification. In 2.D.2 I present suggestive evidence that the rollout of 9 Duolingo courses to German decreased course participation in traditional German language courses, but had no effect or even a slight positive effect on the number of exams taken. This is in line with online language learning substituting for costly learning, but not for certification, which may be needed for visa or employment.

In most countries in the Eurppean Union, schools instruct pupils in one or more foreign languages. Data on the share of pupils learning any given foreign language by education level is available through *Eurostat*. In 2.D.3, I study how Duolingo course availability has impacted in-school language learning. Reassuringly, before Duolingo course availability, there are no discernible pre-trends in the share of students learning available foreign languages. Moreover, I find that across school levels, low-cost language learning increased the share of students by 1–2 percentage points up to five years after course introduction. The effect size is increasing over time, which is in line with gradual adoption of Duolingo courses over time. Excluding the most learnt foreign language across the EU, English, the effects are still positive but smaller in magnitude, and marginally significant in most cases. These results also imply that low-cost language learning likely increases the total online and offline learning intensity of available languages, improving foreign language skills.

2.6 Migration Aspirations and Flows

In this section, I study whether the staggered introduction of low-cost language learning has impacted migration patterns in accordance with the model predictions of section 2.4.1. As comprehensive bilateral migration flow data is only available on a yearly level for a limited number of countries, this section will mostly rely on migration intentions as elicited in the Gallup World Poll (GWP). To complement evidence the with evidence on migration flows, I additionally study migration flows to OECD countries and global scholarly migration flows.

2.6.1 Data

I use the 2007–2022 vintages of the Gallup World Poll (GWP), which is a representative survey of about 1,000 individuals per year in more than 150 countries. Besides many questions concerning demographic, economic and social issues, it includes a question on whether one would like to emigrate if one had the opportunity, as in Chapter 1 of this thesis. The question's wording is *Ideally, if you had the opportunity, would you like to move permanently to another country, or would you prefer to continue living in this country?*, and the answer options are **yes**,**no**, **don't know** and

prefer not to answer.[29] In the latter two cases I discard the observations. If individuals mention that they desire to emigrate, they are also asked where they would like to migrate.[30] Importantly, bilateral migration intentions as elicited in the GWP are strongly correlated to migration flows across migration corridors and are predictive of subsequent migration flows (Tjaden, Auer and Laczko, 2019). Based on the answers to these questions, I construct the stock of people in country o aspiring to emigrate to country d from the origin country's population at t, the total number of GWP respondents N_{ot} and the share of respondents aspiring to emigrate from o to d, N_{odt}: $M_{odt} = pop_{ot} N_{odt} / N_{ot}$.[31] To construct the odds ratio of migrating to d over staying in o, I calculate $\frac{M_{odt}}{M_{oot}} = N_{odt} / (N_{ot} - \sum_{d'} N_{odt})$. I weight this procedure using the individual-level sampling weights provided by Gallup. The GWP comprises 2,003,036 interviews from 166 jurisdictions of whom 22.8% desire to emigrate.[32] Jurisdictions are visited on average 11.5 times across the 16-year period. 5% of those aspiring to migrate do not indicate a preferred country or indicate a region (e.g. "African country") or jurisdictions without information on languages spoken. I regard those individuals as stayers, to ensure that the sum of migrants and stayers used for calculation in the equals unity.

I complement the data on migration intentions with actual migrant flow data from two different sources. The OECD International Migration database records yearly bilateral migration from virtually all countries in the world to 37 OECD countries. This data consists of collected national statistics about inflows of migrants, in some countries administered by nationality, in others by country of birth. I construct the odds of migration by using the origin-country population from the World Bank. I focus on the time period from 2007 to 2019, due to the large influence of the Covid-19 pandemic on international mobility barriers. A particularly mobile population are academic scholars. Akbaritabar, Theile and Zagheni (2024) constructed a dataset of yearly scholar migration between all countries in the world over time from the OpenAlex (version 2024_v1) database of published scholarly articles (Priem, Piwowar and Orr, 2022). In this dataset, a scholar is counted as a migration from o to d in t if her main affiliation on a paper published in t is an institute in d, and the last available main affiliation prior is from an institution in o. I construct the odds of scholarly migration using the total number of scholars publishing in o at t, also provided by Akbaritabar, Theile and Zagheni (2024). I use the 2007 – 2019 data to study the effect of exposure to Duolingo on scholar migration.

To complete the dataset with additional information on country- and country pair-level, I use the database by Conte, Cotterlaz and Mayer (2022). This dataset includes all important variables to estimate gravity models up to and including 2022: trade flows, trade agreements, geographical distances, macroeconomic indicators, from a variety of original sources.

2.6.2 Aspirations

Three-way Fixed Effects. Table 2.6.1 shows the main estimation results of the model in Equation 2.9. As hypothesized by the model, foreign exposure to Duolingo increases migration intentions strongly across specifications. Inclusion of origin-year fixed effects increases the point estimate, which could be driven by the downward bias exercised by multilateral resistance as discussed in section 2.4.4. In line with the model, domestic exposure to Duolingo is negatively correlated to the log odds of migration, although the estimates are on the brink of significance.

[29]Figure 2.D.5 shows that the share of world population that aspires to migrate has increased since 2010. During the Covid-19 pandemic migration aspirations by 3–4 percentage points, but reverted to previous trends soon after. Chapter 1 of this thesis finds that at least 2–3 percentage points of this increase can be attributed to the rollout of mobile internet networks.

[30]The question's working is: *To which country would you like to move?* to which respondents can give an answer which is codified to a country by the interviewer if possible.

[31]I obtain yearly data on population pop_{ot} from the World Bank.

[32]Several jurisdictions in the GWP are absent in the dataset of Ginsburgh, Melitz and Toubal (2017) nor is there reliable information on languages spoken from the CIA World Factbook. These are omitted from analysis.

The inclusion of controls on the pair level over time barely influences the estimates. Because of the unavailability of data for one origin and eight destinations, I subsequently discuss results without the control variables.

TABLE 2.6.1: The Effect of Duolingo Courses on Bilateral Migration Aspirations (2007 – 2022)

	(1) $\frac{M_{odt}}{M_{oot}}$	(2) $\frac{M_{odt}}{M_{oot}}$	(3) $\frac{M_{odt}}{M_{oot}}$	(4) $\frac{M_{odt}}{M_{oot}}$
DL_{odt}	0.267***	0.374***	0.306***	0.352***
	(0.065)	(0.080)	(0.062)	(0.068)
DL_{oot}	-0.225		-0.269*	
	(0.160)		(0.140)	
Observations	123263	123263	98019	98019
Unique origin countries	153	153	152	152
Unique destination countries	196	196	188	188
Unique dyads	9439	9439	9439	9439
Origin-destination FE	✓	✓	✓	✓
Origin-year FE		✓		✓
Destination-year FE	✓	✓	✓	✓
Controls			✓	✓
Weidner-Zylkin correction		0.386***		0.360***
		(0.060)		(0.058)

Notes: PPML estimation of a Gravity model of migration without (odd columns) and with origin-year fixed effects (even columns). The dependent variable is the ratio of the total number of people desiring to emigrate from origin country o to destination country d in year t over the total number of people not desiring to emigrate from country o in year t. Trade controls include a dummy for joint EU membership, a dummy for a WTO trade agreement between two origin and destination country, as well as the log of trade flows from the origin to the destination country. Because data on trade flows is not available for all destinations, columns 3 and 4 have less observation than columns 1 and 2. The bottom row reports results from the bias correction of Weidner and Zylkin (2021) for three-way fixed effects Poisson regression. Data on migration intentions originates from the Gallup World Poll and the Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

One unit of foreign Duolingo exposure, corresponding to a language course enabling communication between everyone in the origin country to everyone in the destination county, increases the log odds of migration by 45% (column 2). In comparison, one unit of domestic Duolingo exposure *decreases* the log odds of migration by 25% (column 1). In line with the model, this suggests that the introduction of low-cost language learning courses to languages rewarded on foreign labor markets increase migration intentions, whereas the introduction of courses to languages that are rewarded on domestic labor markets decrease intentions. The net effect of the availability of low-cost language learning depends on the relative return to language skills at home and abroad.

One may be wondering whether courses with higher returns have stronger effects. To study this, I subdivide treatment intensity in four bins of exposure and include indicators for these bins in equation 2.9 instead of a continuous variable for both foreign and domestic exposure. Table 2.D.4 presents the results. I find that the effect for foreign exposure is close to zero for small levels of exposure, but increases monotonically with higher treatment intensity. The effect for domestic exposure is driven by those with a very large exposure. These are countries where two languages are spoken by a large proportion of the population. In such cases, knowledge of both languages may enable one to earn a large premium on the labor market.

Table 2.D.5 shows the main results using the Duolingo exposure calculated using official languages at the destination. The effects are statistically significant, but smaller. This could be driven by the strong effects for non-official but widely spoken languages among prospective migrants, such as English.

At first glance, the effect sizes may seem very large. However, one has to consider that these are relative average effects on the odds ratio of migration. About 20% of individuals in GWP desire to emigrate, and as there are about 200 alternative destination countries, the bilateral stock of aspiring migrants is about 0.1% on average. A 45% increase implies that the stock of aspiring migrants across a dyad receiving full treatment increases by about 0.045 percentage points. As average foreign Duolingo exposure is 0.26 in 2022 and there are 196 potential destinations in the main estimation sample, this suggest a persuasion rate of about 2.3% of the population.[33] This simple back-of-the-envelope calculation is an upper bound of the share of people that changed their migration intention due to Duolingo course availability, for two reasons. First, as most exposure intensities are considerably lower than unity, the average per-unit effect is closer to 37.4%. Second, it disregards multilateral resistance effects. For example, if a language course to English becomes available, this increases migration aspirations to all English-speaking destinations. However, it increases migration aspirations with less than the point estimate to a given English-speaking destination as other English-speaking countries exert a downward influence, as language-sharing destinations likely are close substitutes.[34]

Importantly, these results do not imply that total emigration intentions on the origin-level would increase by 45%. A positive point estimate in a gravity model could imply (a combination of) two effects. First, the availability of language learning may have shifted the preferred destination of aspiring migrants, or induced migration intentions among those who would not have desired to migrate in absence of a course. To study which effect prevails, I collapse the data at the origin-year level and regress the emigration rate on the domestic and average foreign Duolingo exposure and country and year fixed effects. Table 2.D.3 reports the results. Although statistically insignificant, the effects suggests that moving from no to full foreign Duolingo exposure increases the emigration rate by about 4 percentage points. Columns 3 and 4 report results using a migration stock-weighted measure of average foreign Duolingo exposure. Latter results are similarly sized, but statistically significant, indicating that courses in more attractive destinations increase the emigration intention rate. If the effect of 45% would be fully driven by newly aspiring migrants, this would explain 10 percentage points. Hence, less than half of the increase in migration aspiration odds can be explained by new attraction, and the rest by diversion effects.

Event study. To assess the plausibility of the parallel trends assumption by considering pre-trends, to study the dynamic effect and to alleviate the concerns that my results are driven by negative weights in staggered difference-in-differences settings (Goodman-Bacon, 2021), I estimate an event study estimator for large increases in foreign Duolingo exposure across a dyad. As the Nagengast and Yotov (2023) estimator requires a binary treatment, I define an event as a dyad receiving an increase of 50 percentage points in foreign Duolingo exposure measure, using all other dyads as a control group. Figure 2.D.6 shows that the foreign Duolingo exposure increases strongly, and that pre-trends, driven by treatment and control units receiving small amounts of exposure, are negligible.

Figure 2.6.1 shows that the main event study results. I find no evidence of pre-trends before strong increases in exposure, increasing confidence in the conditional parallel trends assumption.

[33] This is a considerable share of the population that used Duolingo. By the end of 2022, more than 500 million people are registered on Duolingo (see e.g. `https://blog.duolingo.com/2022-duolingo-language-report/`), which is about 6.2% of world population. As the GWP does not survey the full world population, the propensity to take up Duolingo

[34] Based on gravity model estimates, one can simulate the effects of a language course on the distribution of migrants taking multilateral resistance effects into account. see e.g. Guichard and Machado (2024).

I find a gradual increase in treatment effect in the first three years after treatment, which dampens thereafter.[35] The slow onset of treatment effects is in line with the gradual adoption of courses.

FIGURE 2.6.1: Event Study around Large Increases of DL_{odt} on Bilateral Migration Aspirations (2007-2022)

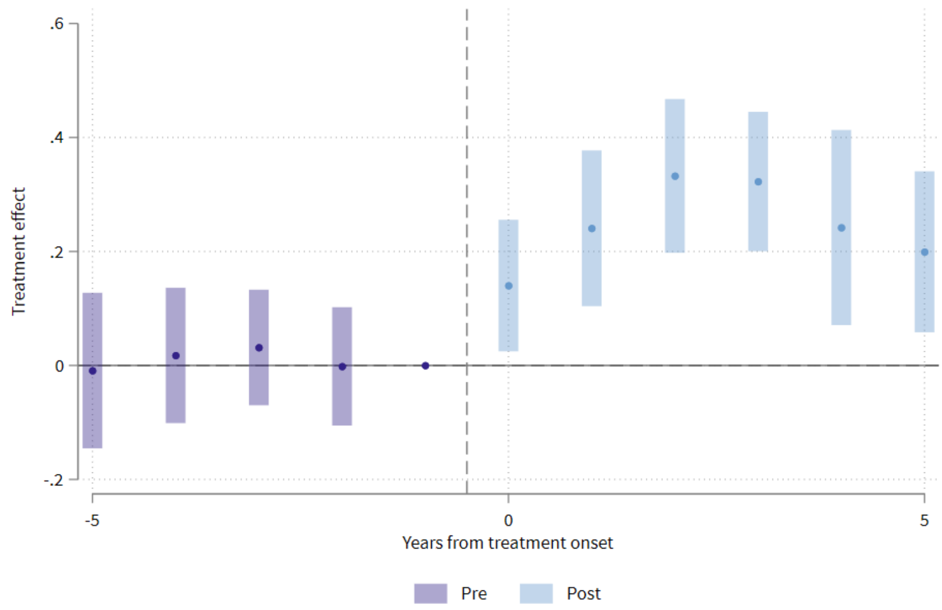

Notes: PPML regressions of the heterogeneity-robust event study estimator by Nagengast and Yotov (2023) of migration aspiration odds on a binary indicator for whether a country pair has experienced increase in foreign Duolingo exposure exceeding 50 percentage points, including origin-destination pair, origin-year and destination year fixed effects. An event is defined as an increase in Duolingo exposure of more than 50 percentage points. 2.D.6 shows that treated units on average experience a stable increase of about 60 percentage relative to the control group. Estimates are shown for the 5 years before and after an event. Shaded blue bars indicate 95% confidence intervals based on standard errors clustered at native language level. See notes to Table 2.6.1 for information on the data and sample.

The Role of English

The dominant role of English is clear from the model: the equilibrium level of language learning $l_{oS}^{T}*$ is an increasing function of the migration probability-weighted returns to the language skill. In the case of English, the domestic contribution to this is large across countries. As shown in Figure 2.B.8, this is reflected in learner numbers. Moreover, English takes a special role as a source language as many non-native speakers can use English as a language of instruction.

To study the special role of English, I split out the contribution of English and all other languages, calculating two exposure measures for both foreign and domestic exposure. I do this separately for English as a source language and a target language. Table 2.6.2 shows the results. Columns 1 and 2 show that the effect of foreign Duolingo exposure is significant for both English and other languages as source languages, but the former is considerably stronger. This could be driven by the fact that many English speakers are non-native speakers themselves, who are higher educated and at the same time have a larger propensity to migrate. Columns 3 and 4 shows that the effect is also strongest for English as a target language. However, it is considerable smaller and insignificant for other languages. This is in line with the model, as can be seen from equation 2.8: expected returns to English are considerably larger than for other languages,

[35] I also perform a similar event study using large increases in domestic Duolingo exposure, controlling for foreign Duolingo exposure. The results, presented in Figure 2.D.7, suggest that domestic exposure decreases migration aspirations gradually in the first three years after exposure.

not in the last place because of returns on the domestic labor market. Hence, uptake of language learning is stronger for English, which may lead to a larger increase in bilateral migration, given the same level of foreign Duolingo exposure.

As the effect for foreign Duolingo exposure is driven by English as a target language, one could be concerned that the effect is driven by one or few English-speaking countries. In Table 2.D.9 I exclude five high-income native English-speaking destination country at a time and all at the same time. Although the point estimates are somewhat smaller when excluding the US and few destination countries, which is not the case.

The results for domestic language exposure in 2.6.2 show a similar pattern: the effects are strongly negative for English as a source- and target-language. This suggests that countries where English is spoken by many people, the opportunity to learn English, or to learn another widely used language using English, strongly decreases migration intentions.

TABLE 2.6.2: The Role of English as a Source and Target language

	(1) $\frac{M_{odt}}{M_{oot}}$	(2) $\frac{M_{odt}}{M_{oot}}$	(3) $\frac{M_{odt}}{M_{oot}}$	(4) $\frac{M_{odt}}{M_{oot}}$
$DL_{odt}^{S=EN}$	0.729***	0.827***		
	(0.155)	(0.186)		
$DL_{odt}^{S\neq EN}$	0.207***	0.251***		
	(0.073)	(0.097)		
$DL_{oot}^{S=EN}$	-0.623***			
	(0.112)			
$DL_{oot}^{S\neq EN}$	0.038			
	(0.172)			
$DL_{odt}^{T=EN}$			0.219***	0.543***
			(0.084)	(0.093)
$DL_{odt}^{T\neq EN}$			0.257***	0.109
			(0.095)	(0.087)
$DL_{oot}^{T=EN}$			-0.705***	
			(0.216)	
$DL_{oot}^{T\neq EN}$			0.129	
			(0.175)	
Observations	123263	123263	123263	123263
Origin-destination	✓	✓	✓	✓
Origin-year FE		✓		✓
Destination-year FE	✓	✓	✓	✓

PPML regressions based on the sample and specification of column 1 and 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Column 1 and 2 report results from models where the Duolingo exposures are calculated separately for courses with- and without English as source language. Column 3 and 4 report results from models where the Duolingo exposures are calculated separately for courses with- and without English as target language. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: See notes to Table 2.6.1

Heterogeneity

Table 2.6.3 shows the interaction of foreign Duolingo exposure with several other variables. As a validation, I confirm that the effect is stronger for those countries searching more intensively for Duolingo. As most users use Duolingo on mobile devices, mobile internet access should increase the availability of Duolingo at the extensive (number of users) intensive (intensity of learning) margin. I interact the foreign Duolingo exposure with the country-year level share of individuals covered by mobile networks, finding that access to mobile networks more than doubles the effect size. However, I find that the country-level number of broadband users is irrelevant to the effect size. Column 3 and 4 show that the effect size is smaller for linguistically closer countries and countries sharing a language. If countries' languages are linguistically close, learning costs were already plausibly lower, so the introduction of online language learning did not decrease costs that much. If countries already share a language, benefits to learning another destination-country language may be very low and other language learning opportunities to learn the shared language may be widespread. Although insignificant, column 5 and 6 shows that the effect is slightly stronger for higher-income origins and considerably stronger for higher-income destination countries, which is in line with larger gains from migration in high-income destination countries. In Column 7 I find that larger pre-existing migration networks reduce the effect. This could be driven by the reduced need for language learning if there is a diaspora in the destination. This implies that the availability of low-cost language learning expands the choice set of aspiring migrants, raising interest to destinations with low pre-existing migrant networks.As Duolingo does not offer certification, learning on Duolingo may be less of a relative cost reduction in language learning costs in countries with language requirements for immigrants. Table 2.D.2 test whether the presence of language requirements for residency moderate the effect of foreign Duolingo exposure. Although the effect is 5% lower to destinations with language requirements, it is not statistically significant.

TABLE 2.6.3: Effect Heterogeneity of Foreign Duolingo Exposure

	(1) $\frac{M_{odt}}{M_{oot}}$	(2) $\frac{M_{odt}}{M_{oot}}$	(3) $\frac{M_{odt}}{M_{oot}}$	(4) $\frac{M_{odt}}{M_{oot}}$	(5) $\frac{M_{odt}}{M_{oot}}$	(6) $\frac{M_{odt}}{M_{oot}}$	(7) $\frac{M_{odt}}{M_{oot}}$
DL_{odt}	0.318***	0.154*	0.451***	0.373***	0.378***	0.212	0.604***
	(0.078)	(0.092)	(0.096)	(0.079)	(0.079)	(0.137)	(0.108)
$DL_{odt} \times GTI_o^{Duolingo}$ (2006-2022) [0,1]	0.729**						
	(0.306)						
$DL_{odt} \times 3G_{oy}$ [0,1]		0.332**					
		(0.152)					
$DL_{odt} \times \text{Broadband}_{\text{oy}}$ [0,100]		-0.000					
		(0.005)					
$DL_{odt} \times \text{AP15 Linguistic proximity}_{\text{od}}$ [0,1]			-0.505**				
			(0.218)				
$DL_{odt} \times \text{Shared official language}_{\text{od}}$ {0,1}				-0.416**			
				(0.170)			
$DL_{odt} \times GDPpc_{oy}$ (1 s.d.)					0.062		
					(0.065)		
$DL_{odt} \times GDPpc_{dy}$ (1 s.d.)						0.153	
						(0.102)	
$DL_{odt} \times \text{Above-median migrant stock}_{\text{od,2005}}$							-0.255**
							(0.123)
Observations	121180	89759	114137	122242	117908	115121	122279

Notes: PPML regressions based on the sample and specification of column 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Every column introduces an interaction effect between foreign Duolingo exposure and a moderator. Due to limited availability of data on the moderator the sample size varies. GTI is obtained from Google Trends, 3G data is obtained from Collins Bartholomew, Broadband subscription data from ITU, Linguistic Proximity from Adserà and Ferrer (2021), a dummy for sharing an official language from Conte, Cotterlaz and Mayer (2022), GDP per capita (PPP) from the World Bank and the stock of migrants in 2005 from the UN International Migrant Stock database. For ease of interpretation and to prevent a spurious correlation with time trends, I standardized the measure of GDP for every year of the data. Using the migration stock in 2005, I calculate the median number of migrants by origin country. Using the median, I construct a binary indicator taking value one if a particular destination housed an above-median amount of migrants. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Robustness

I perform a series of additional robustness tests on the main results of Table 2.6.1. First, I examine four potential risks to identification and thereafter the sensitivity of the main results to changes in the sample.

Identification. Although I argue that course enrollment is plausibly exogenous to trends in migration intentions across language areas, one can use features of Duolingo rollout to construct exogenous instruments for Duolingo exposure. Section 2.D.4 discusses two such instruments: one based on Duolingo's tendency to develop courses for language dyads with many speakers of the source and target languages and one based on the reduced cost of course development when there are other courses developed for both the source and target language. Table 2.D.1 reports the results. Both instruments predict foreign Duolingo exposure well, although the latter is weaker, which increases estimated standard errors considerably. In both cases the effect size is larger than the baseline PPML results. In part, this could be driven by measurement error in the Duolingo exposure. In particular, as Duolingo courses have become more extensive over time and as more people have access to a desktop or smartphone, the actual exposure to low-cost learning may have been gradually increasing over time, which is captured by these instruments linearly increasing over time.

As mentioned in 2.4.4, residual multilateral resistance could be present in the model with three-way fixed effects. If preference shocks for destinations are correlated within a nest or group of countries, multilateral resistance terms are common to that nest (by year and origin country). Hence, Table 2.D.8 shows the results for specifications with origin-year-destination fixed effects for three types of nests: World Bank's 7 global regions, World Bank's 4 country groups, and membership of the EU. As these specifications are demanding as they introduce many fixed effects to the model, I report the total number of FEs estimated in the table. Although the fixed effects based on the World Bank regions reduce the point estimate somewhat, the estimate remains and large across all specifications.

As discussed in section 2.4.4, the Duolingo exposure measures are strongly correlated across countries speaking the same languages, which could lead to an underestimation of the standard error if there is intra-cluster correlation in migration intentions. To allow for arbitrary correlations within origins and within destinations with the same native language, I cluster standard errors on different levels in Table 2.D.6. In particular, in columns 5 and 6 I cluster standard errors by most spoken language in the origin (63 clusters) and most spoken language in the destination (70 clusters). The results are only slightly less significant than the baseline results clustered by origin- and destination country. However, clustering on the most spoken language may be too granular. For example, many former colonies may house many speakers of one of the world languages for which many Duolingo courses are rolled out, although the most spoken language is a native language. As migration patterns in such countries could be correlated to each other and their colonial hosts, this could be a concern. Hence, I modify the condition for most spoken language in the following way: I count an origin (destination) country into the language group of the most spoken language that is a source (target) language in any Duolingo course. This yields considerably fewer clusters: 22 on the origin level and 30 on the destination level. Nevertheless, columns 7 and 8 of Table 2.D.6 show that the standard errors only become slightly larger, confirming prior results.

As mentioned in 2.4.2, if Duolingo modules are developed due to increased demand for language learning, this is most likely driven by the countries with the most source- and target speakers. To test whether the countries with most speakers by language are not driving the results, I perform three exercises. First, I remove the contribution of the source language in the origin country with most speakers, for every source language, from the Duolingo exposure. Second, I omit the contribution of the target language in the destination country with most speakers, for every target language. Third, I do both at the same time. Table 2.D.7 shows the results, which have changed little compared to Table 2.6.1. Hence, it shows that the effects are not (just) driven by the countries for whose markets Duolingo courses are most likely developed.

Sample. The first year in the data is 2007, which means that 6 years of data before the first course introduction is included in the main estimation sample. Although this choice of starting year is driven by the starting date of the GWP interviews, it is still somewhat arbitrary, and one may be concerned that such a long pre-period, including the 2007-2008 financial crisis and its aftermath, could strongly affect the diff-in-diff estimates. Table 2.D.10 shows that it is not the case; using later time periods decreases the point estimate on foreign Duolingo exposure. Moreover, the last column of Table 2.D.10 shows that omission of the Covid-19 period and subsequent years has only a limited effect on the estimate.

To study how sensitive the results are to single Duolingo courses, I recalculate the main measure of foreign Duolingo exposure omitting a single course at a time. Figure 2.D.8 report both the point estimates and p-values from this exercise. Results are mostly insensitive, with the exception of one course: Spanish to English. The coefficient on foreign Duolingo exposure decreases to 0.23, but remains significant (p=0.006). This is not completely surprising, as it is the most popular course on Duolingo with more than 50 million learners (see Figure 2.B.7). To ensure that the estimates are not driven by single origin- or destination countries, I omit one country at a time.

The results, shown in Figure 2.D.9 and 2.D.10, indeed confirm that the effects are not driven by any single country.

2.6.3 Flows to OECD Countries

Previous section has established that low-cost language learning has a strong effect on bilateral migration intentions. However, did these translate into increases in migration flows? Unfortunatally, there is no dataset containing information about yearly gross migration flows between countries. To nevertheless provide a partial answer to this question, I study the impact on yearly bilateral migration flows to OECD countries. In similar vein to 2.6.1, I present event study estimates in Figure 2.6.2 around increases in foreign Duolingo exposure of 50 percentage points or more. The results are not conclusive. Although the instantaneous and first post-treatment effect are positive and close to significant, the first pre-treatment estimator is negative, suggesting that migration flows increased between the second and first year before treatment. If anything, these results would suggest a considerable effect directly upon treatment, which seems implausible given the delayed response for migration intentions in 2.6.2 and the time it take for language learning and migration preparations.

I interpret these results with caution for two additional reasons. First, the uncertainty of the estimates is large in comparison. The standard error of the average effect if 0.12, which is too large to detect a reasonably sized effect. Second, an important limitation of the OECD migration dataset is that it is incomplete on the destination level. This is problematic in a setting where treatment is correlated across same-language destinations which are partially inside and partially outside the dataset. For example, if a course becomes available to the language all same-language destinations are affected and exercise a downward spillover effect on each other. However, because these observations are outside the dataset it is impossible to account for these multilateral resistance effects.

FIGURE 2.6.2: Event study of Migration Odds to OECD Countries around Large Increases in Duolingo_{odt} (2007-2019)

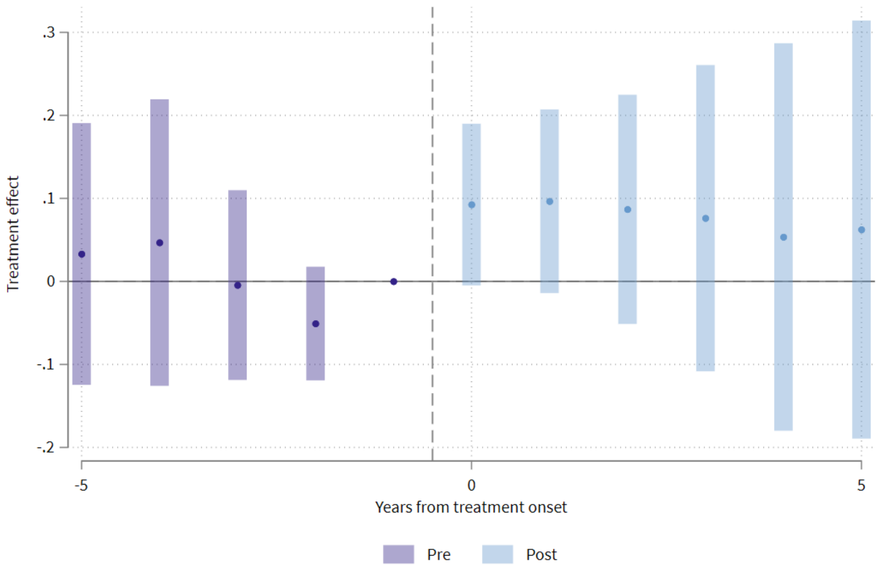

Notes: PPML regressions of the heterogeneity-robust event study estimator by Nagengast and Yotov (2023) of migration flow odds on a binary indicator for whether a country pair has experienced an increase in foreign Duolingo exposure exceeding 50 percentage points, including origin-destination pair, origin-year and destination year fixed effects. Estimates are shown for 5 years before and after an event. Shaded blue bars indicate 95% confidence intervals based on standard errors clustered at native language level. Data on migration flows originates from the OECD and the Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses.

2.6.4 Global Flows of Scholars

To nevertheless estimate the effect on global migration flows, I turn to data on a specific type of migration flows: academic scholars moving between institutions. Table 2.6.4 shows the results from the estimation of a gravity model of scholarly migration odds, net of three-way fixed effects. I find that foreign Duolingo exposure increases bilateral flows by slightly more than 4%, which is just significant at the 5% level. As most scholars in English-speaking countries are publishing in English or are immigrants less proficient in the local language who are more mobile than natives, courses from English may be particularly fruitful. Effects are indeed stronger for courses from English: full foreign Duolingo exposure increases flows with more than 7.5%. Basic levels of language skills could aid high-skilled workers to take up jobs abroad, even though they do not need the local language for their job: For example, it could facilitate other aspects of life including finding housing or interacting with authorities. Partitioning the exposure measure over English as a target language, I do not find any significant effect.

TABLE 2.6.4: The Effect of Duolingo on Scholarly migration flows (2007-2019)

	(1) $\frac{M^{scholar}_{odt}}{M_{oot}}$	(2) $\frac{M^{scholar}_{odt}}{M^{scholar}_{oot}}$	(3) $\frac{M^{scholar}_{odt}}{M^{scholar}_{oot}}$
DL_{odt}	0.042**		
	(0.021)		
$DL^{S=EN}_{odt}$		0.074**	
		(0.032)	
$DL^{S\neq EN}_{odt}$		0.023	
		(0.025)	
$DL^{T=EN}_{odt}$			0.031
			(0.033)
$DL^{T\neq EN}_{odt}$			0.039
			(0.025)
Observations	167670	168024	168024

PPML regressions of the odds of scholar migration based on the sample and specification of column 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Column 2 and 3 use the same procedure as Table 2.6.2, splitting the Duolingo exposure by English and other languages for source- and target languages, respectively. Data on scholarly migration flows originate from Akbaritabar, Theile and Zagheni (2024) and the Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. The estimation sample concerns 189 unique origin and 194 unique destination countries. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.7 Migrants' Language Skills, Selection and integration

One group that has large incentives to take up low-cost language learning are prospective migrants. Greater pre-arrival learning in expectation of a future move most likely improves the language skills of migrants upon arrival. However, the access to low-cost language learning may also particularly facilitate migration for low-skilled individuals who only need basic skills for employment and who had less access to language leaening before arrival. Hence, a priori the expected effect of pre-arrival language learning opportunities is ambiguous. In addition, better language skills before arrival can be used to search more effectively for employment opportunities in the destination. After arrival, migrants can take up low-cost language learning while residing in the host country. Unless this strongly affects selection into return migration, this imaginably improves their language skills. All of these may influence subsequent migrant integration.

To study these effects, I turn to two large-scale survey data sets for two distinct geographic settings between 2007 and 2022: the EU Labor Force Survey (EU LFS) in the EU and American Community Survey (ACS) in the US. Both data sets include information on respondents background characteristics, including country of birth and year of immigration, as well as questions on economic integration. The ACS includes questions on self-assessed language proficiency at the time of interview throughout, but has no question about language skills upon arrival, which the 2021 EU LFS does have.

Both settings have considerable variation in the availability of Duolingo courses over immigrant origin (and destination for the EU) countries over time- English is the most learned second language on Duolingo and the language with the most courses as target language on Duolingo (22).[36] However, it concerns only a single target language for a single destination country, which limits the variation and renders the identifying assumptions weaker. The EU LFS does not face this problem, as it (including the United Kingdom, Switzerland, and Norway) hosts many languages that are target of Duolingo courses (62), as well as several countries that are not. A drawback of the EU LFS is that it does not distinguish the exact origin country of birth of immigrants, but rather one of 24 broad origin country groups.

In the following, I present the main strategies to identify the effect of pre- and post-arrival exposure to Duolingo on migrants' language skills, selection and integration. Because this identification strategy is arguably stronger for the EU, I present the results using the EU LFS in sections 2.7.3 and 2.7.4. Nevertheless, in section 2.7.5 I discuss results for the US using the ACS, which are further presented in detail in section 2.D.6, and compare the findings.

2.7.1 Empirical strategies

Upon arrival. To study the effect of low-cost language learning on migrants' characteristics and language skills upon arrival, I estimating the following specification, which uses the staggered introduction of Duolingo courses:

$$y_{iodct} = \beta DL_{oc}^{T_d} + \phi_{od} + (\psi_{oc}) + \theta_{dc} + \xi_t + \epsilon_{iodct} \tag{2.10}$$

Here, y_{iodct} denotes the outcome of individual i who migrated from origin country o to destination country d in year c interviewed at year t. In some cases, this outcome is realized at year c and faithfully reported in t. An example is the language skills upon arrival. In other cases, I only observe the contemporaneous value of y. For these, I rely only on those interviewed within the first year after arrival, to ensure that outcomes are as close to their values upon arrival as possible. The main independent variable of interest is the Duolingo exposure $DL_{oc}^{T_d}$ to the native language of d. ϕ_{od} are dyadic fixed effects, which capture origin-destination pair specific immigrant characteristics. For example, this controls for the linguistic distance between main languages in countries, which are an important determinant of migrant integration. ψ_{oc} and θ_{dc} are origin-cohort and destination-cohort fixed effects, capturing unobserved heterogeneity in origin-specific and destination-specific characteristics of migration cohorts, such as selection into emigration and destination-country immigration policies. ξ_t captures year-specific factors.

This constitutes a triple differences in cohort time with a fully continuous treatment. For the US, there is only one destination country d, which removes the possibility of including origin-cohort FE (and I include origin FE instead) and destination-cohort FE simply become cohort FE. In this case, it becomes a simple two-way fixed effects specification. To nevertheless account for origin-specific cohort quality, I include a comprehensive vector of time-varying controls to eliminate differences in cohort quality explained by observable factors. To identify β as the causal effect of the availability of low-cost language learning, the outcomes need to fulfill parallel trends in levels for all levels of treatment intensity. The plausibility of this assumption can be examined by considering whether there are pre-trends in outcomes between strongly treated and untreated units.

As in section 2.6, I use the foreign exposure to Duolingo, proxying for the returns of a language skills. However, as the surveys ask about language skills in the main native question, here I construct the Duolingo exposure $DL_{oc}^{T_d}$ as the probability that a Duolingo course enables

[36] As English is the second most studied language in the USA on Duolingo and in 2017 between 2 and 6% of inhabitants across the 50 states used Duolingo, it is plausible that many immigrants use it to improve their English skills after arrival. https://blog.duolingo.com/the-united-states-of-languages-an-analysis-of-duolingo-usage-state-by-state/

communication between a random person in the origin country and a random person in the destination country who speaks the native language. For the EU, the origin indicator o is replaced by the origin group indicator o_g and I calculate the aggregated exposure at the origin group level, weighting with origin-specific cohort sizes using the yearly bilateral flow data introduced in section 2.6.3:

$$DL_{o_g c}^{T_d} = \frac{1}{\sum_{o \in o_g} N_{odc}} \sum_{o \in o_g} N_{odc} DL_{oc}^{T_d} \tag{2.11}$$

Here, $N_{o_d c}$ are the gross flow of immigrants from o to d arriving in c. For some origin regions the weighting is not restrictive because a single language dominates (such as North Africa or Latin America) or because migrants mostly originate from one country in a global region (e.g. Vietnamese in Czech Republic). However, for others this introduces measurement error. Section 2.D.5 discusses this further and provides an estimate of the degree of attenuation bias introduced due to the aggregation: it is about 25%.

After arrival. After an immigrant has arrived in the host country she can continue learning on Duolingo or start a Duolingo course, if a relevant course is available. To isolate this post-arrival effect from that of pre-arrival exposure to Duolingo, I sketches the identification problem in Figure 2.D.11. For sake of simplicity, I consider a binary treatment, but the logic also extends to settings with continuous treatment intensity. Each panel of Figure 2.D.11 shows the availability of a relevant low-cost language course before and after arrival based on the time of arrival and time of interview relative to course rollout. I illustrate the exposure regimes someone falls into (upper) and the resulting pre- and post-migration exposure (lower panel). I do this in two cases: one where the migrant is interviewed within the first year upon arrival (left) and in the second year after arrival (right). In the former case, those arriving before rollout had no access to the course and those arriving after rollout had access to the course course before moving, but as they arrived recently had no time yet to study after arrival. In the latter case, those arriving at least two years before course rollout had no opportunity to use Duolingo, whereas those arriving closer to the rollout date have a gradually longer time window to have used the course after arrival. This provides variation in pre-arrival To capture this, I calculate the post-arrival exposure as the average exposure to Duolingo since arrival:

$$DL_{otc}^{T_d,post} = \frac{1}{t-c} \sum_{\tau=1}^{t-c} DL_{o(t+1-\tau)}^{T_d} \tag{2.12}$$

Here, I denote time since arrival in the destination by $t-c$. $DL_{o(t-\tau)}^{T_d}$ is the Duolingo exposure from o to d in year $(t-\tau)$. For those interviewed during the year of arrival ($t = c$), I set $DL_{otc}^{T_d,post}$ to 0, as they had limited time to take up the course after arrival. To estimate the effect of pre- and post-arrival exposure jointly, I estimate the following model:

$$y_{iodct} = \sum_i \beta_i DL_{oc}^{T_d} \times \mathbb{1}(t-c=i) + \gamma DL_{otc}^{T_d,post} + \phi_{od(t-c)} + (\psi_{oc}) + \theta_{dc} + \xi_t + \epsilon_{iodct} \tag{2.13}$$

Here, notation follows that of equation 2.10. Compared to equation 2.10, I flexibly estimate the effect of pre-arrival exposure on outcomes for each number of completed years since arrival. This is important, as an initial effect of pre-arrival learning may diminish over time relative to the unexposed immigrants, who may catch-up. As pre-arrival and post-arrival exposure are correlated, flexibly capturing dynamic effects of pre-arrival exposure is also important to ensure that the post-arrival exposure does not spuriously capture pre-arrival effects. Additionally, I include pair-by-time since arrival fixed effects $\phi_{od(t-c)}$. These control for general and origin-by-destination-specific integration patterns over time, such as cultural and linguistic distance, and selection into return migration. To make the sample representative of the EU and US, I weight all regressions using the representative yearly weights provided by the EU LFS and ACS, respectively.

To interpret γ as the causal effect of post-arrival exposure on immigrant characteristics and outcomes, stronger exposed immigrants should have had similar integration patterns than weaker or unexposed immigrants in absence of post-arrival Duolingo exposure, conditional on pre-arrival exposure. Furthermore, to pin down the causal effect of learning on outcomes I also need to assume that post-arrival exposure does not predict. This could be violated if selection into migration changes due to immigrants anticipating future increases in Duolingo exposure (i.e. due to the future rollout of a course). The latter is particularly unlikely for those migrants who had no access to Duolingo before arrival. To isolate the effect of post-treatment exposure in absence of pre-arrival exposure, I can estimate equation 2.13 on the subsample with $DL_{oc}^{T_d} = 0$ (i.e. among those who were completely unexposed upon arrival). However, the effect of post-arrival exposure may be different than the additional effect of post-arrival learning when pre-arrival learning was present. If migrants already were exposed before leaving, the effect of post-arrival learning is plausibly weaker as low-hanging learning gains have already been exploited.

An important caveat is that both data sets are repeated cross-sections of individual interviews. Hence, differential selection into answering the survey, or differential selection into return migration may affect the estimates of γ. This is potentially important as levels of linguistic and economic integration may affect return migration decisions (Dustmann, 2003). Differential selection on observables can be partially evaluated using the effect of arrival measures on age, sex and educational attainment (among older individuals).

2.7.2 Data

The EU Labor Force Survey (LFS) are harmonized surveys conducted by the national statistical agencies of EU countries as well as some non-EU countries. The surveys include many questions on demographic characteristics and labor market participation. For an individual's main job, it includes a variable on the monthly income decile. For migrants it includes a variable on the global region of birth, as well as the years of residence. I use the surveys between 2008 and 2021, as prior to 2008 information on the years of residence was unavailable. The EU LFS fielded add-on surveys in 2014 and 2021 covering part of the sample. The add-on modules asked the reason for migration,[37] whether one participated in a language course,[38] and language skills upon arrival in 2021, with the following answer options: hardly or none, beginner, intermediate, advanced, or mother tongue.

I restrict the sample to those who arrived at age 18 or older, to capture those who have opted to migrate themselves. Moreover, as I am interested in labor market outcomes, I restrict the sample to those who are aged 60 or below at the time of interview. Moreover, I focus on immigrants who are interviewed up to and including 9 years of arrival as longer times. In line with previous analysis, I consider immigrants who have arrived in 2007 and thereafter. To construct a measure of Duolingo exposure at the origin group-destination-year level, I use the OECD bilateral migration data as discussed in section 2.7.1. Unfortunately, migration data for Cyprus, Greece, Ireland, Malta and Romania is sparse, and I drop these destinations. The full sample includes 668737 individuals from 24 origin regions in 23 destination countries.

[37] Employment – job found before migrating, Employment – no job found before migrating, no job found before migrating, Family reasons, Education or training, Retirement (2021), International protection or asylum, Other

[38] Yes – general language course (2021), Yes – work-specific language course (2021), Yes (2014), No – because language courses were not available or affordable (2021), No – because language skills were sufficient (2021), No – was not necessary (2014), No – for other reasons

I construct several different datasets: those answering the question on language skills upon arrival in 2021, those answering questions on reasons for migration in 2014 and 2021, all of those interviewed within one year of migration, and a comprehensive dataset of all individuals up to and including the fifth of arrival. Descriptive statistics of the three samples are shown in Table 2.D.11.

2.7.3 Language skills and integration upon arrival

TABLE 2.7.1: The Effect of Duolingo Exposure on Language Skills upon Arrival, Migration Reasons and Characteristics

	(1)	(2)	(3)	(4)	(5)
Panel A: Language skills upon arrival (2021)					
	At least beginner	At least intermediate	At least advanced	Mother tongue	Did a language course
$DL_{oc}^{T_d}$	0.198***	0.151***	0.047	-0.031	-0.006
	(0.044)	(0.037)	(0.033)	(0.022)	(0.053)
Observations	19254	19254	19254	19254	18803
R^2	0.32	0.43	0.50	0.74	0.31
Mean dep. var.	0.506	0.343	0.253	0.170	0.405
Panel B: Reason for migration (2014 & 2021)					
	Employment, job on arrival	Employment, no job on arrival	Family	Education	Refugee
$DL_{oc}^{T_d}$	0.062***	-0.012	-0.008	-0.004	-0.045
	(0.024)	(0.028)	(0.040)	(0.019)	(0.028)
Observations	64648	64648	64648	64648	64648
R^2	0.13	0.17	0.10	0.13	0.42
Mean dep. var.	0.203	0.218	0.373	0.067	0.080
Panel C: Observable Characteristics in first year after arrival (2008–2021)					
	Primary educated	Secondary educated	At least tertiary education	Female	Age
$DL_{oc}^{T_d}$	0.039	-0.006	-0.033	-0.090**	-0.211
	(0.030)	(0.035)	(0.044)	(0.037)	(0.617)
Observations	37088	37088	37088	50444	50444
R^2	0.18	0.11	0.17	0.05	0.09
Mean dep. var.	0.216	0.311	0.473	0.542	32.675
Origin group-year FE	✓	✓	✓	✓	✓
Destination-year FE	✓	✓	✓	✓	✓
Origin group-Destination FE	✓	✓	✓	✓	✓

Notes: OLS estimations of the model of equation 2.10. Panel A is from the 2021 add on sample, Panel B is from the 2014 and 2021 add-on samples, Panel C is from the full 2008-2021 LFS. "At least a beginner" in column 1 of panel A is 1 if a respondent indicates to have had at least beginner level language skills and 0 if she answers to have had hardly any or no language skills. The subsequent levels in column 2–4 indicate binary indicators for higher minimum levels of language skills. "did a language course" is a binary indicator for whether an immigrant did a language course after arrival. Panel B reports the levels of migration reasons, where the omitted category is "other". Panel C reports current educational attainment for respondents at least 25 years of age and a binary indicator for female and an integer variable for age. The measure of Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. Standard errors reported in parentheses are two-way clustered: on the country group of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 2.D.13 shows the results from an event study of three levels of language skills around large increases in Duolingo exposure. Due to the relatively low number of observations I bin the

cohorts relative to arrival. I find no evidence of pre-trends in language skills upon arrival, but a considerable increase of language skills after arrival. Panel A of Table 2.7.1 reports estimation results from the model of equation 2.10. On a mean of about 51%, I find that full Duolingo exposure increased the language skills of migrants by 20 percentage points. As Duolingo exposure in 2021 is about 40 percentage points across the EU migrant pool, this implies that the share of respondents with at least beginner skills, or an increase of around 15%. Moreover, the share of individuals reporting at least intermediate language skills has also improved considerable, but there is no effect on advanced language skills or speaking the native language as your mother tongue. The latter is reassuring, as the availability of language learning should not impact those speaking the language natively. Moreover, as Duolingo enables acquisition of basic language skills, it probably has no large effects on advanced skills. Ultimately, there is no effect on the propensity to have done a language course in the origin country, suggesting that Duolingo neither crowds out taking language courses in the host country, but also does not boost it. Panel C of 2.7.2 explores the effect on reasons for taking or not taking a language course. Respondents are much more likely to answer that they did not take a course because they possessed sufficient language skills upon arrival.

Does this large increase also affect the reasons for immigration? It turns out that it does. Panel B of Table 2.7.1 shows that migrants are 6 percentage points more likely to have arrived for employment with having a found a job before arriving. This suggest that better language skills enable migrants to find employment on a distance. Panel C studies the selection of immigrants. I find no significant change in the skill composition in terms of formal education among those migrants aged 25 or above, although results suggest that the pool of migrants became lower skilled. Moreover, I do find that the proportion of men increases strongly.

Figure 2.7.1 examines the heterogeneity of the effect by reason for migration. Immigrants arriving for education have the strongest language skills: almost 80% have some skills upon arrival. Language skills are particularly low for refugees: about 30% have basic language skills. I find that the treatment effects are largest for immigrants who arrive for employment with a job upon arrival and "other" migrants. Moreover, the effect at beginner levels of skills are also significant for family and educational migrants. Skills among migrants who are less prepared upon arrival, those arriving for employment without a job and refugees, are unaffected by exposure to low-cost language learning.

Panel A and B of Table 2.7.2 reports results for economic integration. In the first year after arrival, Duolingo exposure increases the probability to work by 10 percentage points, which goes at the expense of all other categories. In term of job characteristics, I find no difference in the propensity to be self employed or to work on a temporary contract. In line with the increase in employment, working hours increase by 15% on average. However, using the information on income deciles based on monthly wage income, it seems that migrants do not earn more. In fact, these estimates suggest that hourly wages declined. This could be partially driven by the worsening skill composition of the immigrant pool as suggested by the results in Table 2.7.1.

FIGURE 2.7.1: Migration Reason-specific Language Skills and the Effect of Duolingo

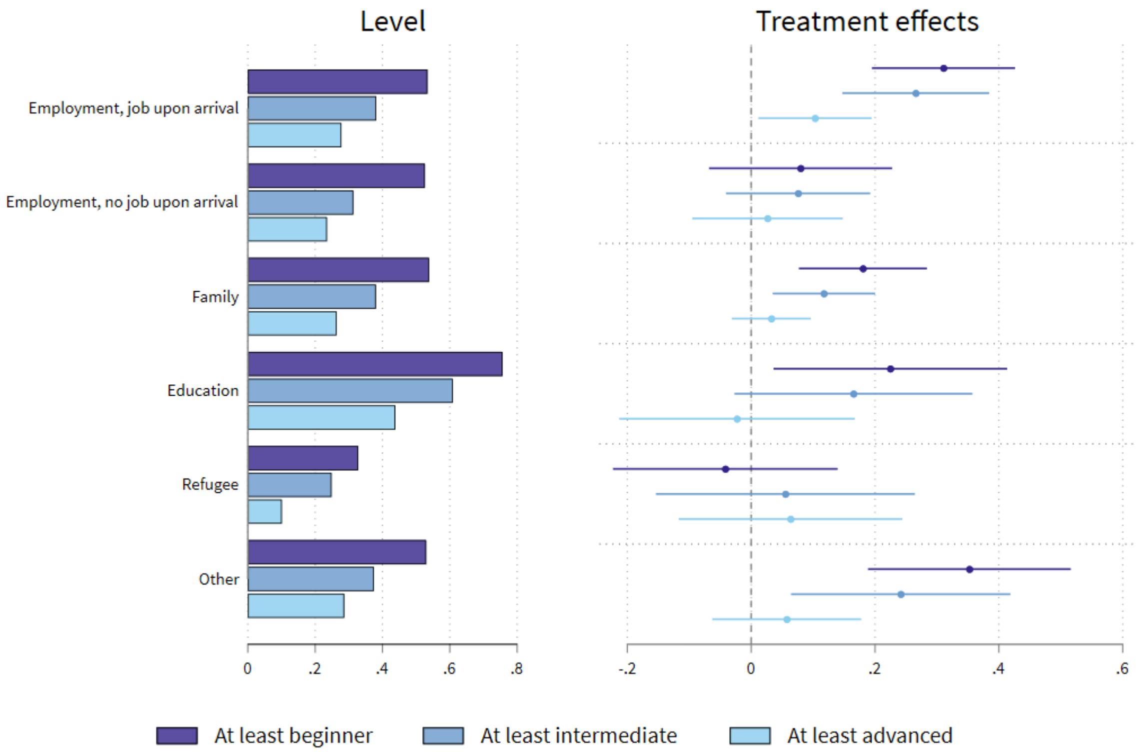

Notes: Levels of language skills upon arrival by migration reason among those who arrived before any Duolingo exposure. (left) and the treatment effect of Duolingo exposure by migrant group (right). OLS estimations of the model of equation 2.10 including indicators for the levels of reasons for migration (not shown) and an interaction between the pre-arrival Duolingo exposure and all levels of reasons for migration. The estimates on the interaction terms are shown with 95% confidence intervals based on two-way clustered standard errors: on the country group of origin and country of destination level. For information on the data and outcome variables and migration reasons, see notes to Table 2.7.1.

TABLE 2.7.2: The Effect of Duolingo Exposure on Additional Outcomes

	(1)	(2)	(3)	(4)	(5)
Panel A: Main activity in first year after arrival					
	Employed	Unemployed	Retired	Student	Stay-at-home
$DL_{oc}^{T_d}$	0.104**	-0.012	-0.016*	-0.037	-0.023
	(0.049)	(0.044)	(0.009)	(0.036)	(0.038)
Observations	45216	45216	45216	45216	45216
R^2	0.23	0.11	0.04	0.15	0.14
Mean dep. var.	0.496	0.141	0.040	0.146	0.155
Panel B: Employment in first year after arrival					
	Self-employed	Temporary work	log of hours usually worked	Income decile	Lowest income decile
$DL_{oc}^{T_d}$	-0.018	0.025	0.143**	-0.305	-0.055
	(0.031)	(0.041)	(0.064)	(0.373)	(0.046)
Observations	26699	24028	25989	16000	16000
R^2	0.12	0.17	0.11	0.24	0.14
Mean dep. var.	0.093	0.324	3.559	4.937	0.134
Panel C: Language courses (2021)					
	Yes, General	Yes, Work-specific	Not available of affordable	Not because skilled enough	Not for other reasons
$DL_{oc}^{T_d}$	-0.008	-0.003	-0.039	0.085**	-0.035
	(0.054)	(0.016)	(0.027)	(0.036)	(0.039)
Observations	19128	19128	19128	19128	19128
R^2	0.28	0.13	0.15	0.41	0.16
Mean dep. var.	0.356	0.053	0.107	0.345	0.140
Origin group-year FE	✓	✓	✓	✓	✓
Destination-year FE	✓	✓	✓	✓	✓
Origin group-Destination FE	✓	✓	✓	✓	✓

Notes: OLS estimations of the model of equation 2.10. Panel A and B consider those interviewed in the first year after arrival in the full 2008-2021 LFS, Panel C is from the 2021 add-on sample. The outcomes in Panel A are the mutually exclusive categories of main activity within the first year after arrival, Panel B shows other job characteristics. The EU LFS does not include labor income but reports the within country-year income decile. In Column 4 I report OLS regressions using the integer income decile between 1 and 10 as outcome variable, and column 5 uses a binary indicator for the lowest income decile. The outcomes in Panel C are the levels of the categorical question on whether someone has taken a traditional language course after arrival and the reasons why if yes and not. The measure of Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. Standard errors reported in parentheses are two-way clustered: on the country group of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.7.4 Integration after arrival

To study integration in the first year after arrival, I restrict the sample to those in the first five years after arrival, for two reasons. First, as after five years many immigrants have considerable language skills and it seems implausible that the introduction of low-cost language learning affects is helpful. Second, the longer after arrival, the stronger return migration impacts the estimates. Unfortunately, questions on language skills are only included in 2014 and 2021, which makes a detailed analysis of the impact on language skills over time infeasible. Instead, Table 2.7.3 shows the result of pre- and post-arrival exposure on several employment-related outcomes. Column 1 shows the effect upon arrival as shown in previous tables. Column 2-4 report results from the full sample of immigrants within the first 5 years of arrival. To study whether the initial gains in outcomes fade out over time, I interact the Duolingo exposure with the log of years since arrival plus one in column 2. Column 3 introduces the post-arrival treatment intensity and Column 4 includes an interaction of pre-arrival intensity with dummies of years since arrival (not shown). The last column limits the sample to those who had no exposure to Duolingo before arrival.

Panel A of Table 2.7.3 shows that employment gains from pre-arrival exposure slowly fade out. Exposure to Duolingo after arrival positive affects employment. This effect is stronger among those who migrated when Duolingo was not yet available: the probability to be employed increased by almost 11 percentage points. This estimate is remarkably similar to the effect from pre-arrival exposure in Column 1.

TABLE 2.7.3: The Effect of Duolingo Exposure on Language Skills after arrival

	(1) Upon arrival	(2)	(3) Full	(4)	(5) Arrival before any Duolingo Exposure
Panel A: Main activity: employment					
$DL_{oc}^{T_d}$	0.104**	0.103***	0.090**	0.070*	
	(0.049)	(0.034)	(0.035)	(0.039)	
$DL_{oc}^{T_d} \times \log(t-c+1)$		-0.037	-0.063**		
		(0.026)	(0.026)		
$DL_{otc}^{T_d,post}$			0.062***	0.048*	0.114***
			(0.023)	(0.026)	(0.041)
Observations	45216	398697	398697	398697	234743
R^2	0.23	0.15	0.15	0.15	0.14
Mean dep. var.	0.496	0.558	0.558	0.558	0.550
$DL_{oc}^{T_d} \times (t-c)$ FE				✓	

Notes: OLS estimations of the model of equation 2.13, with the following outcomes: Panel A shows results for a binary indicator for having employment as one's main activity. Column 1 shows the effect upon arrival as shown in previous tables. Column 2-4 report results from the full sample of immigrants within the first 5 years of arrival. To study whether the initial gains in outcomes fade out over time, I interact the Duolingo exposure with the log of years since arrival plus one in column 2. Column 3 introduces the post-arrival treatment intensity and Column 4 includes an interaction of pre-arrival intensity with dummies of years since arrival (not shown). The last column limits the sample to those who had no exposure to Duolingo before arrival. The measure of Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. Standard errors reported in parentheses are two-way clustered: on the country group of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.7.5 Language skills in the US

Section 2.D.6 shows and discusses a similar analysis for the US as presented above for the EU. Contrary to the EU case, I find that pre-exposure Duolingo availability does not increase language skills in the US. Moreover, I find suggestive evidence that selection into migration becomes more negative in terms of education, although statistically insignificant. However, full exposure to Duolingo after arrival increases language skills of immigrants, and improves employment rates by 4 percentage points. The stark differences between the EU and US results for pre-arrival exposure could be explained by several reasons. First, language skills upon arrival are arguably better in the case of the US than for the EU as English is more widely spoken abroad than any of the other European languages. However, as the EU LFS and ACS ask questions with different answer options, it is not possible to test difference in the levels of immigrants' language skills directly. Second, English courses were often the first available courses, when Duolingo was not yet popular in a given linguistic region of origin. On the contrary, the courses to European languages were often rolled out after English and in later years when Duolingo was already established and amassed many users. This could have led to a quicker and stronger learning response in the EU. Third, European countries have stricter language requirements than the US, which may provide stronger motives for learning before arrival.

Altogether, I find employment effects of 10 percentage points for full pre-arrival exposure and post-arrival exposure in the EU, and no effect for pre-arrival exposure and a 4 percentage point effect for post-arrival exposure in the US. The magnitude of some of these effects is surprisingly large, especially given the potential attenuation bias due to aggregation of exposure measures as discussed above for the EU. Comparing my estimates to those from the literature, these estimates seem large. However, it is hard to compare these estimates directly, as I do not have reliable information about the share of migrants taking up Duolingo. Foged, Hasager and Peri (2024) finds that extensive language training programs (200 additional hours of instruction) in Denmark improve short-run employment rates of refugees by 3–4 percentage for women and 8 percentage points for men.[39] Lochmann, Rapoport and Speciale (2019) finds similarly large effect sizes among immigrants in France, which also includes non-refugees. However, these are effect sizes of actual uptake, and my estimates are intention-to-treat estimates with an unknown take-up rate.

The calculated Duolingo exposures based on communication probabilities are imperfect measures of whether a Duolingo course is relevant or not across a migration corridor, which could lead to overestimation of the estimate of the effects of course availability. For example, as many courses to European languages are using English as a source language, an underestimation of the share of English speakers among the pool of prospective migrants could lead to an exaggeration of the effect size. This can be driven by two different factors: first, the migrant pool may have considerably larger English skills than the whole native population and second, the data I use on the share of speakers from (Ginsburgh, Melitz and Toubal, 2017) relies predominantly on data from 2012 and before.

2.8 Conclusion

Language proficiency is paramount to the success of immigrants. The ability to learn a host country's language enable the acquisition of skills that increase labor market earnings and lower migration costs. Moreover, the ability to learn a language may also foster interest in a country's language and culture, which subsequently. Low-cost language learning offered on online platforms facilitate language acquisition among individuals without prior access to language learning, deem investing in a language course too risky, or perceive the cultural distance as large. Using the staggered introduction of language courses on Duolingo and the global distribution of spoken languages, I provide causal evidence of the availability of low-cost language learning on migration intentions and flows, selection into migration, the language skills of immigrants and subsequent economic integration. As a first step, I aim to understand how the availability of low-cost language availability affected language learning and skills. I find that availability of Duolingo improved scores on (predominantly) passive components of English-language tests and increased school-based instruction of learnable languages.

To study the effect of language course availability on bilateral migration intentions and flows, I construct a measure of how beneficial taking up a language course is for a prospective migrant across a corridor. This foreign Duolingo exposure captures the probability course take-up of newly enables two randomly picked individuals in two countries to communicate to eachother. Moreover, I additionally construct a measure of domestic language exposure capturing the usability of a language course at home. Using these, I find that foreign Duolingo exposure strongly

[39] In line with this paper, I find that initial employment gains relative to unexposed migrants become smaller after several years.

increases migration intentions within the first 3 years after course introduction, by on average 45%. Additional analysis shows that more than half of this increase is driven by diversion of migration intentions between destinations. However, I do not find strong evidence that this translates to increases in the extent of migration flows to OECD countries. Using a recently published dataset on global scholarly migration flows, I find that courses from English to other languages increase by about 7% suggesting that the availability of low-cost language learning at least can change the composition of the migrant pool.

To further study the effects conditional on migration, I turn to the European Union Labor Force Survey (EU LFS). As many migrants lack basic language skills upon arrival, they have large economic and social incentives to engage in language learning. I find that availability of an appropriate language course before arrival strongly increases the probability to possess beginner-level language skills by 20 percentage points, and to possess intermediate-level skills by 15 percentage points, but find no effects at the highest level of skills or on the share of mother tongue speakers. Moreover, I find that the proportion of migrants arriving for employment reasons with a job on arrival increases, which is suggestive of the fact that improved language proficiency increases the ability to search and find a job before arrival. Considering heterogeneity in the effect sizes by migration motive, I find that the effect is driven by those arriving with a job upon arrival, family migrants and those moving for education, which are groups who often had time to prepare before migration. I find that migrants who are unlikely to (be able to) prepare before arrival, economic migrants without pre-arrival jobs and refugees, to be unaffected. More migrants are in work due to more jobs being found before arrival, but this advantages diminishes after 3 years as unexposed migrants catch up. Despite the lowered threshold to language learning, I do not find evidence that low-cost language learning considerably worsened the educational composition of the migrant pool.

Although the identification strategy is weaker for a similar analysis among U.S. migrants, I find that no discernible pre-trends in migrant characteristics exist before exposure to Duolingo. Contrary to the EU, I do not find any effect on language skills upon arrival in the US. This could be driven by an offsetting effect through changing selection: the estimates suggest that the share of immigrants with a tertiary education has decreased by several percentage points. Exposure to Duolingo after arrival to the US improves language skills and integration outcomes. The differences in results between the EU and US could be explained by various factors, including the lack of basic skills among many EU migrants, the absence of language requirements in the US for many visa types and the relative popularity of Duolingo in the US compared to typical origin countries.

The findings of this paper suggest that availability of low-cost language learning increases the pool of potential migrants considerably and that it increases language skills at the low end of the distribution upon arrival. Policy makers concerned with addressing worker shortages or with improving integration could take this information to the heart by facilitating language learning opportunities abroad. However, contrary to the availability of costly traditional language learning as in (Jaschke and Keita, 2021), low-cost language learning does not increase the average educational attainment of migrants. As many refugee hosting countries spend vast resources on integration courses due to low initial language knowledge, the positive impact of digital learning methods on language skills can motivate the targeted development of digital language courses for major refugee origin languages.

Altogether, this paper shows how one aspect of the internet can improve the integration of migrants. However, less remains known about the overall effect of the internet on integration. Through improved information provision, immigrants may become better informed and make better destination choices, improving integration (Porcher, 2020). However, much in line with the results of this paper, migrants' skills upon arrival have improved during the internet era. Yet, due to the availability of internet in migrants' origin country, migrants spend more time online and less time with natives, worsening linguistic integration after arrival (Yarkin, 2024). Fruitful avenues for further exploration in this literature are the role of remote work on migrant selection and the role of migration experiences transmitted through social media on migration decisions.

Ultimately, the rich variation in low-cost language learning explored in this paper can be used by other scholars trying to understand how language learning opportunities affect other aspects of integration, such as social integration through time use with natives or intermarriage rates. Moreover, it can be employed to study how it affects other bilateral ties between countries, including interest in foreign cultures, tourist visits and trade patterns.

Appendix

2.A Model details and extentions

2.A.1 Total migration

The derivative of the probability to migrate to language skills is:

$$\frac{\partial \mathbb{P}_{od}}{\partial l_{oS}^{T}} = \frac{e^{\mu_{od}+l_{oS}^{T}b_{oSd}^{T}}\left[\left(\sum_{d'} e^{\mu_{d'}+l_{oS}^{T}b_{od'T}}\right)b_{oSd}^{T} - \left(\sum_{d'} b_{od'T}e^{\mu_{d'}+l_{oS}^{T}b_{od'T}}\right)\right]}{\left(\sum_{d'} e^{\mu_{d'}+l_{oS}^{T}b_{od'T}}\right)^{2}} = \mathbb{P}_{od}\left(b_{oSd}^{T} - \sum_{d'} \mathbb{P}_{od'}b_{od'T}\right) \lesseqgtr 0 \quad (2.1)$$

A change in language skills increases the probability to migrate to destination d if the return to the language skill in the destination is larger than the average return across potential locations, weighted with migration probabilities. As most people do not migrate ($\mathbb{P}_{oo}$ is large w.r.t. $\mathbb{P}_{od}$, where $d \neq o$), the domestic return to language skills plays a prominent role in equation 2.1. If language skills are only rewarded in one destination, larger language skills increase the probability to migrate to that destination. However, if the language skill is rewarded more in other destinations or at home, this may decrease migration flows.

Moreover, larger language skills do not need to imply larger *total* emigration. equation 2.2 shows the condition for which this is the case. Total emigration increases if the migration probability-weighted foreign returns exceed domestic returns. Hence, if migration probabilities are low enough, and a language skills is moderately valued on the domestic labor markets, total emigration decreases. This is likely to be the case for many countries where English is rewarded on domestic labor markets but migration links to English-speaking destinations are weak (e.g. due to high moving costs).

$$\frac{\partial \sum_{d\neq o} \mathbb{P}_{od}}{\partial l_{oS}^{T}} = \mathbb{P}_{oo}\sum_{d\neq 0} \mathbb{P}_{od}b_{oSd}^{T} - (1-\mathbb{P}_{oo})\mathbb{P}_{oo}b_{oSo}^{T} > 0 \implies \frac{1}{\sum_{d\neq 0} \mathbb{P}_{od}}\sum_{d\neq 0} \mathbb{P}_{od}b_{oSd}^{T} > b_{oSo}^{T} \quad (2.2)$$

2.A.2 Derivation of equation 2.7

The marginal benefit of language skills generally is a function of the level of language skills through the migration probabilities. In the low migration limit, the deterministic part of utility of staying exceeds the utility of migrating. This is a reasonable assumption, as most people never migrate across borders during their lifetime. Using (1) the low migration limit to eliminate the skill-dependence of the denominator of the migration probability and (2) the property that $e^x \approx (1+x)$ if x is small, equation 2.6 can be written as:

$$\sum_d \mathbb{P}_{od} b^T_{oSd} \approx \sum_d b^T_{oSd} \frac{e^{\mu_{od}}}{e^{\mu_{oo}}} e^{l^T_{oS} b^T_{oSd}} = b^T_{oSd}(1 + b^T_{oSd} l^T_{oS})\mathbb{P}_{od}(0) \tag{2.3}$$

$\mathbb{P}_{od}(0)$ denotes the approximate migration probability if $l = 0$. Due to the exponential form of the migration probabilities, the benefits from larger language skills are increasing in language skills returns in the low migration limit. Plugging this expression into equation 2.6 gives the following:

$$l^T_{oS}\left(2c_T - \sum_d \mathbb{P}_{od}(0)(b^T_{oSd})^2\right) = b^T_{oSd}\mathbb{P}_{od}(0) \tag{2.4}$$

Rearranging gives the following expression for equilibrium levels of language skills:

$$l^{T*}_{oS} \approx \frac{\mathbb{P}_{oo} b^T_{oSo} + \sum_{d \neq o} \mathbb{P}_{od}(0) b^T_{oSd}}{2c_{oST} - \sum_d \mathbb{P}_{od}(0)(b^T_{oSd})^2} \tag{2.5}$$

The second term in the denominator can be assumed small compared to $2c_{oST}$. To see why this is the case, I estimate c_{oST} and compare it to the second term in the denominator. Returns for English in European countries, where 30-70% (l) of the population learns English, vary between 10-50% (b, see section 2.4.1 for a discussion on these). Equating marginal costs and benefits in absence of migration using the mid-point of the ranges for b and l gives an estimated cost of $c_{oST} = b/2s = 0.3/(2 * 0.5) = 0.3$. As migration probabilities are low, and returns to foreign languages abroad do not exceed 50% (see section 2.2), this term is much smaller than $2c_{oST}$ in the low migration limit. Using this approximation, I arrive at the right hand side of equation 2.7:

$$l^{T*}_{oS} \approx \left(\mathbb{P}_{oo} b^T_{oSo} + \sum_{d \neq o} \mathbb{P}_{od}(0) b^T_{oSd}\right) \frac{1 + \eta_{oST} Duolingo_{ST}}{\kappa_{oST}} \tag{2.6}$$

2.A.3 Calculation of the proxy for returns to skills

α_{cL} denotes the number of language L in country c. If I assume that all languages are equally and randomly distributed among a country's population (e.g. if $\alpha_{cL} = 0.5$ and $\alpha_{cL'} = 0.5$, 25% of people speak none of L and L'. 25% speak only L, 25% speak only L' and 25% speak both), I can calculate the probability that two randomly chosen individuals can communicate. I denote the product of the number of speakers of a language in the origin and destination as $\alpha^2_l = \alpha_{ol}\alpha_{dl}$ and the number of languages spoken in either country by N. For convenience, l is an ordered index of languages. Using the law of total probability, I can write the probability two randomly picked individuals, one in the origin and one in the destination, as a function of α_{cL}'s. Here, I show the first $k = 4$ out of $k = N$ terms:

$$\mathbb{P}(comm_{od}) = \sum_l \alpha^2_l - \sum_{l > l'} \alpha^2_l \alpha^2_{l'} + \sum_{l > l' > l''} \alpha^2_l \alpha^2_{o'} \alpha^2_{l''} - \sum_{l > l' > l'' > l'''} \alpha^2_l \alpha^2_{o'} \alpha^2_{l''} \alpha^2_{l'''} \dots \tag{2.7}$$

For a large number of common languages, this implies that there are many terms. For N languages, the k$^{\text{th}}$ term contains are $\binom{N}{k}$ elements. As the largest number of shared languages between any two countries in the data is 6, I calculate the terms up to and including k = 6. Using this result, and setting $\alpha_{oS} = 1$, I obtain the probability conditional on the individual in o speaking S, $\mathbb{P}(comm_{od}|S)$. Additionally, setting $\alpha_{oT} = 1$, I can calculate $\mathbb{P}(comm_{od}|S \wedge T)$. Using those, the availability of learning T from S (sloppily denoted by $DL_{S\to T}$) expands the set of people to communicate with by:

$$\mathbb{P}(comm_{od}|DL_{S\to T}, S) = \mathbb{P}(comm_{od}|S \wedge T) - \mathbb{P}(comm_{od}|S) \tag{2.8}$$

The total probability of a Duolingo course from S to T facilitating communication between two randomly chosen individuals in o and d is then simple found by:

$$\mathbb{P}(comm_{od}|DL_{S\to T}) = \mathbb{P}(comm_{od}|DL_{S\to T}, S)\mathbb{P}(S) = \mathbb{P}(comm_{od}|DL_{S\to T}, S)\alpha_{oS} \tag{2.9}$$

2.B Descriptives on Language Learning and Duolingo

2.B.1 Online courses

FIGURE 2.B.1: Share of Population Doing any Online Course across the EU

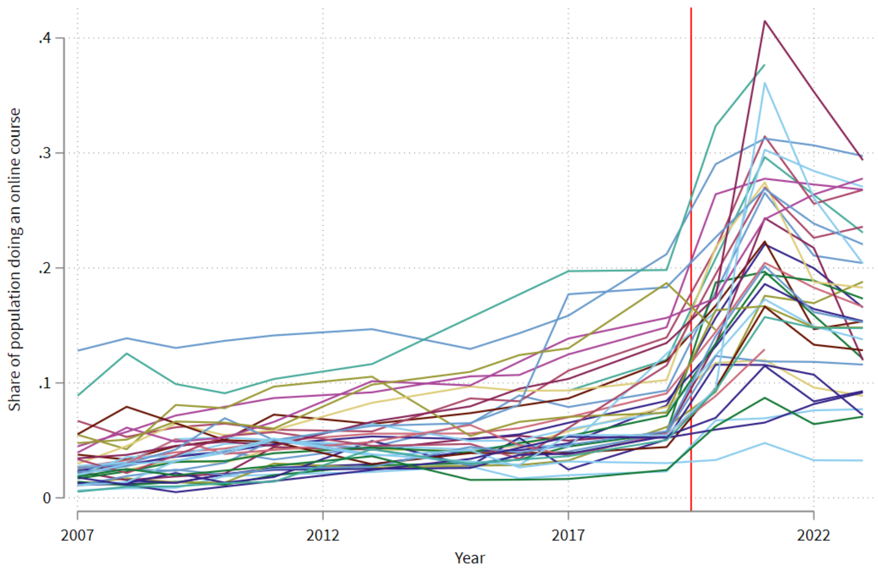

Notes: The share of population doing any online course between 2007 and 2023, by EU member state. The vertical line denotes the onset of the Covid-19 pandemic. Data stems from the EU survey on the use of Information and Communication Technologies (ICT) in households and by individuals and is available in eurostat table isoc_ci_ac_i.

2.B.2 Course content

FIGURE 2.B.2: Tasks on Duolingo

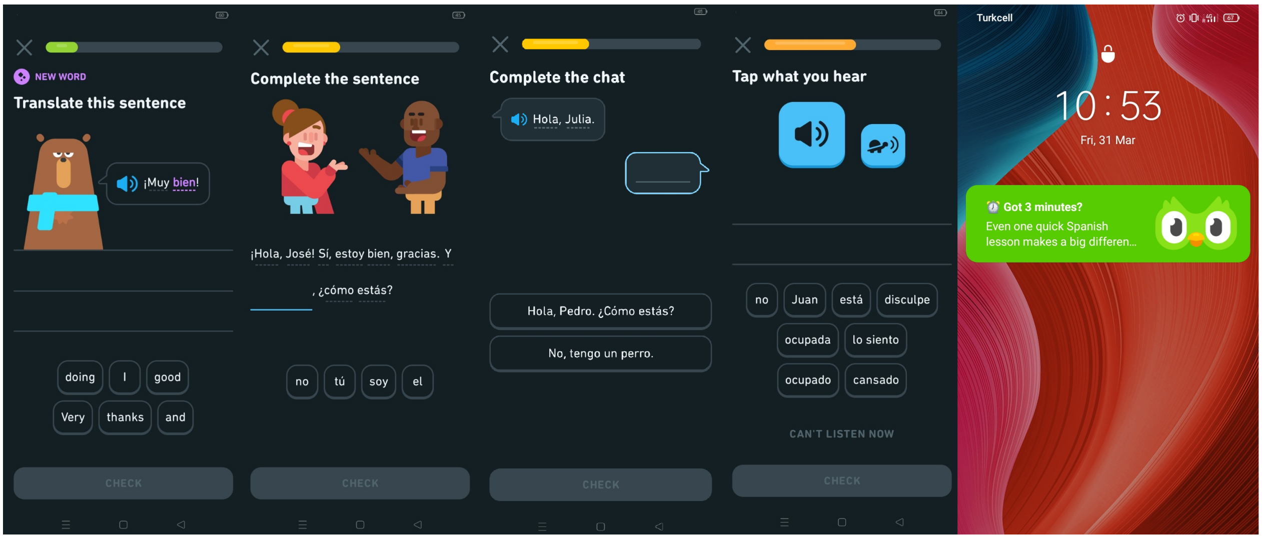

Notes: Example of typical tasks on Duolingo for the English to Spanish course.

2.B.3 Duolingo Courses

FIGURE 2.B.3: Available Courses as of 2022

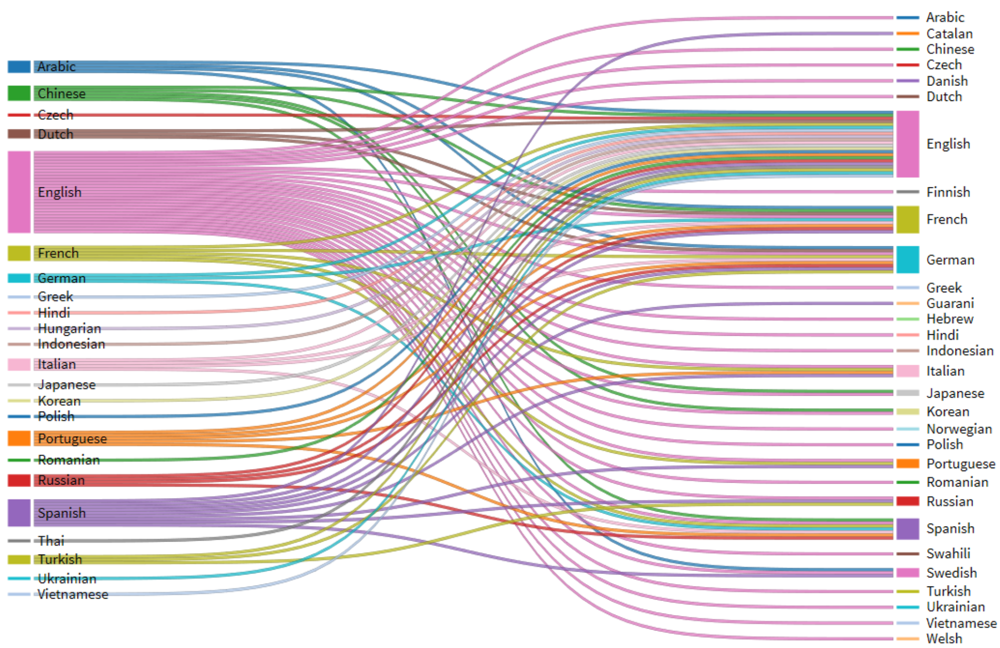

Notes: Sankey diagram of all available courses on Duolingo. In total, this comprises 84 modules over 23 source languages and 30 target languages. Information on courses and rollout dates is obtained from Fandom and verified using the Duolingo website.

FIGURE 2.B.4: Most Studies Language by Country on Duolingo in 2021

Notes: Most studied language by country in 2021. Data from the 2021 Duolingo Language Report.

FIGURE 2.B.5: Percentage of Learners Learning English, Spanish, or French across the World in 2020

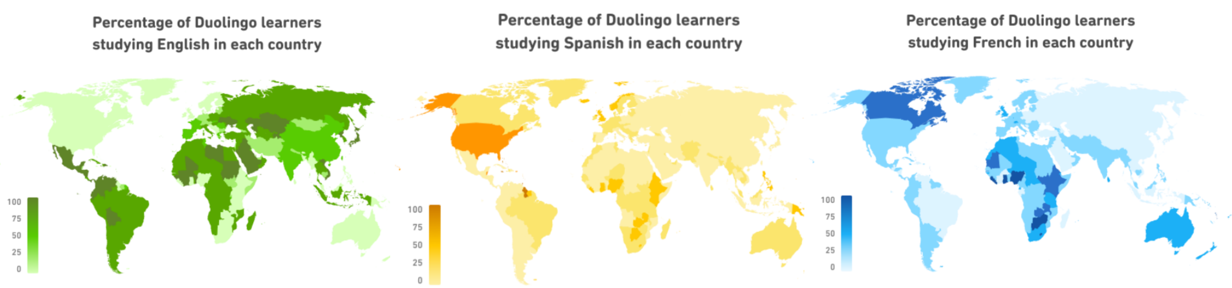

(A) English (B) Spanish (C) French

Notes: Most studied language by country in 2021. Data from the 2020 Duolingo Language Report.

FIGURE 2.B.6: Monthly Active Users on Duolingo between 2012 and 2023

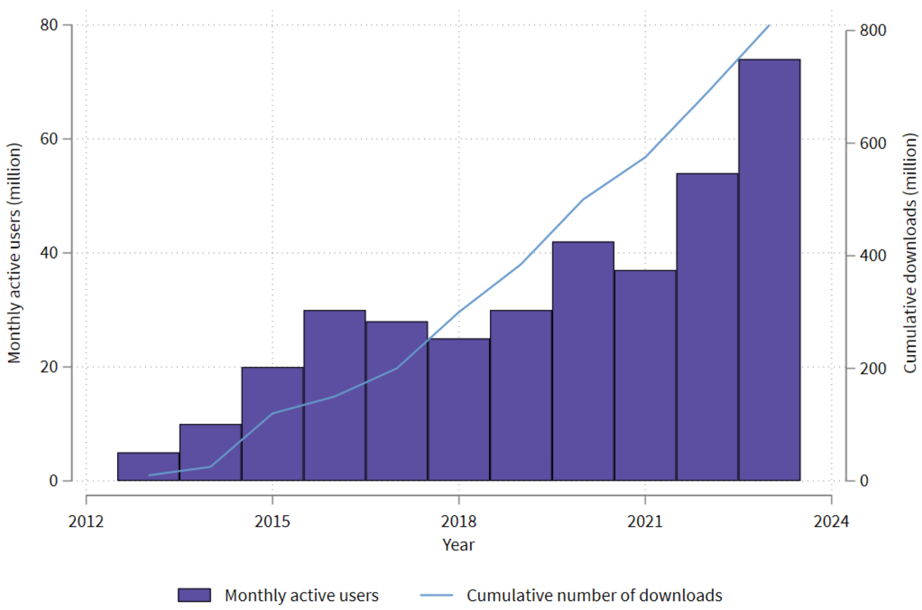

Notes: Monthly active users and the number of downloads of the Duolingo mobile application. For the definition of a monthly active users, see notes to Figure 2.3.1. Numbers on the Monthly Active Users and cumulative downloads are obtained from `https://www.businessofapps.com/data/duolingo-statistics/`.

2.B.4 Uptake of Courses

FIGURE 2.B.7: Number of Learners by Course for 25 Most Popular Courses

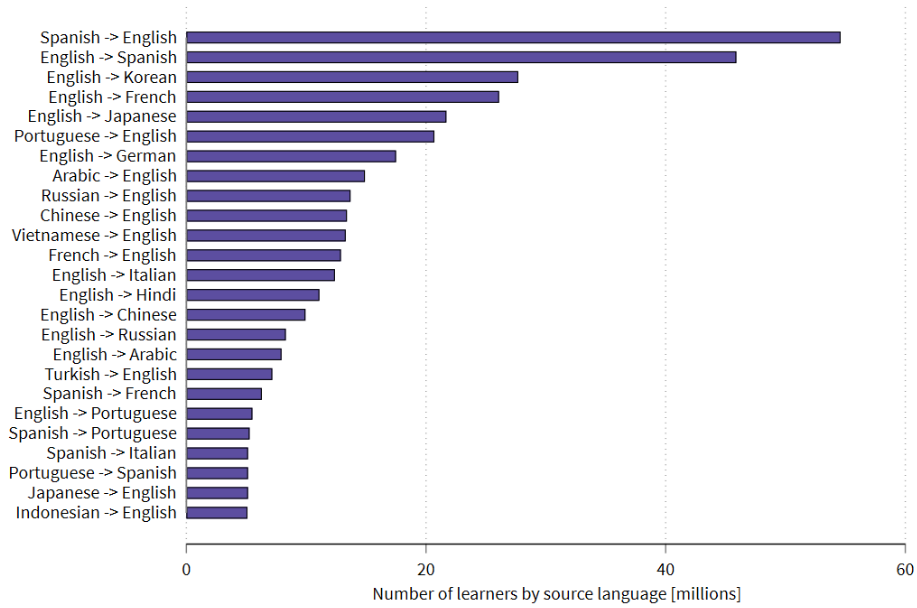

Notes: Total numbers of learners by course, as indicated on the Duolingo website on October 6, 2022.

FIGURE 2.B.8: Number of Duolingo Learners by Source and Target Language

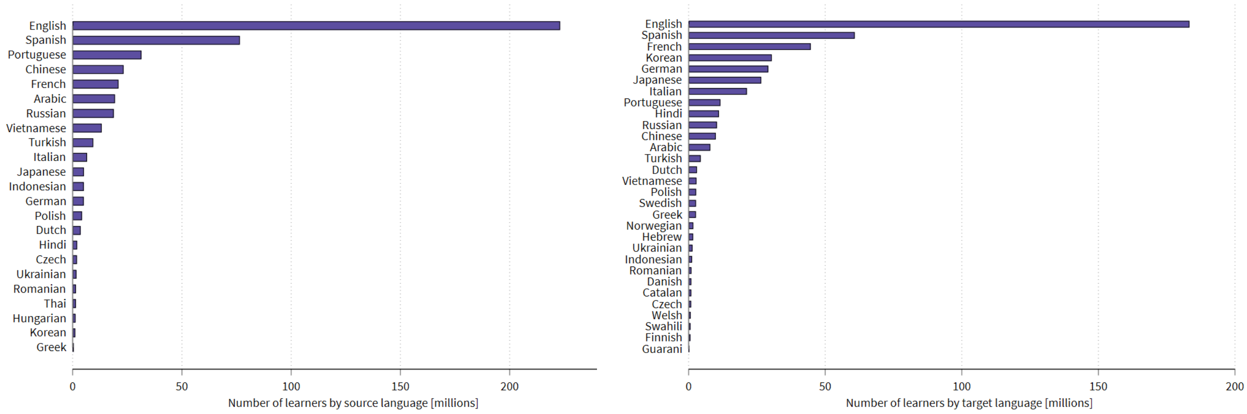

(A) Source language (B) Target language

Notes: Total numbers of learners by (a) source and (b) destination language indicated on Duolingo on October 6, 2022. To calculate the total number of learners by source and target language, I sum over all courses. This sum represents the total number of instances someone started learning a specific target language from a specific source language on Duolingo. In practice, users may initiate multiple courses from the same source language, so the numbers in (a) are higher than the total unique individuals using a specific source language. Likewise, the numbers in (b) represent the total number of learner-language attempts for a specific target language.

FIGURE 2.B.9: Relative Number of Duolingo Learners by Source and Target Language

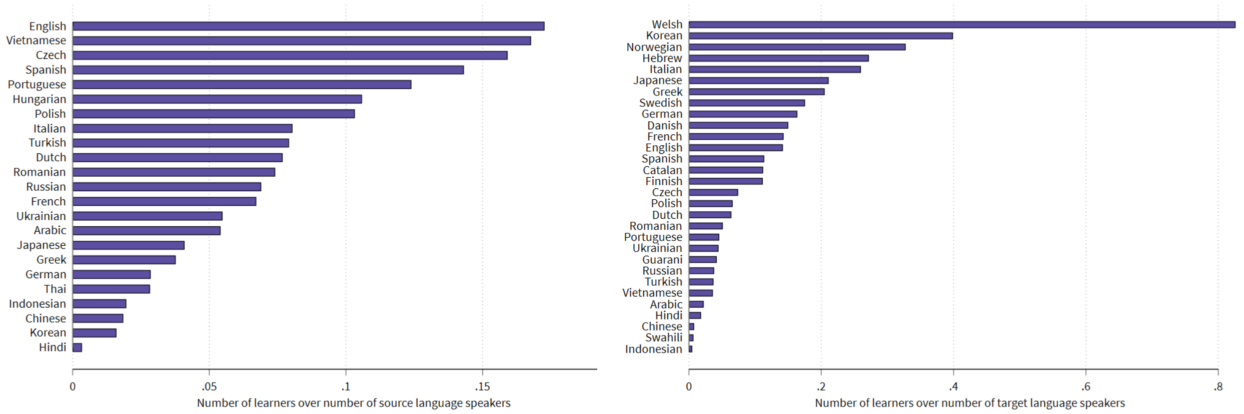

(A) Source language (B) Target language

Notes: See notes to Figure 2.B.8 for a description of the calculation of the total number of learners. To calculate the number of learners relative to speakers, I divide the numbers reported in Figure 2.B.8 by the total number of speakers by language from Ginsburgh, Melitz and Toubal (2017).

TABLE 2.B.1: Determinants of the Number of Learners of Duolingo Courses

	(1) users	(2) users	(3) users
Source language speakers	0.008***	0.004***	
	(0.001)	(0.000)	
Target language speakers	0.008***	0.004***	
	(0.001)	(0.000)	
Source speakers $\times$ Target speakers (100 million)		0.001***	0.002*
		(0.000)	(0.001)
Observations	84	84	52
Source and Target FE			✓

Notes: OLS regressions of the number of learners on Duolingo, as measured of the number of learners by language course, on the number of speakers of the source and the target language. Column (3) introduces fixed effects on the source- and target language level, which drops 32 courses where either the source or target language is a singleton. Standard errors are clustered two-way on the source and destination language. Data on learners is obtained from the Duolingo platform in October 2022.

TABLE 2.B.2: Internet Traffic to Duolingo by Global Region

Global region	Share of traffic
North America	27%
South America	12%
Western Europe	11%
Eastern Europe	8%
Northern Europe	7%
East Asia	6%
Southern Europe	5%
South East Asia	5%
Central America	4%
Other	15%

Notes: Data has been obtained from Semrush (https://de.semrush.com/website/duolingo.com/overview/) in June 2024. Other includes Africa, Middle East and Turkey, and Oceania. Traffic from these regions is too low to analyze in isolation, but together accounts for about 15% of all traffic.

2.B.5 Visualizing Duolingo Exposure

FIGURE 2.B.10: Foreign Duolingo Exposure by Directed Country Pair (2012-2023)

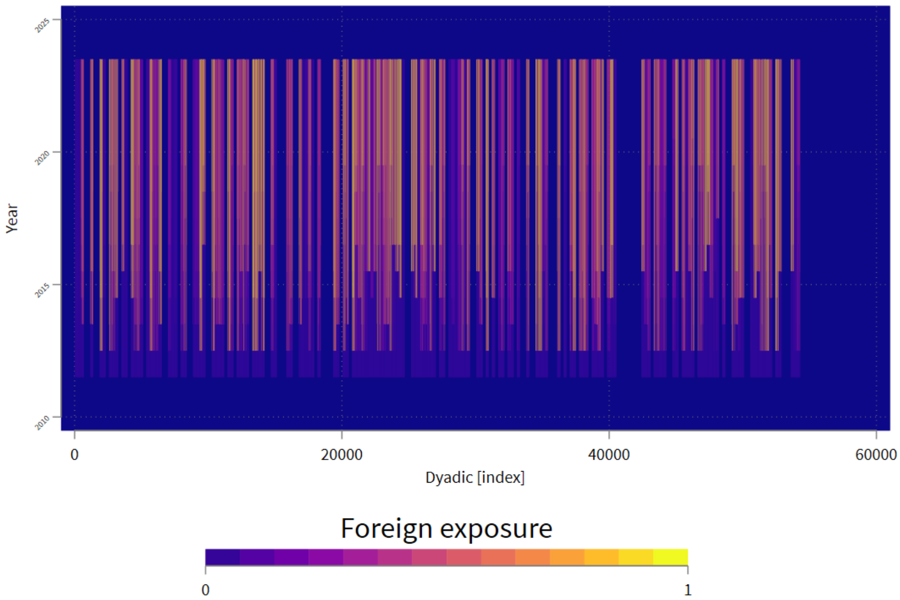

Notes: Dyadic foreign exposure to Duolingo DL_{odt} between 2012 and 2023. Brighter colors indicate a larger exposure.

FIGURE 2.B.11: Average Foreign Duolingo Exposure by Origin Country (2012-2023)

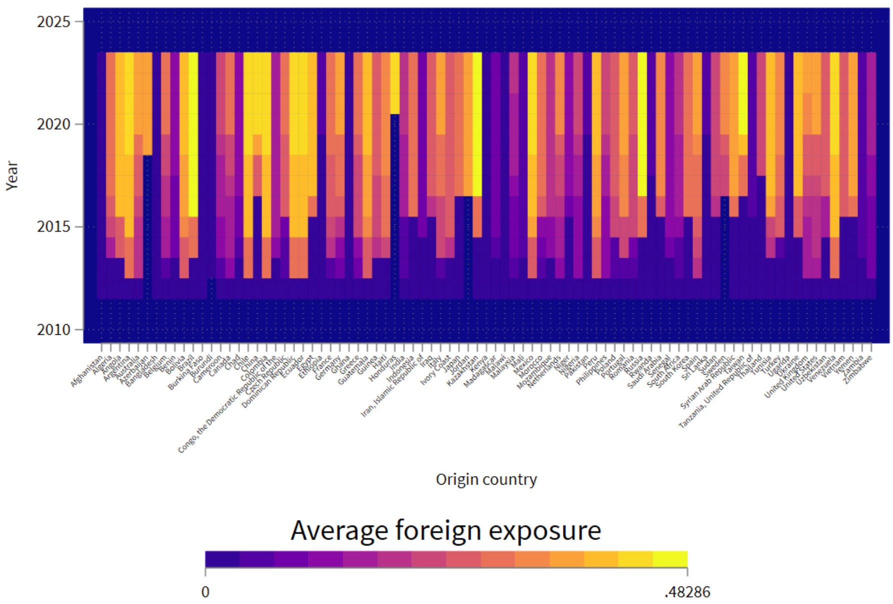

Notes: Simple-weighted foreign exposure to Duolingo $\frac{1}{N}\sum_d DL_{odt}$ between 2012 and 2023 across origin countries. Brighter colors indicate a larger exposure.

FIGURE 2.B.12: Variation in Foreign Duolingo Exposure to and from the Netherlands (2012-2023)

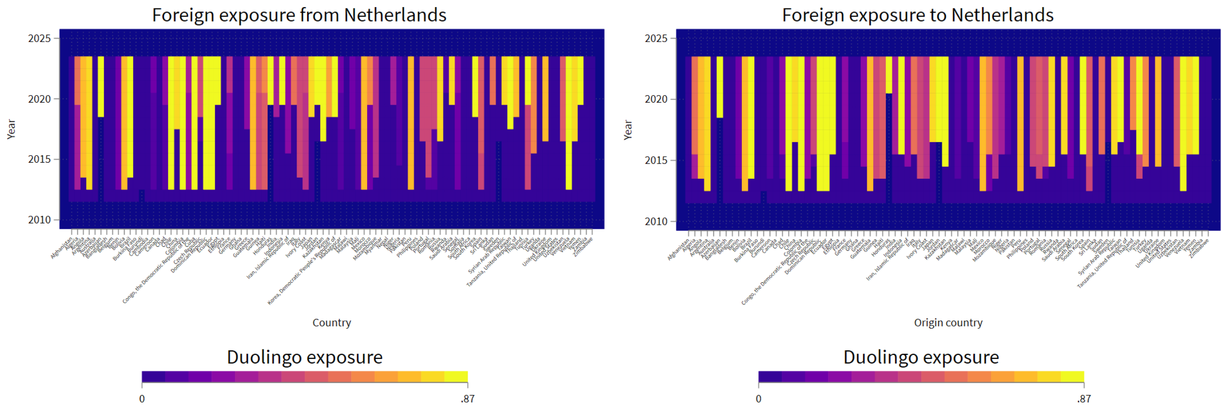

(A) From the Netherlands (B) To the Netherlands

Notes: Foreign exposure to Duolingo within an (a) origin and (b) destination country (the Netherlands) to/from all other countries with more than 10 million inhabitants (for legibility). The Netherlands is used as an illustration as it features multiple spoken language to and from which courses are rolled out at different times. Hence, there is variation in treatment timing across (a) destinations and (b) origins and in several instances treatment changes more than once. Brighter colors indicate a larger exposure.

FIGURE 2.B.13: Domestic Duolingo Exposure by Origin Country over Time (2012-2023)

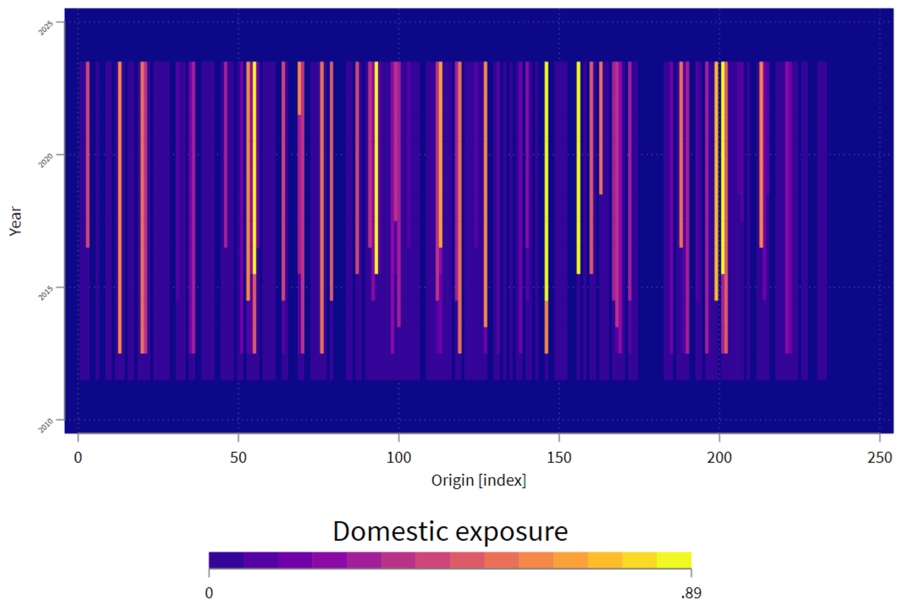

Notes: Dyadic domestic exposure to Duolingo DL_{oot} between 2012 and 2023. Brighter colors indicate a larger exposure.

2.C Google Trends

Google is the most-used search engine, with a global market share of about 90% between 2009 and 2024 (Allcott et al., 2024; Statcounter, 2024). Google provides *Google Trends*, a platform which enables users to query the relative search interest of a search term relative to all search activity on *Google Search*. Users can query the Google Trends Index (GTI), which is a measure of relative search intensity for a search term (i) by region for a given time period or (i) over time for a given region.[40][41] Importantly, it is not possible to directly query the relative search interest in a search term over time across countries. The relative search interest is normalized to 100 for the highest relative intensity within a query of type (i) or (ii), and all other data points get an integer score 0-100 relative to the highest relative intensity.

In the following, I discuss how, despite these limitations, a panel dataset measuring relative search intensity can be constructed. I denote (i) the interest by region as $GTI^{\widetilde{T}}_{o(2006-2022)}$ and (ii) the interest over time for a given geographic region (e.g. the whole world, a country, or a subnational region) as $GTI^{\widetilde{T}}_{\widetilde{o}t}$. Here, the available regions are 240 countries and territories. T is the term or topic, o is the geographic region of interest, and t is the time period of interest. The variables with a tilde indicate that search interest is not scaled *across* that dimension, whereas the absence of a tilde indicates that it is not scaled. For example, $GTI^{\widetilde{T}}_{\widetilde{o}t}$ is obtained through querying the GTI for every combination of origin region and search term. Hence, it is normalized to 100 for every origin-search term combination, and is uninformative about the relative search interest across origins and terms. Using the interest by regions and across time for the same search terms, geographic areas and time period, one can construct and index that is normalized across geographic regions and time: $GTI^{\widetilde{T}}_{ot} = \frac{1}{100} * GTI^{\widetilde{T}}_{\widetilde{o}t} \times GTI^{\widetilde{T}}_{o(2006-2022)}$. This enables us to compare relative search intensity across regions over time.[42] To proxy relative search interest for the search term Duolingo and its commonly used transliterations (in Arabic, Cyrillic, Japanese, Korean, and Mandarin scripts), I use $GTI^{Duolingo}_{ot}$. To proxy relative search interest for languages, I use $GTI^{\widetilde{T}}_{ot}$, which is *not* scaled across languages. The reason for this is that if one would scale across languages, the full variation in GTI^{T}_{ot} would be driven by languages with large absolute levels of search interest, limiting variation to mostly English and Spanish. I query search interest for all Topics that are used in Duolingo courses, all other languages with more than 50 million speakers according to the 2022 Ethnologue, and the following smaller European languages: Bulgarian, Lithuanian, Albanian, Latvian, Estonian, Slovak, Slovenian, Serbian, and Croatian. This gives a sample of 65 languages. Although I do not use it for the event studies in section 2.5, in the following section I explain how to obtain the interest scaled across search terms GTI^{T}_{ot} using a process called *Anchorbanking*.

Anchorbanking across terms and topics

As one can query up to five terms or topics, one can identify the normalized relative search intensity across these terms or topics. In case one has a set of terms or topics $\mathcal{T}$ exceeding five, one has to query the search terms in overlapping folds of 5 (the first fold includes the first five terms, the second fold includes the fifth to the ninth term). After querying, one rescales the GTI of the terms from the second fold onwards using the ratio of the time-averaged Indices of the overlapping term in the first over that in the second fold, and repeats this procedure for all subsequent folds. After this procedure all the GTIs are normalized to the highest value in the first fold. Furthermore, as long as one has topics and terms all across the distribution of relative search intensity, this also allows circumventing the rounding problem between terms. For this it is not sufficient to just query with overlap, but also to re-order after rescaling and repeating the

[40]Search terms can be simply a set of words, or a Topic. Google Trends *topics* are language-agnostic and include synonyms and common misspellings. This has the large advantage that it captures search behavior of without the need of translating and accounting for different grammatical forms.

[41]The temporal frequency available on Google Trends depends on the period of interest. If the period of interest is more than 5 years, the frequency is monthly. If it is between 8 months and 5 years, it is weekly. If it is shorter than 8 months, the frequency is daily. As I am interested in longer periods, for my purpose the monthly frequency suffices and I always query time series between 2006 and 2022.

[42]This approach is still limited by the rounding of GTI on integers. Hence, regions with less interest of around two orders of magnitudes smaller than the most interested region are strongly subject to rounding errors. This may lead to noisy results for low interest regions.

procedure of querying with overlap and rescaling. For my purpose, repeating this process two or three times suffices to obtain a distribution of GTI_{ot}^{T} that barely changes upon another repetition.

Figure 2.C.1a shows the global interest in Duolingo over time $GTI_{world,t}^{Duolingo}$. Before its existence, search interest was virtually zero, increasing rapidly between 2013 and 2015. Figure 2.C.1b shows the aggregated interest in languages using the sum of the fully scaled search interest $GTI_{ot}^{languages} = \sum_T GTI_{world,t}^{T}$, showing a strong increase from 2016 onwards.

FIGURE 2.C.1: Worldwide Google Trends Index for Duolingo and Languages over time

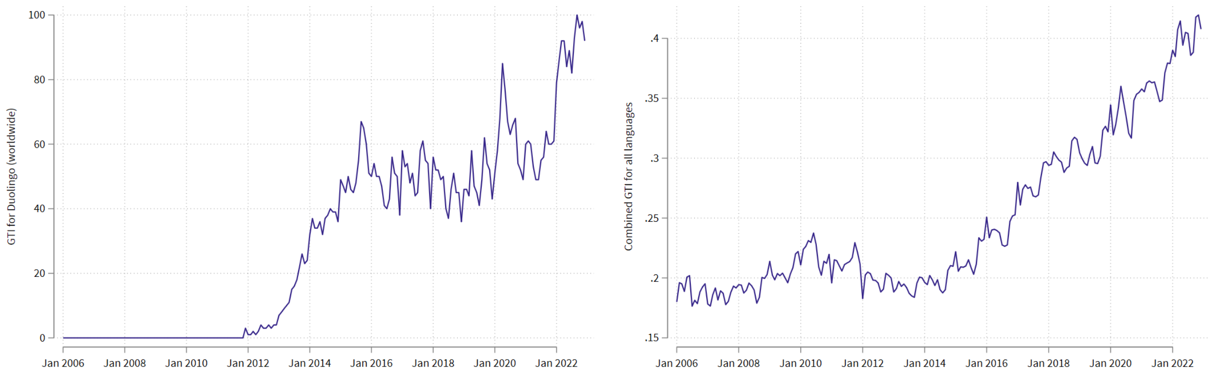

(A) Duolingo (B) Languages

Notes: Global relative search intensity for (a) Duolingo and its transliterations and (b) relative search intensity for all 65 queried languages. Data obtained by repeatedly querying Google Trends.

2.D Additional results

2.D.1 TOEFL and GRE Test Scores

FIGURE 2.D.1: The Effect of Duolingo Rollout on Component Scores of the English language (TOEFL) test (2007-2021)

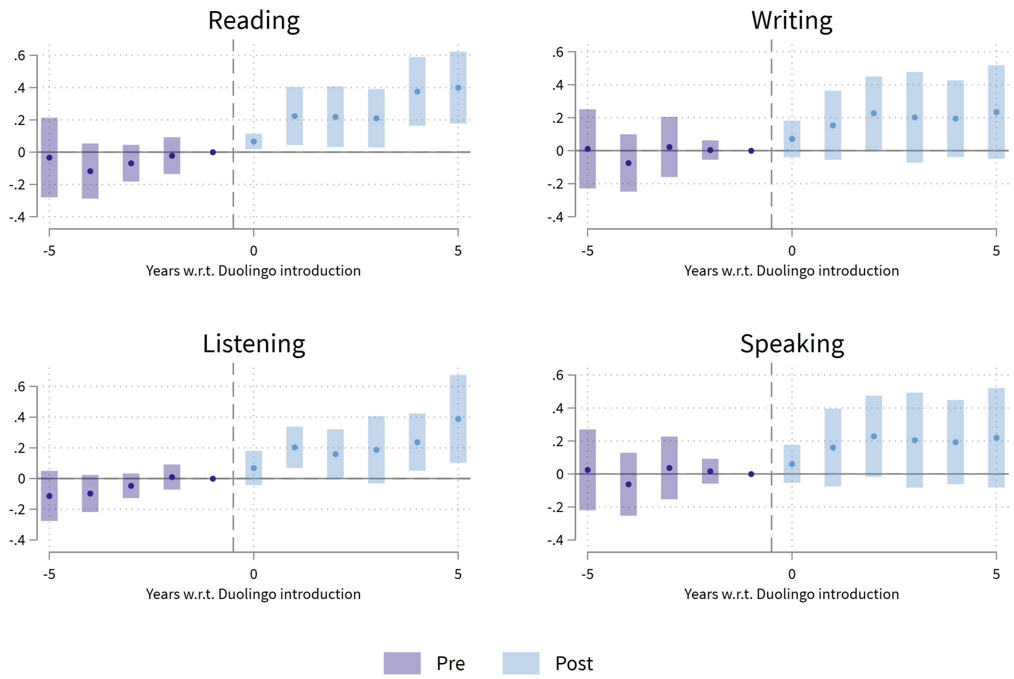

Notes: See notes to Table 2.5.2. The results in this table report results for countries with below-median scores in the respective aspect in 2010. N = 942 from 61 languages, of which 13 are treated.

FIGURE 2.D.2: The Effect of Duolingo Rollout on GRE Test Takers and Scores (2011-2022)

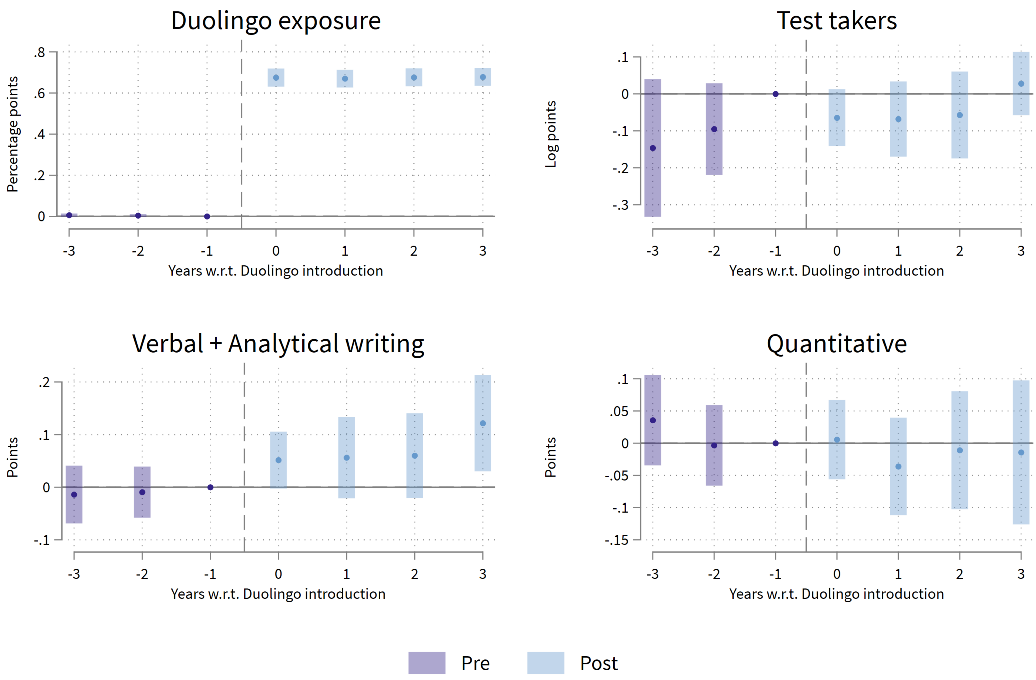

Notes: Results from Wooldridge (2023) event study estimators around increases in Duolingo exposure exceeding 20 percentage points. The panels report results using the following specifications and outcomes: a linear model of the Duolingo exposure as an outcome (upper left), (upper right) a Poisson model of the number of test takers, (lower left) a linear model of average scores for the verbal and analytical writing part of the GRE test and (lower right) a linear model of average scores for the quantitative parts. Effect sizes in the bottom two figures are standardized. The number of test takers and scores are reported as averages by country of citizenship. N = 1,510 from 148 distinct origin countries, of which 89 are treated. Shaded blue bars indicate 95% confidence intervals based on cluster-robust standard errors at the country of citizenship level. Data is obtained from the GRE annual reports between 2011 and 2022.

2.D.2 Does Duolingo crowd out traditional adult language learning?

It is a priori unclear how the introduction of Duolingo courses affect traditional language learning. On the one hand, potential learners may use Duolingo instead of traditional in-class language. On the other hand, Duolingo may spur language learning at basic levels and generate interest in destination language countries and culture and increase in-class course participation, particularly at higher proficiency levels. To provide some evidence on this, I turn to data on the number of German language course and exam takers at German language learning institutes (*Goethe* institutes) across more than 90 countries outside of Germany. By combining information from these language learning institutes with the staggered rollout of 9 Duolingo courses between 2014 and 2020, I can study whether the availability of low-cost language learning affected German language course and exam participation.

Participants of courses at Goethe institutes are predominantly high-skilled young adults (77% are aged 35 or below), almost 60% are female, and almost half are still in education. Education and cultural interest are the predominant reasons for learning, and only a minority share for migration-related reasons (Huber, Sommerfeld and Uebelmesser, 2022). I obtain the number of exams and course takers by country from Uebelmesser, Sommerfeld and Weingarten (2022) from 2007-2014 and collect the same information from Goethe's yearly reports between 2016 and 2022 on the global region level. In the latter, data is aggregated for 12 global regions in the latter, who house about 12 Goethe institutes each on average. Using the sum of registrations and exam takers by country in the year 2014, I construct a matrix X_{rc} of weights of every country c in every region r. Using this matrix and the time-varying Duolingo exposure by origin country for German, I construct a weighted exposure to Duolingo courses on the global region level. To study the impact of Duolingo availability on the number of course participants and test takers, I estimate a Poisson model event study using the Wooldridge (2023) estimator using increases in the Duolingo exposure of more than 20 percentage points. The results are reported in Table 2.D.3.

Although imprecisely estimated, the results in Figure 2.D.3 indicate that introduction of a relevant Duolingo course decreases the number of registrations, but the number of exams is even increasing if anything. This is consistent with language learners substituting Goethe courses with learning on Duolingo, but exams to a lesser extent. This is not surprising, given that Germany requires proof of language skills for some types of residence, and German higher education institutes for admission.

FIGURE 2.D.3: The Effect of Duolingo on Institutional German Learning

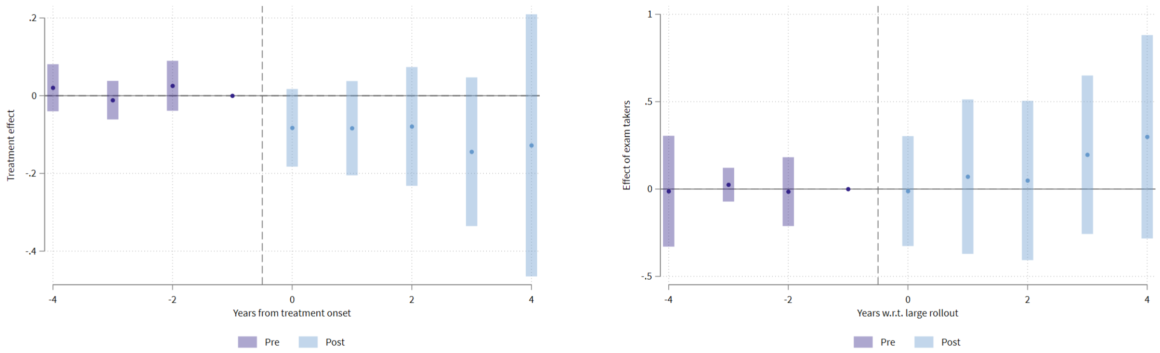

(A) Course participants

(B) Exam takers

Notes: Results from Wooldridge (2023) event study PPML estimators around increases in average region-level Duolingo exposure exceeding 20 percentage points. The text provides a discussion about how Duolingo exposure is constructed on the regional level. Standard errors are clustered at the origin region level. N = 192 with 12 unique origin regions, of which 6 are treated. Data on participants and test takers originate from Uebelmesser, Sommerfeld and Weingarten (2022) and the yearly reports (*Jahrbuch*) of the Goethe Institute.

2.D.3 Does low-cost language learning affect in-school instruction in the EU?

The availability of low-cost language learning may affect in-class instruction in various ways. First, the availability of a language course may foster Second, it may induce interest in the particular language among pupils, who continue to study the language further in class. To study whether these effects are at play, I resort to data on foreign language learning across the EU. The EU consistently reports numbers on the share of pupils learning specific foreign languages for three stages of education. English is by far the most learned foreign language across the EU: 73% of pupils learn it in primary education, 92% in lower secondary and 82% in upper secondary education. Hence, I show results with and without English included. Figure 2.D.4 shows the results. The upper panel suggests that across levels Duolingo exposure increases the number of pupils. Without English, the results are smaller in magnitude, but results still seem to be positive.

FIGURE 2.D.4: The Effect of Duolingo Courses on In-school Language Learning

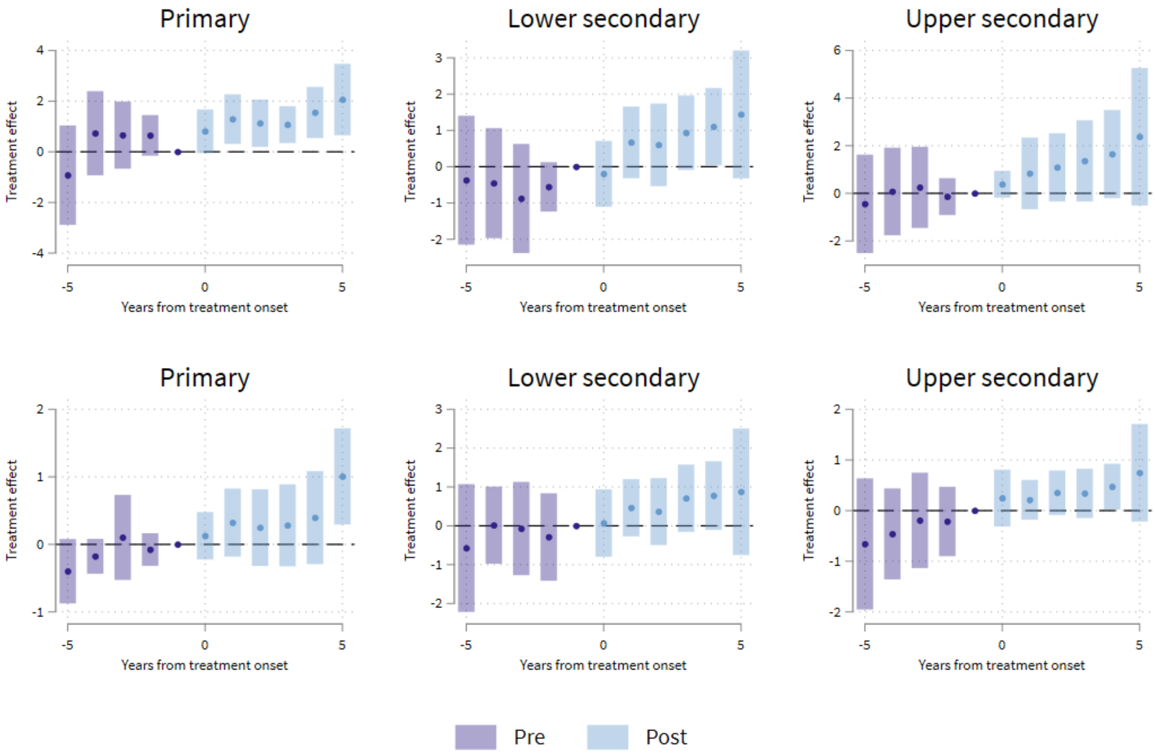

Notes: Results from linear Wooldridge (2023) event study estimators around increases in Duolingo exposure exceeding 50 percentage points. The Duolingo exposure is constructed by the maximum of the share of native-language speakers of the source language of a Duolingo course. I use the share of native language speakers in this case as school-aged children are unlikely to speak foreign spoken languages already. The dependent variable is the share of pupils studying a specific language. The columns report results for primary, lower secondary and upper secondary students separately. The upper row presents results including English, the bottom row excluding English. N = 9,890 from 34 countries and 28 languages, 879 country-language pairs of which 100 treated in the upper panel, N = 9,458 from 34 countries and 27 languages, 846 country-language pairs of which 84 treated in the bottom panel. Shaded blue bars indicate 95% confidence intervals based on two-way cluster-robust standard errors at the country and language level. Data is obtained from Eurostat table educ_enrl1tl for 2007–2012 and educ_enrllng1 for 2012–2022. In case a country-language pair has an observation in both in 2012, I take the simple unweighted average between both.

2.D.4 Migration Intentions

Pattern over time

FIGURE 2.D.5: Share of GWP Respondents Desiring to Emigrate

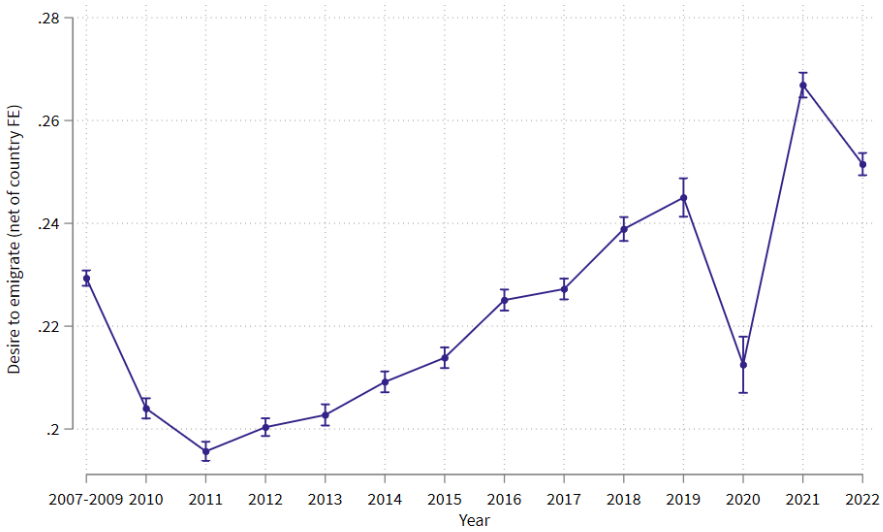

Notes: Share of individuals answering positively to the question on desiring to emigrate over time, net of country fixed effects. The graph reports coefficients of a regression of a dummy for desiring to emigrate on country fixed effects and year fixed effects. I report the year fixed effects, adding back the constant. Because of the limited number of countries visited by GWP in the first years, I aggregate the years 2007–2009 in a single category. 95% confidence intervals are reported based on standard errors are clustered at the country level. Data are from the Gallup World Polls between 2007 and 2022. N = 2,014,359.

Control Function Approach

Although it is plausible that Duolingo courses were not rolled out in anticipation of trends in international migration, one may still be concerned about the presence of unobserved confounding between the availability of DDuolingo courses and bilateral migration (intentions) over time. To further mitigate such concerns, I construct two distinct instruments for foreign exposure to Duolingo and re-estimate the main results using a control function approach.

As demand for language learning requires source language learners being interested in target languages, Duolingo's incentives to roll out languages courses are increasing in the product of the source- and target language size. To see this, assume there are two distinct source languages S_1 and S_2 with 2 and 1 million speakers respectively and two distinct target languages T_1 and T_2 with 2 and 1 million speakers respectively. I assume that the probability that a source language speaker is interested in a target language is proportional to its applicability, e.g. the number of speakers. Consequently, the potential market size for a $S_1 \rightarrow T_1$ course is four times as large as a $S_2 \rightarrow T_2$ course. Hence, the potential market size of a course varies on the bilateral level and is an increasing function of the product of the number of speakers. As witnessed in Table 2.4.1, Duolingo indeed prioritized courses between languages with many speakers, after controlling for the number of source- and target language speakers. Moreover, the slope of the number of learners by course per source language speaker is also increasing in the number of target language speakers, as witnessed in Table 2.B.1. Hence, I instrument the rollout of Duolingo courses by the interaction between the log number of speakers of the source and target language, interacted with a linear time trend starting in 2012 (the introduction year of the first courses). I opt to construct the instrument using the log of speakers rather than the the absolute number, as otherwise the vast amount of variation is driven by country-pairs speaking the largest two languages: one speaking English and the other speaking Mandarin.

$$Z_{STt} = \begin{cases} Z_{STt} = log(N^S)log(N^T)t & \text{if } t \geq 2012 \\ 0 & \text{otherwise} \end{cases} \tag{2.10}$$

The second instrument is based on Duolingo's propensity to roll out courses to the same source and target languages as existing courses. For example, if Duolingo rolls out courses from a specific source languages, they have language-specific knowledge that facilitates the development of further courses using that language as a source language. The same logic applies to target languages. To construct a bilateral instrument for course roll-out, I interact the source language and target language propensity, omitting the focal source- and target language. The latter would generate a mechanical correlation between the instrument and the Duolingo exposure.

$$Z_{STt} = \begin{cases} \left(\sum_{S' \neq S} Duolingo_{ST}\right) \left(\sum_{T' \neq T} Duolingo_{ST}\right) t & \text{if } t \geq 2012 \\ 0, & \text{otherwise} \end{cases} \tag{2.11}$$

Here, $Duolingo_{ST}$ is a binary indicator for whether there is a Duolingo course from language S to language T by 2022. In both cases, I aggregate the instrument in similar vein as the main exposure to Duolingo:

$$Z_{odt} = \max_{S,T} \alpha_{oS} \alpha_{dT} Z_{STt} \tag{2.12}$$

Using either instrument, I estimate the following linear first stage:

$$DL_{odT} = Z_{odt} + \phi_{ot} + \theta_{dt} + \psi_{od} + \epsilon_{odt} \tag{2.13}$$

As the gravity model is a non-linear model, naive 2SLS estimation is invalid. Instead, one should resort to a control function procedure (Wooldridge, 2015). To implement this, I estimate equation 2.13, obtain the residuals $\hat{\epsilon}_{odt}$, and include these residuals in the baseline gravity model. Any variation in the potentially endogenous regressor unexplained by the instrument and the fixed effects is absorbed by this residual. Hence, inclusion of the residual in the second stage controls for the endogenous variation.

The remaining identification assumption behind the first control function strategy is that trends in migration aspirations between countries both speaking widely spoken languages have followed similar trends than between countries where one of the languages spoken is smaller, after controlling for three-way fixed effects. This implies that differential trends for countries speaking source- and target languages of varying sizes are absorbed by the origin-year and destination-year fixed effects. The remaining identification assumption behind the second control function approach is that outcome trends would have followed similar paths between country pairs both speaking languages that have other Duolingo courses than country pairs that do not, in absence of strong Duolingo exposure between the treated country pair.

Results. Table 2.D.1 shows the results of the control function approach. Columns 1 and 3 shows that both instruments strongly predict Duolingo exposure, but that the language size instrument is particularly strong. Moreover, I find that the IV results are in line with the OLS results, only somewhat larger. The results suggest that, if anything, the endogeneity bias is negative, and that the effect size exceeds 80%.

The Role of Language Requirements

TABLE 2.D.1: Control Function Estimates of the Effect of Duolingo of Migration Aspirations

	(1)	(2)	(3)	(4)
	Language size		Course propensity	
	DL_{odt}	$\frac{M_{odt}}{M_{oot}}$	DL_{odt}	$\frac{M_{odt}}{M_{oot}}$
Z_{odt}	0.001***		3.609***	
	(0.000)		(0.776)	
DL_{odt}		0.602***		1.141***
		(0.189)		(0.366)
Control function		-0.351*		-0.855**
		(0.191)		(0.367)
Observations	121698	121477	121698	121477
Unique origin countries	153	153	153	153
Unique destination countries	196	196	196	196
Unique dyads	10,527	10,527	10,527	10,527
Kleibergen-Paap F-statistic	444.3		21.7	
Origin-destination FE	✓	✓	✓	✓
Origin-year FE	✓	✓	✓	✓
Destination-year FE	✓	✓	✓	✓
Estimator	OLS	PPML	OLS	PPML

Notes: Control function estimation of the effect on Duolingo exposure on migration odds for two different instruments. Odd columns report estimates from a linear first stage and even columns report the results of the three-way gravity model estimated by PPML. Columns (1) and (2) estimate this procedure using the instrument based on language size as described in equation 2.10 and columns (3) and (4) estimate this procedure using the instrument based on the propensity to roll out languages to particular courses as described in equation 2.11. For notes on the data and estimation sample, see notes to Table 2.6.1. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level.

As Duolingo effectively enables low-cost language learning without certification, this could reduce migration barriers more for destination countries without language requirements. (which typically require a costly certificate). I therefore draw on the MIPEX database, which provides information on migration policy for 56 destination countries. I retrieve the MIPEX indicator for language requirements for permanent residence (originally taking values 0, 50, or 100) and convert it to a 0-2 scale. Figure 2.D.2 shows the effect of the interaction of Duolingo exposure with the MIPEX indicator. I find that the effect is strongest for those countries without requirements, although the interaction is not statistically significant.

TABLE 2.D.2: Heterogeneity of Results by Destination-country Language Requirements

	(1) $\frac{M_{odt}}{M_{oot}}$	(2) $\frac{M_{odt}}{M_{oot}}$	(3) $\frac{M_{odt}}{M_{oot}}$
DL_{odt}	0.374***	0.237***	0.286***
	(0.080)	(0.080)	(0.097)
DL_{odt} × Permanent residence language requirements (0-2, MIPEX)			-0.056
			(0.042)
Observations	123263	41699	41699

PPML regressions based on the sample and specification of column 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Columns 2–4 restrict the sample to those 56 destination countries where MIPEX is available. Column 3 introduces an interaction between an indicator for language requirements for permanent residency, taking values 0 (no language requirements), 1 (some language requirements) or 2 (strict language requirements). The number of observations are slightly different from Table 2.6.1 because of 3 countries not speaking any language present on Duolingo required for Column 7 and 8. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.D.3: The Effect of Duolingo on Total Emigration

	(1) $\frac{M_{ot}}{P_{ot}}$	(2) $\frac{M_{ot}}{P_{ot}}$	(3) $\frac{M_{ot}}{P_{ot}}$	(4) $\frac{M_{ot}}{P_{ot}}$
$DL_{ot}^{foreign}$	0.036	0.042		
	(0.037)	(0.037)		
$DL_{ot}^{domestic}$		-0.035*		
		(0.019)		
$DL_{ot}^{foreign}$ (weighted)			0.048**	0.047**
			(0.020)	(0.020)
$DL_{ot}^{domestic}$				-0.029
				(0.018)
Observations	1757	1757	1757	1757
Average dependent variable	0.219	0.219	0.219	0.219

Notes: OLS regression with country- and year fixed effects of migration rates on foreign- and domestic Duolingo exposure. Columns 1 and 2 use unweighted averages of Duolingo exposure, whereas columns 3 and 4 uses a weighted average measure of foreign exposure to Duolingo, using the bilateral stock of migrants in 2005 as weights. See notes to Table 2.6.1 for the sample and data. Standard errors reported in parentheses are clustered on the country of origin level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Event study

FIGURE 2.D.6: Change in Duolingo Exposure around Increases in Exposure Exceeding 50 pp

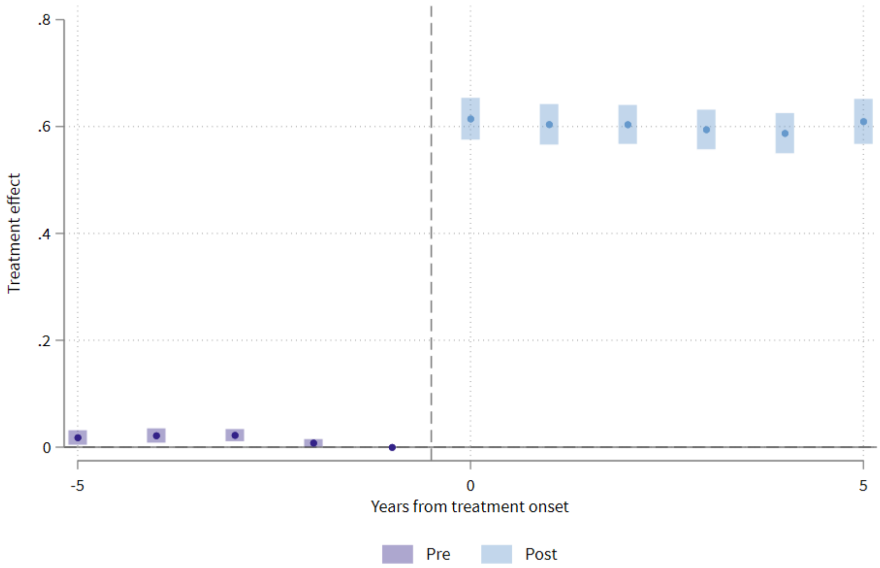

Notes: OLS regression of the heterogeneity-robust Wooldridge (2023)-estimator with three-way fixed effects of the treatment exposure on a binary indicator for whether a origin-destination pair has experienced an increase in foreign Duolingo exposure exceeding 50 percentage points. See notes to Table 2.6.1 for information on the data and sample.

FIGURE 2.D.7: Event Study around Large Increases in Domestic Duolingo Exposure

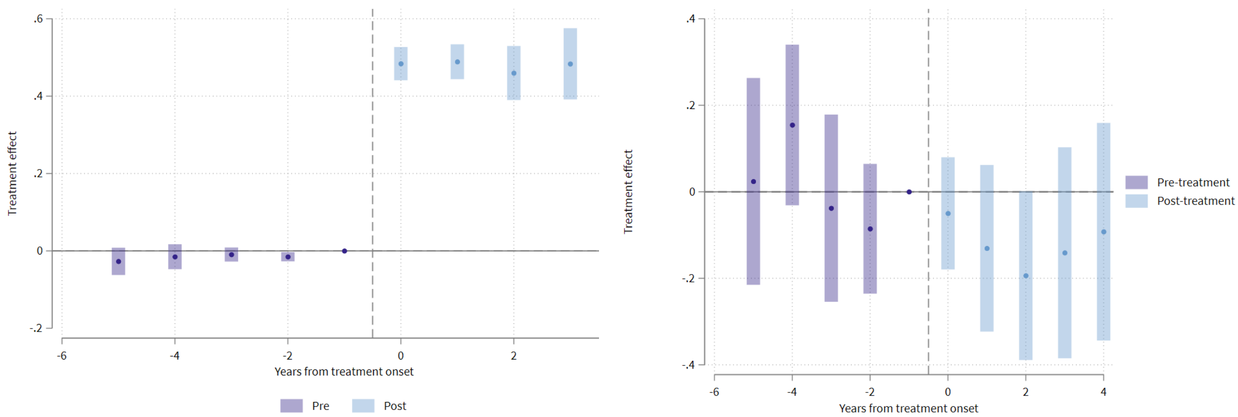

(A) Domestic Duolingo Exposure (B) Migration Aspiration Odds

Notes: (a) OLS regression of the heterogeneity-robust Nagengast and Yotov (2023)-estimator with three-way fixed effects of the treatment exposure on a binary indicator for whether a country pair has experienced increase in domestic Duolingo exposure exceeding 50 percentage points. See notes to Table 2.6.1 for information on the data and sample. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. (b) PPML regressions of the heterogeneity-robust event study estimator by Nagengast and Yotov (2023) of migration aspiration odds on a binary indicator for whether an origin country pair has experienced increase in domestic Duolingo exposure exceeding 50 percentage points, including origin-destination pair and destination year fixed effects. An event is defined as an increase in Duolingo exposure of more than 50 percentage points. Estimates are shown for the 5 years before and after an event. Shaded blue bars indicate 95% confidence intervals based on standard errors clustered at native language level. See notes to Table 2.6.1 for information on the data and sample.

Alternative Treatment Definitions

TABLE 2.D.4: Testing Monotonicity of the Effect

	(1)	(2)
$0 < DL_{odt} \leq 0.25$	-0.066	0.004
	(0.048)	(0.054)
$0.25 < DL_{odt} \leq 0.5$	0.024	0.158***
	(0.058)	(0.058)
$0.5 < DL_{odt} \leq 0.75$	0.132**	0.180**
	(0.055)	(0.070)
$0.75 < DL_{odt} \leq 1$	0.170***	0.297***
	(0.061)	(0.080)
$0 < DL_{oot} \leq 0.25$	0.011	
	(0.068)	
$0.25 < DL_{oot} \leq 0.5$	-0.139	
	(0.103)	
$0.5 < DL_{oot} \leq 0.75$	0.077	
	(0.117)	
$0.75 < DL_{oot} \leq 1$	-0.369***	
	(0.105)	
Observations	123263	123263
Origin-year fixed effects		✓

Notes: Gravity model estimated by PPML without (column 1) and with (column 2) origin-year fixed effects. See notes to Table 2.6.1 for information on the data and sample. Instead of numerical values for both treatments, these results include binary indicators for four bins of treatment intensity per exposure variable and omits the regressors reported in Table 2.6.1. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.D.5: Using Official Target Languages in the Destination

	(1)	(2)
DL^{off}_{odt}	0.167***	0.198***
	(0.056)	(0.061)
DL^{off}_{oot}	-0.161	
	(0.376)	
Observations	123655	123655
Origin-destination FE	✓	✓
Origin-year FE		✓
Destination-year FE	✓	✓

PPML regressions based on the sample and specification of column 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Exposure measures are calculated by only taking into account target languages that are official languages in the destination country. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Robustness on Empirical Approach

TABLE 2.D.6: Different ways of clustering standard errors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Level of clustering:	Pair		Origin & Destination		Main origin & destination language		Main Duolingo origin & destination language	
DL_{odt}	0.267***	0.373***	0.267***	0.373***	0.267***	0.373***	0.267***	0.373***
	(0.065)	(0.052)	(0.065)	(0.080)	(0.060)	(0.094)	(0.058)	(0.109)
DL_{oot}	-0.225***		-0.225		-0.225		-0.225	
	(0.075)		(0.160)		(0.157)		(0.169)	
Observations	123180	123180	123180	123180	123180	123180	123180	123180
Number of clusters	10641	10641	153	153	63	63	22	22
Number of clusters (2)			194	194	70	70	30	30
Origin-destination FE	✓	✓	✓	✓	✓	✓	✓	✓
Origin-year FE		✓		✓		✓		✓
Destination-year FE	✓	✓	✓	✓	✓	✓	✓	✓

PPML regressions based on the sample and specification of column 1 and 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Standard errors, reported in parentheses, are clustered on different levels. Columns 1 and 2 cluster on the country pair level, 3 and 4 on the origin- and destination level (as Table 2.6.1), 5 and 6 on the main origin- and main destination language and 7 and 8 on the main origin- and main destination language that is available in any course on Duolingo. The number of observations are slightly different from Table 2.6.1 because of 3 countries not speaking any language present on Duolingo required for Column 7 and 8. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.D.7: Omission of Exposure Contribution in Countries with Most Speakers by Language

	(1)	(2)	(3)	(4)	(5)	(6)
	Omission of contribution to exposure in:					
	Origins with most source speakers		Destinations with most target speakers		Both	
DL_{odt}	0.276***	0.421***	0.326***	0.346***	0.388***	0.408***
	(0.059)	(0.083)	(0.100)	(0.123)	(0.080)	(0.123)
DL_{oot}	-0.209		-0.254		-0.240	
	(0.155)		(0.169)		(0.162)	
Observations	123180	123180	123180	123180	123180	123180
Origin-destination FE	✓	✓	✓	✓	✓	✓
Origin-year FE		✓		✓		✓
Destination-year FE	✓	✓	✓	✓	✓	✓

PPML regressions based on the sample and specification of column 1 and 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Columns 1 and 2 includes alternative measures of Duolingo exposure, excluding the contribution of the source language in the origin country with most speakers, for every language. Columns 3 and 4 includes alternative measures of Duolingo exposure, excluding the contribution of the target language in the destination country with most speakers, for every language. Columns 5 and 6 includes alternative measures of Duolingo exposure, excluding both types of contributions. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.D.8: Controlling for Origin-destination-nest-year Fixed Effects

Nest:	(1) none	(2) WB region	(3) WB income	(4) EU members	(5) All three
DL_{odt}	0.365***	0.264***	0.346***	0.363***	0.231***
	(0.080)	(0.090)	(0.098)	(0.065)	(0.085)
Observations	111251	101005	103788	111198	95942
Number of Fixed Effects	14120	21500	18389	15884	30484
Number of groups		7	4	2	7+4+2

PPML regressions based on the sample and specification of column 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. The sample sample is somewhat smaller than the baseline sample of Table 2.6.1 due to missing information of WB income groups for some jurisdictions. Column 1 shows that Every column additionally includes origin-nest-year fixed effects where the nests are given by the column header Column 4 uses the three. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Robustness on Sample

TABLE 2.D.9: Omitting a High-income Native-English Destination Country at a Time

Omission of destination:	(1) AU	(2) CA	(3) UK	(4) US	(5) IE	(6) All
DL_{odt}	0.349***	0.349***	0.359***	0.327***	0.349***	0.295***
	(0.069)	(0.071)	(0.071)	(0.067)	(0.068)	(0.082)
Observations	96448	96442	96434	96435	96746	90429

PPML regressions based on the sample and specification of column 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Columns 1-5 each remove a destination country, column 6 removes all countries at once. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.D.10: Different Sample Periods

Time period:	(1) 2008–2022	(2) 2009–2022	(3) 2010–2022	(4) 2011–2022	(5) 2007–2019
DL_{odt}	0.349***	0.328***	0.311***	0.277***	0.332***
	(0.077)	(0.077)	(0.076)	(0.077)	(0.061)
Observations	117919	111894	102948	93437	93611

PPML regressions based on the sample and specification of column 2 of Table 2.6.1. See notes to Table 2.6.1 for the estimation strategy, data and sample. Every column restricts the sample to the years in the column header. Standard errors reported in parentheses are two-way clustered: on the country of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE 2.D.8: Omission of a Duolingo course at a time

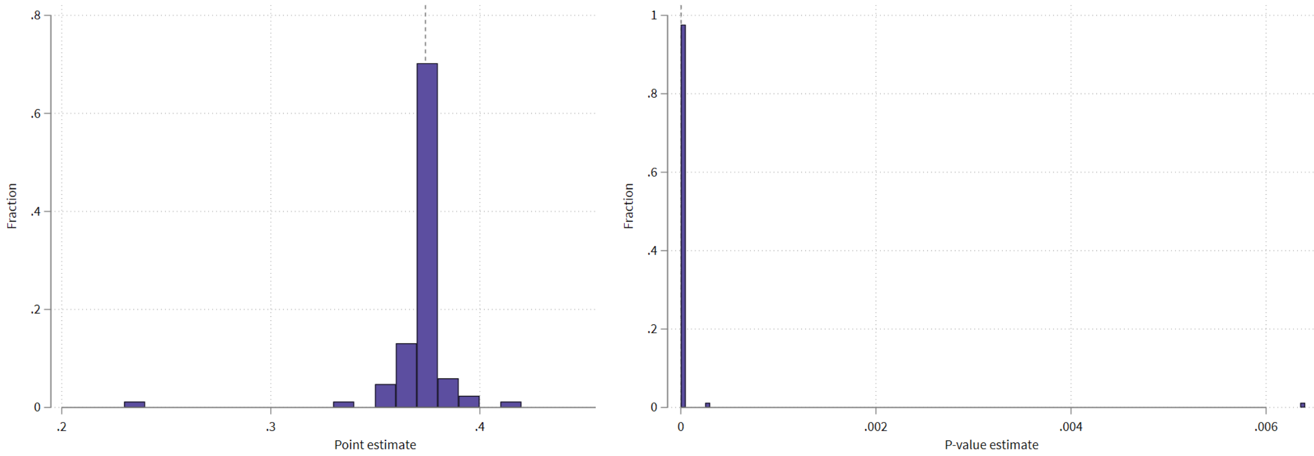

(A) Point estimates (B) p-values

Notes: Re-estimation of the results reported in column 2 of Table 2.6.1, omitting a language course at a time in the construction of DL_{odt}. (A) shows the point estimates on foreign Duolingo exposure, (B) shows the p-values of the coefficient.

FIGURE 2.D.9: Omission of An Origin Country at a Time

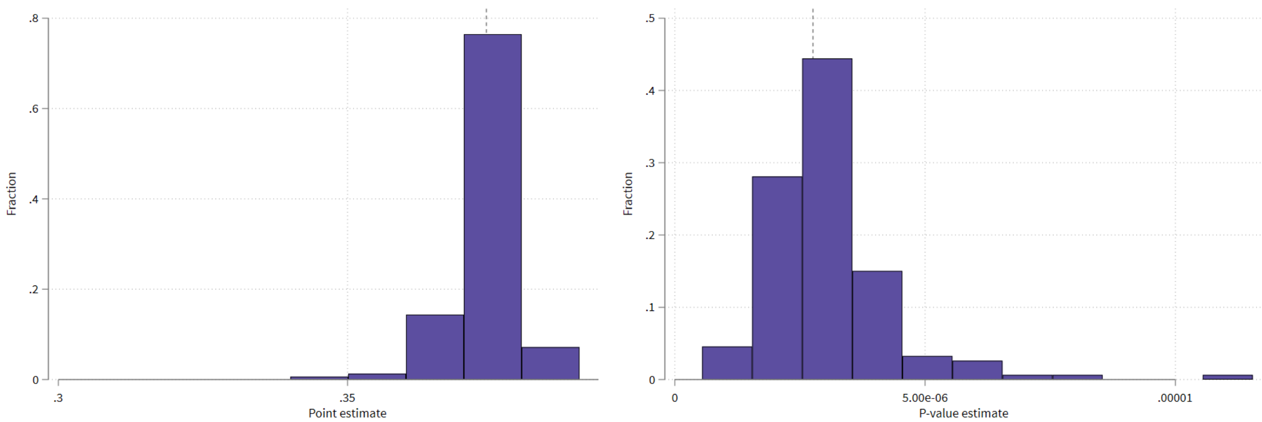

(A) Point estimates (B) p-values

Notes: Re-estimation of the results reported in column 2 of Table 2.6.1, omitting an origin country at a time. (A) shows the point estimates on foreign Duolingo exposure, (B) shows the p-values of the coefficient.

FIGURE 2.D.10: Omission of A Destination Country at a Time

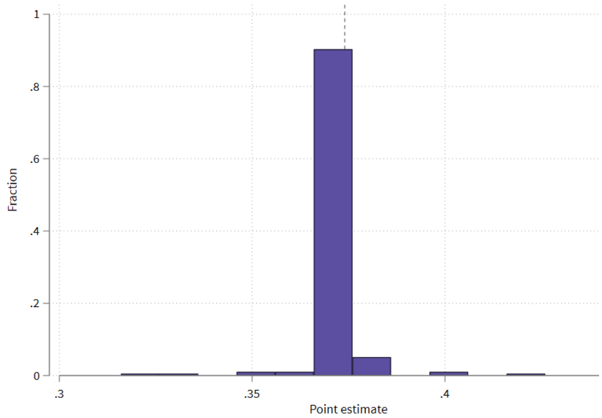

(A) Point estimates

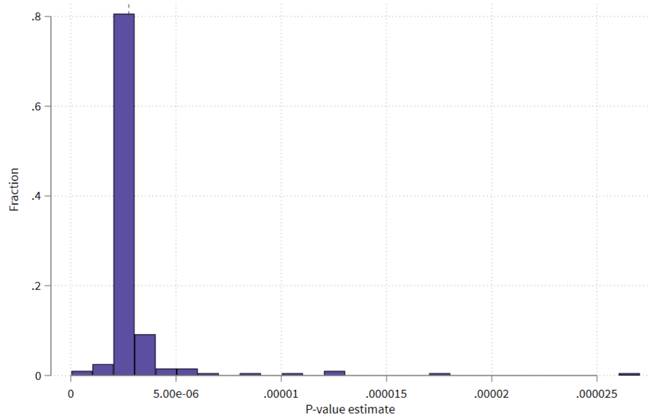

(B) p-values

Notes: Re-estimation of the results reported in column 2 of Table 2.6.1, omitting a destination country at a time. (A) shows the point estimates on foreign Duolingo exposure, (B) shows the p-values of the coefficient.

2.D.5 Migrant Integration in the EU

Identification of pre- and post-arrival effects

FIGURE 2.D.11: Identification of Pre- and Post-treatment Exposure

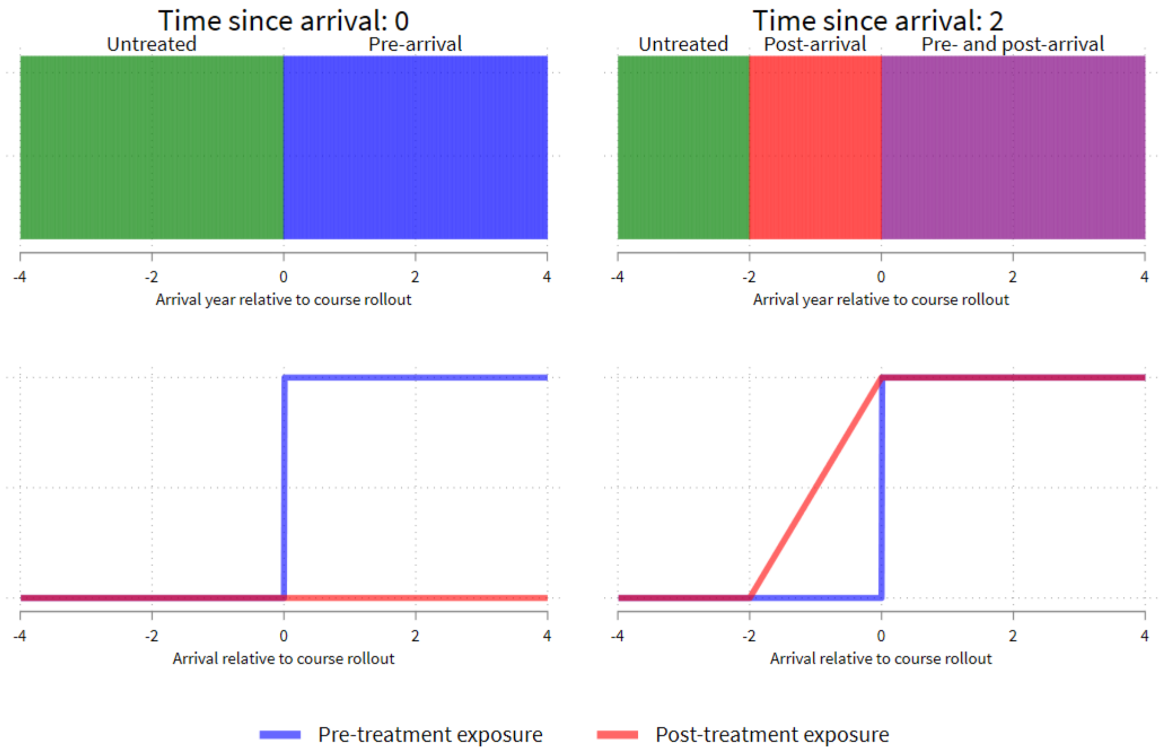

Notes: This figure illustrates the identification of the pre-arrival and post-arrival effects of Duolingo availability on migrant outcomes. The x-axis in all 4 panels represents the time of arrival of the migrant relative to the roll-out of a language course. The left column considers a migrant interviewed within the first year of arrival and the right column a migrant interviewed two years after arrival. The upper row identifies three regimes: untreated, only post-arrival learning and both pre- and post-arrival learning. The bottom row shows the pre- and post-arrival treatment intensity.

Attenuation bias due to aggregation

As the group-level exposure is aggregated from country-level exposure, it is a noisy measure of individual-level availability of low-cost language learning. To quantify the degree to which this noise attenuates point estimates, I examine the extent of Figure 2.D.12 shows the scatterplot of the Duolingo exposure $DL_{o_g dt}$ and the migration-flow weighted within-group variance. In absence of measurement error all circles lie on the x-axis, if classical measurement error was maximal all lines would lay on the hump-shaped line. The figure shows that across average exposure levels there is considerably less intra-group variance than maximal. About three quarters of the variance in the Duolingo exposure by origin country, destination and arrival year is driven by origin country *group*, destination and arrival year. This suggest that attenuation bias plays only a small but not negligible role. Using the formula for attenuation bias under classical measurement error from Pischke (2007) I estimate that I underestimate the true effect by about 25%:

$$\beta \approx \frac{\sigma^2_{\Delta x} + \sigma^2_{\Delta u}}{\sigma^2_{\Delta x}} \hat{\beta} \approx \frac{5}{4} \hat{\beta} \tag{2.14}$$

Descriptives and additional results

TABLE 2.D.11: Descriptive Statistics of Main LFS Samples

	Language skills upon arrival (2021)			Upon arrival ($t-c=0$)			Full sample		
	mean	s.d.	N	mean	s.d.	N	mean	s.d.	N
Female	0.53	0.50	19341	0.53	0.50	51464	0.53	0.50	668737
Age	37.47	9.03	19341	31.34	9.30	51464	35.23	9.07	668737
Primary educated	0.40	0.49	19250	0.46	0.50	49598	0.37	0.48	655011
Secondary educated	0.29	0.45	19250	0.20	0.40	49598	0.29	0.45	655011
Tertiary educated	0.31	0.46	19250	0.34	0.47	49598	0.33	0.47	655011
Time since arrival	6.71	4.08	19341	0.00	0.00	51464	4.65	3.55	668737
Main activity: employment	0.62	0.49	19289	0.46	0.50	46127	0.59	0.49	631765
Pre-treatment Duolingo exposure	0.17	0.23	19279	0.20	0.26	50464	0.10	0.20	664591
Language upon arrival: at least advanced	0.31	0.46	19341	0.23	0.42	607	0.31	0.46	19341
Language upon arrival: at least intermediate	0.41	0.49	19341	0.32	0.47	607	0.41	0.49	19341
Language upon arrival: at least beginner	0.56	0.50	19341	0.51	0.50	607	0.56	0.50	19341
Reason: employment, job before arrival	0.16	0.37	18911	0.23	0.42	2801	0.18	0.38	64780
Reason: employment, no job before arrival	0.21	0.41	18911	0.13	0.34	2801	0.22	0.41	64780
Reason: family	0.37	0.48	18911	0.36	0.48	2801	0.34	0.47	64780
Reason: education	0.10	0.30	18911	0.15	0.36	2801	0.10	0.30	64780
Reason: refugee	0.10	0.30	18911	0.03	0.18	2801	0.10	0.30	64780

Notes: Descriptive statistics of the three samples used for the EU LFS.

FIGURE 2.D.12: Measurement Error due to Group-level Duolingo exposure

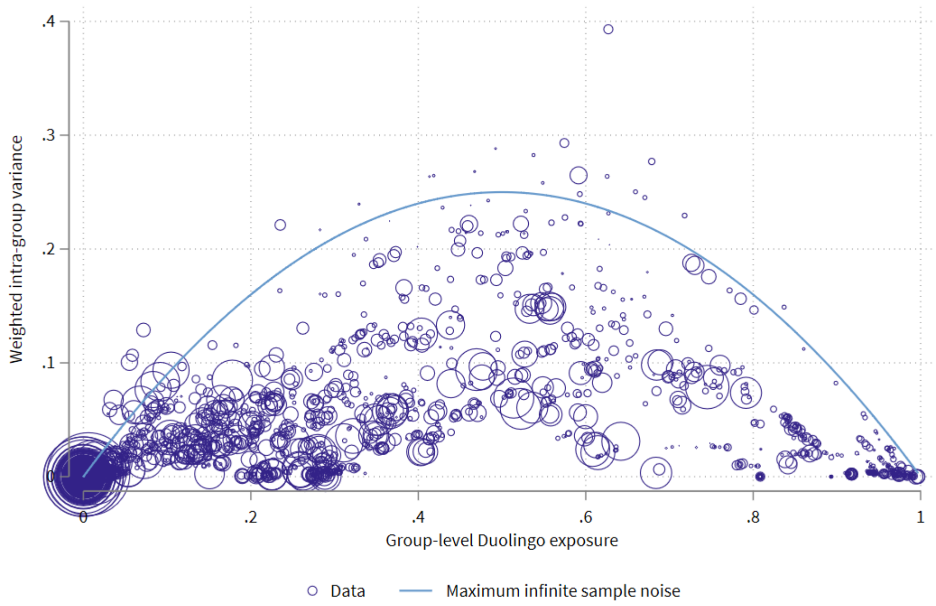

Notes: Scatter plot between the intra-origin group variation and the mean level Duolingo exposure on the origin group by destination by year of arrival level. N = 3,643. Empty circles' surface area is proportional to the weighted number of observations in the full estimation sample.

FIGURE 2.D.13: Event study of Language Skills upon Arrival around the Large Increases in Duolingo Exposure

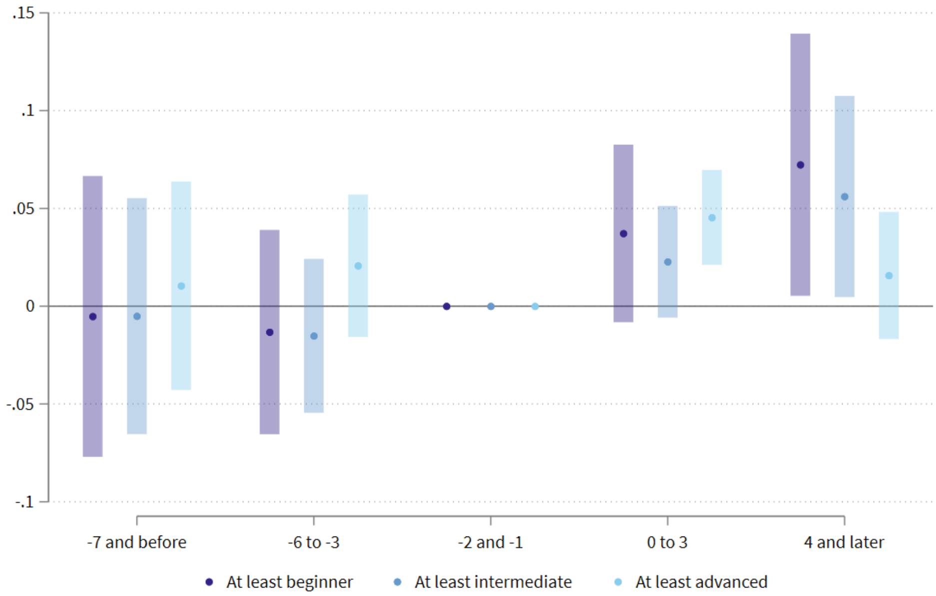

Notes: OLS event study of language skills around large increases (at least 50 percentage points) in Duolingo exposure. Shaded blue bars indicate 95% confidence intervals based on two-way cluster-robust standard errors at the origin group and destination country level.

2.D.6 Migrant Language Skills and Integration in the US

I follow the same empirical strategy as section 2.7. However, because there is no variation in target languages across the US, the main equation can not be estimated with origin-by-year fixed effects for the US. Variation in origin-specific "cohort quality" may be large (Borjas, 1985), and entry wages of US migrants in more recent cohorts is decreasing (Borjas, 2015). Hence, these estimates would be particularly sensitive to changes in cohort quality among migrants from specific native languages areas. Careful assessment of differential exposure between treated and untreated countries of origin before the availability of relevant Duolingo courses is crucial.

Data

The American Community Survey (ACS) is a large yearly household survey fielded by the US Census Bureau among more than 3 million people each year. Respondents are randomly selected each year and are legally obliged to answer, providing information about themselves and other members of their household. The ACS collects information of a range of relevant demographic characteristics and economic outcomes, as well as information on an individuals' migration history and country of birth, the language spoken at home and a self-assessment of the contemporaneous language skills of all household members, with the following answer options: Only English, very well, well, not well, does not speak English.[43] Although the language spoken at home could be used to construct an exposure measure, I choose to use the country of birth as the language spoken at home could be endogenous. I restrict the sample to those who arrived on ar

[43]Self-reported language skills in the ACS have been shown to strongly correlate to actual language proficiency (US Census Bureau, 2015). However, some studies have found that active learners under-assess their learning gains, e.g. Ma and Winke (2019)

after age 18 and those currently 59 or younger who immigrated to the US in the past 10 years. In line with the other datasets, I start the analysis in 2007 until the last available year, 2022.

FIGURE 2.D.14: Distribution of English Language Skills Upon Arrival in the US

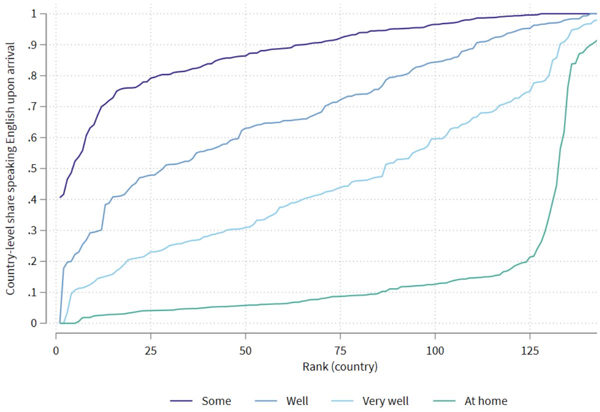

(A) Unweighted

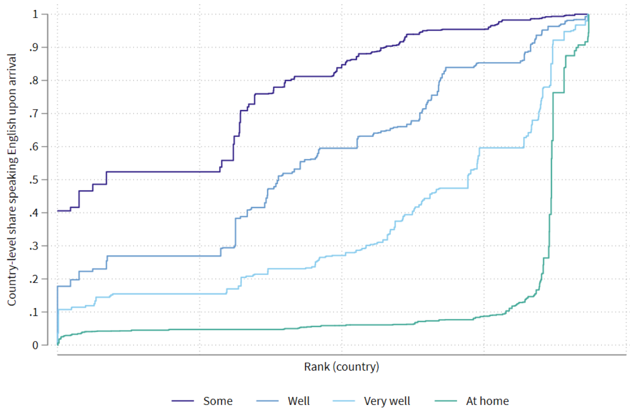

(B) Weighted

Notes: Distribution of the level of language skills upon arrival in the US between 2007 and 2022, by origin country (left) and weighted by cohort size (right). The strong difference between both graphs shows that immigrants from several large immigrant origin countries have poort English skills (such as Mexico).

Figure 2.D.14 shows the cumulative distribution of levels of language skills among those within the first year of arrival. This shows that many respondents lack good English language skills before arriving to the US. As discussed in section 2.7, the identification strategy in the US assumes parallel trends in outcomes between treated and untreated origin countries. A potential risk to this identification assumption could be that language skills for English skills have trended differently for example for countries that speak more widely spoken languages (and who are more likely to have received a Duolingo course) than for less widely spoken languages. Figure 2.D.15 shows that pre-trends before large increases in Duolingo exposure are small. However, the results suggest that the share of respondents with at least some college decreased after introduction of Duolingo.

FIGURE 2.D.15: The Effect of Duolingo Exposure on Migrant Outcomes upon arrival

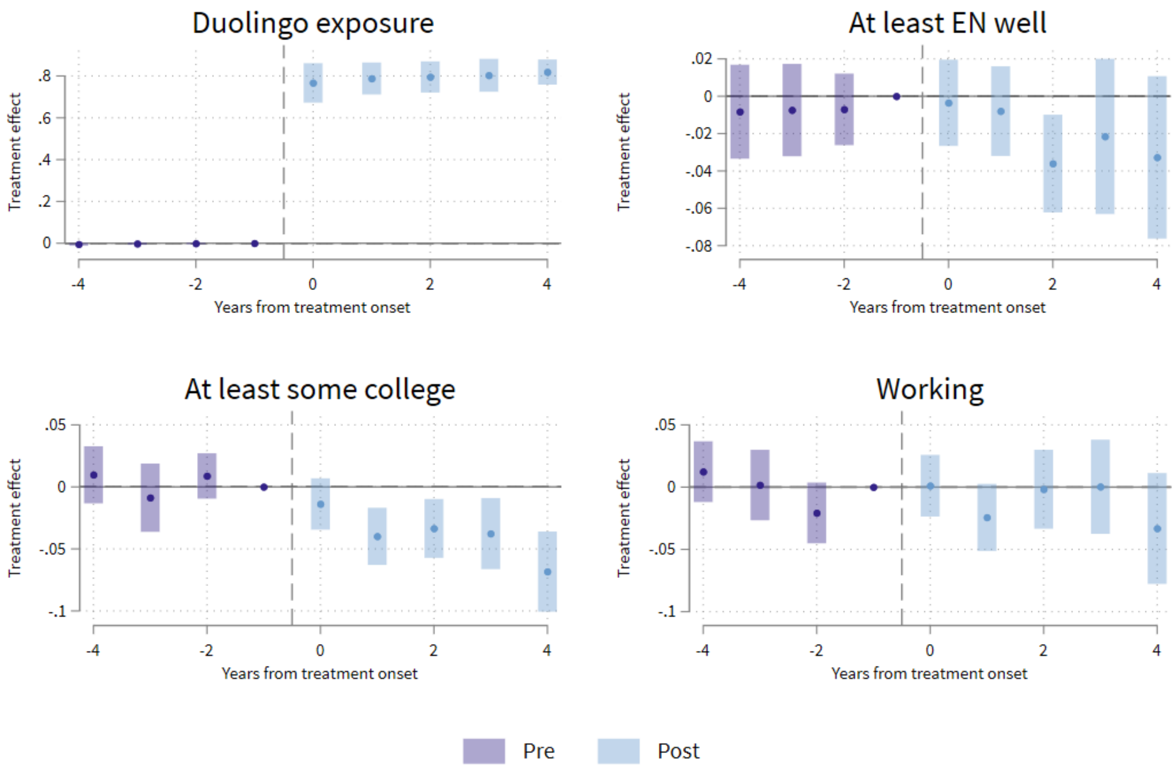

Notes: OLS results from a Nagengast and Yotov (2023) event study estimator around large increases in Duolingo exposure on outcomes among those, including origin and year fixed effects. The panels report results for four different outcomes: the exposure itself (upper left), speaking English at least well (upper right), having at least some college education (lower left) and currently working (lower right) by native language of the test takers. As the unit of observation is the native language level, the Duolingo exposure is binary. Shaded blue areas indicate 95% confidence intervals based on standard errors that are clustered at the origin country level.

Results

Table 2.D.14 examines the language skills, characteristics and employment outcomes upon arrival, with and without country-year controls. The results suggest that the probability to speak English very well in the first year since arrival has decreased, although it is not statistically significant. Moreover, as suggested in Figure 2.D.15, the probability to have college educations has decreases by 2 percentage points on average, but in the two-way fixed effects regression the effect is not significant. Panel C find some suggestive evidence that workers' earnings increase, but that they perform jobs in which English skills are less important. This is suggestive of the fact that the availability of low-cost language learning facilitates immigrants to find jobs in which they do not need to be very proficient in English.

Turning to integration of immigrants beyond the first year after arrival, I analyze the effects of pre- and post-arrival exposure on language skills in Table 2.D.13 and on economic integration in Table 2.D.14. The estimates suggest that the lower share of individuals with very good English language skills upon arrival is only temporary as exposed individuals catch-up. Moreover, the estimates for post-arrival exposure suggest that being able to learn languages after migration increases the probability to speak English at least well with 2 percentage points. Moreover, Table 2.D.14 shows that the introduction of a Duolingo course after arrival increases the probability to be working by 4 percentage points and increases incomes by 7%. The initial lower English-language intensity of immigrants is rapidly catching up to that of unexposed immigrants. In addition, I do not find that exposure to Duolingo after arrival decreases the English intensity of immigrants' jobs. This suggest that the effect of language knowledge does not decrease the English intensity of jobs. This further suggests that selection effects are driving the decreased English intensity of jobs upon arrival.

TABLE 2.D.12: The Effect of Duolingo Exposure on Language Skills upon Arrival in the U.S.

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Language skills						
	At least some		At least well		At least very well	
$DL_{oc}^{T_d}$	0.006	-0.005	-0.006	-0.013	-0.014	-0.024
	(0.011)	(0.011)	(0.013)	(0.014)	(0.009)	(0.015)
Observations	44611	34824	44611	34824	44611	34824
R^2	0.32	0.32	0.35	0.35	0.22	0.22
Mean dep. var.	0.820	0.821	0.619	0.624	0.350	0.352
Panel B: Selection						
	Female		At least 9th grade		At least some college	
$DL_{oc}^{T_d}$	-0.016	-0.022	-0.002	0.005	-0.009	-0.020
	(0.011)	(0.013)	(0.011)	(0.011)	(0.012)	(0.012)
Observations	44611	34824	44611	34824	44611	34824
R^2	0.03	0.03	0.17	0.18	0.24	0.25
Mean dep. var.	0.482	0.476	0.886	0.886	0.451	0.464
Panel C: Integration						
	Working		Log earnings		Occupation: importance English	
$DL_{oc}^{T_d}$	0.029	0.028	0.034	0.109**	-0.942*	-1.010*
	(0.019)	(0.019)	(0.062)	(0.046)	(0.566)	(0.538)
Observations	44611	34824	18650	14673	21157	16202
R^2	0.15	0.17			0.37	0.40
Mean dep. var.	0.418	0.421	31,921	32,512	40.947	40.880
Country-year controls		✓		✓		✓

Notes: PPML (Panel C column 3 and 4) and OLS (all others) estimations of the model of equation 2.10. Panel A, B and C consider those interviewed in the first year after arrival in the full 2007-2022 ACS. Panel A includes three binary indicators for minimum levels of language skills, Panel B includes binary indicator for being female, having at least completed 9th grade and a binary indicator for having at least some college educated. Panel C includes a binary indicator for being in work, yearly labor income in US Dollar and the occupation-level importance of English from ONET. This is a score between 0 and 100 indicating how important a skill is for the job. The skill description for use of English Language: "Knowledge of the structure and content of the English language including the meaning and spelling of words, rules of composition, and grammar".
All even columns include extensive origin-year level controls: log of GDP per capita, unemployment rates, the share of population with tertiary education, the number of conflict deaths from the Global Burden of Disease dataset, the GINI coefficient from the World Inequality Database, the median income from the World Bank, and the share of admissions of legal permanent residents by visa type from Yearbook of Immigration Statistics of the DHS. The measure of Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. Standard errors reported in parentheses are two-way clustered: on the country group of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.D.13: The Effect of Duolingo Exposure on Language Skills after arrival in the USA

	(1) Upon arrival	(2)	(3) Full	(4)	(5) Interview before Duolingo Exposure
Panel A: Speaks at least some English					
$DL_{oc}^{T_d}$	0.006	0.004	0.001		
	(0.011)	(0.009)	(0.010)		
$DL_{oc}^{T_d} \times \log(t-c+1)$		-0.005	-0.004		
		(0.005)	(0.005)		
$DL_{otc}^{T_d,post}$			0.006	0.008	0.002
			(0.006)	(0.007)	(0.008)
Observations	44611	376541	373819	373819	209922
R^2	0.32	0.26	0.26	0.26	0.26
Mean dep. var.	0.820	0.882	0.882	0.882	0.910
Panel B: Speaks English at least well					
$DL_{oc}^{T_d}$	-0.006	-0.004	-0.011		
	(0.013)	(0.010)	(0.009)		
$DL_{oc}^{T_d} \times \log(t-c+1)$		-0.002	0.002		
		(0.008)	(0.008)		
$DL_{otc}^{T_d,post}$			0.019**	0.020**	0.020**
			(0.009)	(0.010)	(0.010)
Observations	44611	376541	373819	373819	209922
R^2	0.35	0.37	0.37	0.37	0.38
Mean dep. var.	0.619	0.679	0.678	0.678	0.715
Panel C: Speaks English at least very well					
$DL_{oc}^{T_d}$	-0.014	-0.020***	-0.023**		
	(0.009)	(0.008)	(0.009)		
$DL_{oc}^{T_d} \times \log(t-c+1)$		0.009*	0.010*		
		(0.004)	(0.005)		
$DL_{otc}^{T_d,post}$			0.006	0.008	0.011
			(0.009)	(0.010)	(0.009)
Observations	44611	376541	373819	373819	209922
R^2	0.22	0.27	0.27	0.27	0.28
Mean dep. var.	0.350	0.416	0.416	0.416	0.437
$DL_{oc}^{T_d} \times (t-c)$ FE				✓	

Notes: OLS estimations of the model of equation 2.13, with the following outcomes: Speaking at least some English (A), speaking English at least well (B) and speaking English at least very well (C). Column 1 shows the effect upon arrival as shown in previous tables. Column 2-4 report results from the full sample of immigrants within the first 5 years of arrival. To study whether the initial gains in outcomes fade out over time, I interact the Duolingo exposure with the log of years since arrival plus one in column 2. Column 3 introduces the post-arrival treatment intensity and Column 4 includes an interaction of pre-arrival intensity with dummies of years since arrival (not shown). The last column limits the sample to those who had no exposure to Duolingo before arrival. The measure of Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. Standard errors reported in parentheses are two-way clustered: on the country group of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 2.D.14: The Effect of Duolingo Exposure on Migrant Outcomes after arrival in the USA

	(1)	(2)	(3)	(4)	(5)
	Upon arrival		Full		Interview before Duolingo Exposure
Panel A: Working					
$DL_{oc}^{T_d}$	0.029	0.051***	0.037***		
	(0.019)	(0.013)	(0.012)		
$DL_{oc}^{T_d} \times \log(t-c+1)$		-0.032***	-0.027***		
		(0.008)	(0.007)		
$DL_{otc}^{T_d,post}$			0.035***	0.032**	0.039***
			(0.011)	(0.012)	(0.010)
Observations	44611	376541	373819	373819	209922
R^2	0.15	0.19	0.19	0.19	0.21
Mean dep. var.	0.418	0.627	0.626	0.626	0.623
Panel B: Log yearly labor income					
$DL_{oc}^{T_d}$	0.034	0.006	-0.029		
	(0.062)	(0.045)	(0.034)		
$DL_{oc}^{T_d} \times \log(t-c+1)$		-0.019	-0.001		
		(0.029)	(0.022)		
$DL_{otc}^{T_d,post}$			0.072	0.084	0.069**
			(0.049)	(0.053)	(0.031)
Observations	18650	236043	234102	234102	130799
R^2					
Mean dep. var.	31920.838	45468.648	45446.903	45446.903	48654.692
Panel C: Occupation-level importance of English language					
$DL_{oc}^{T_d}$	-0.942*	-1.150***	-0.919***		
	(0.566)	(0.424)	(0.308)		
$DL_{oc}^{T_d} \times \log(t-c+1)$		0.669**	0.565**		
		(0.284)	(0.235)		
$DL_{otc}^{T_d,post}$			-0.522	-0.516	-0.114
			(0.435)	(0.456)	(0.406)
Observations	21157	238209	236135	236135	121743
R^2	0.37	0.32	0.33	0.33	0.33
Mean dep. var.	40.947	41.014	41.017	41.017	41.793
$DL_{oc}^{T_d} \times (t-c)$ FE				✓	

Notes: PPML (Panel B) and OLS (Panel A and C) estimations of the model of equation 2.13, with the following outcomes: Being in work, yearly wage income in US Dollar and the occupation-level importance of English from ONET. This is a score between 0 and 100 indicating how important The skill description for use of English Language: "Knowledge of the structure and content of the English language including the meaning and spelling of words, rules of composition, and grammar". Column 1 shows the effect upon arrival as shown in previous tables. Column 2-4 report results from the full sample of immigrants within the first 5 years of arrival. To study whether the initial gains in outcomes fade out over time, I interact the Duolingo exposure with the log of years since arrival plus one in column 2. Column 3 introduces the post-arrival treatment intensity and Column 4 includes an interaction of pre-arrival intensity with dummies of years since arrival (not shown). The last column limits the sample to those who had no exposure to Duolingo before arrival. The measure of Duolingo exposure is constructed using data from (Ginsburgh, Melitz and Toubal, 2017) and the rollout dates of Duolingo courses. Standard errors reported in parentheses are two-way clustered: on the country group of origin and country of destination level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Chapter 3

The Effect of Conflict on Refugees' Return and Integration: Evidence from Ukraine

This chapter is based on joint work with Cevat Giray Aksoy, Yvonne Giesing and Panu Poutvaara

3.1 Introduction

The number of refugees worldwide rose from 27 million in 2021 to 35 million by the end of 2022, the largest annual increase ever recorded.[1] This increase was largely driven by refugees from Ukraine, who, along with Syria and Afghanistan, account for more than half of the world's total refugee population (UNHCR, 2023a). Whether refugees return is important for both their countries of origin and their host countries. The return of refugees to their home countries is crucial for reconstruction efforts, as they often bring valuable skills and resources. For host countries, the potential return of refugees can alleviate the social, economic and political pressures that large influxes can create, such as strained public services, competition for jobs, and challenges to social cohesion. When refugees choose to return, this not only alleviates these challenges, but also promotes a more sustainable approach to migration management.

Although many refugees, particularly those in neighboring countries, initially intend to return when conditions are safe, a substantial number ultimately choose to remain in their host countries (Alrababa'h et al., 2023; UNHCR, 2023b). However, there is a lack of systematic evidence on how refugees' intentions to return change over time, how accurately these intentions predict actual return, and the impact of conflict in refugees' home regions on their return plans, actual return, and integration. This evidence gap arises from the limited availability of longitudinal data that track refugees over time and across countries. The analysis of cross-sectional data is often insufficient to determine the causal effect of conflict on return (intentions), as unobserved heterogeneity among individuals may depend on the intensity of conflict prior to departure.

To address these issues, we launched a longitudinal survey of Ukrainian refugees across Europe in June 2022, following Russia's full-scale invasion of Ukraine on February 24, 2022, which led to the largest refugee crisis in Europe since World War II. We repeatedly ask respondents about their current location, return plans, and integration outcomes, and link this information to geocoded data on conflict intensity in their home municipality. This allows us to estimate the causal effect of local conflict on actual return, return plans, and integration outcomes. In addition, by collecting refugees' expectations about the duration and resolution of the war, we examine how changes in these expectations affect the same set of outcomes.

Our descriptive findings indicate a strong desire among Ukrainians to return home. Initially, around two thirds of Ukrainian refugees intended to return either soon or when it becomes safe, and one in ten planned to settle permanently abroad. Return plans strongly predict actual return: 33 percent of those who initially intended to return soon did so, while none of those who planned to settle permanently outside Ukraine returned. The realized return rate was 2.7 percentage points per 100 days. The net increase in plans to settle outside Ukraine was 1.6 percentage points per 100 days.

Controlling for individual, survey week, and host country fixed effects, we find that the liberation of one's home district increases the probability of returning to one's home municipality by 5.8 percentage points, and a one standard deviation increase in conflict intensity reduces the probability by 1.6 percentage points. We find that the latter has no statistically significant effect on returning to anywhere in Ukraine, suggesting that severe conflict may redirect returns to safer regions within the country. In terms of integration outcomes, the liberation of one's home municipality is associated with a lower likelihood of participating in training, consistent with a shorter time horizon in the destination country, which reduces incentives to invest in host-country specific human capital (Chiswick and Miller, 1994; Cortes, 2004; Adda, Dustmann and Görlach, 2022). However, we find no significant effect of conflict in the municipality of origin on various integration outcomes, such as employment, host country language proficiency, and social or subjective integration.

[1] There are also 5.4 million asylum seekers and 62.5 million internally displaced persons. For a comprehensive global overview, see UNHCR (2023a).

We examine the robustness of our results through several checks. We show that alternative specifications or allowing for spatial correlations in the error structure do not change the results. We also show that the results are largely robust to alternative ways of constructing the treatment, for example when including linear measures of conflict in the principal component analysis. When extending the radius to include deaths further away from one's home municipality, results become insignificant, highlighting that we capture the effect of local conflict rather than at a larger scale. Finally, we check effect heterogeneity and confirm that the results remain similar when, for example, we omit those individuals who come from areas that were already occupied before 2022.

Over time, Ukrainian refugees became less optimistic about their country's victory and the prospect of Russia's withdrawal from all occupied territories by the end of 2024. Our results show that as individuals became more pessimistic, they also became 4.7 percentage points more likely to plan to settle outside Ukraine. However, there is a remarkable resilience in the return plans of the population as a whole. Between September 2022 and January 2023, 71 percent of the panel participants expected Ukraine to regain all of the occupied territories by the end of 2024. By October-November 2023, however, this optimism had fallen to 35 percent. Despite this decline in confidence about territorial liberation, the propensity to return to Ukraine or make plans to do so — either immediately or when conditions are deemed safe — decreased considerably less on the same sample, from 66 percent to 54 percent. Meanwhile, the proportion of respondents who have decided to settle outside Ukraine has increased by 7 percentage points.

Turning to individual-level heterogeneity in return and return plans, 43 percent of those who returned reported that the partner staying behind was the main reason to return, or one of the two main reasons. Results are similar in a regression framework: having a partner back in Ukraine helps to explain 33 percent of total returns in panel regressions. As men aged 18 to 60 are not allowed to leave Ukraine, with certain exceptions, having a partner staying behind concerns almost exclusively female refugees. We also find that men are more likely to plan to settle outside Ukraine. Given the small number of men in our sample, results for men are more uncertain than results for women.

Our results regarding the return intentions of Ukrainian refugees contrast with previous research, which shows that refugees are less inclined to return to their home countries than those who migrate for economic reasons (Cortes, 2004; Camarena and Hägerdal, 2020). This prompts the question: what explains the strong return intentions of Ukrainian refugees?

First, in contrast to civil wars, external threats often catalyze a stronger sense of national identity (Kulyk, 2016; Gehring, 2021; Abramenko, Korovkin and Makarin, 2024). A stronger sense of national identity can greatly increase the emotional cost of living outside one's country. Second, the effective resistance to the Russian invasion has increased Ukrainians' trust in their government and military. This increased trust could not only motivate those who remained in Ukraine to stay, but also induce those who left to return. Moreover, such confidence in government and military institutions could foster optimism about the future of Ukraine. Third, support from the international community, along with the potential for EU accession and NATO membership, is expected to reinforce this optimism. If these factors are strong, they could encourage a large share of Ukrainians to return home.

To explore these potential mechanisms, we use data from the Gallup World Polls and surveys conducted by the Razumkov Center in Ukraine both before and after the full-scale invasion. These two datasets have the advantage of allowing us to examine how trust, confidence in the government and military, and optimism changed after Russia's full-scale attack. An analysis of refugees would not allow us to answer this question, as they were surveyed only after Russia's full-scale attack. Linking trust, confidence, and optimism to the desire to emigrate sheds light on the strength of each mechanism. Our analysis reveals a sharp decline in the share of Ukrainians desiring to live outside Ukraine between 2021 and 2022 (26 percentage points), which cannot be explained by selective outmigration. A sharp increase in confidence in the government and the military and increased optimism help explain 41 percent of the decrease in the desire to emigrate from 2021 to 2022 among Ukrainians living in Ukraine, and a stronger sense of national identity 22 percent. Although these percentages are calculated from two different datasets and cannot be summed directly, they suggest that improved confidence, optimism, and national identity play an important role in the decrease in Ukrainians' desire to emigrate.

Related literature and our contributions

Previous research on refugee crises has established that post-conflict return rates among refugees are low, with returning individuals doing so with significant delays (Camarena and Hägerdal, 2020; Beaman, Onder and Onder, 2022). Factors such as destroyed housing, deteriorating security conditions at home, and personal experiences of violence are identified as major obstacles to return migration (Balcilar and Nugent, 2019; Serdar and Orchard, 2020; Beaman, Onder and Onder, 2022; Alrababa'h et al., 2023). However, this evidence is based on refugees fleeing undemocratic regimes or civil wars, often persecuted by their own governments. Our study extends the current research by examining a case in which refugees are forced to leave a democratic country because of external aggression. This particular focus allows us to uncover new insights into the mechanisms of refugee return and their integration under these circumstances.

In addition, our paper is related to the larger literature that examines the decision to return and the timing of return migration, pioneered by Borjas and Bratsberg (1996). Dustmann (2003) notes that the optimal migration duration of temporary migrants may be a non-monotonic function of the wage differential between home and host country. Adda, Dustmann and Görlach (2022) studies skill acquisition in the host country and how shocks to duration affect skill and wage profiles. Görlach (2023) examines how return and repeat migration depend on financial constraints in the context of Mexican migration to the United States. We add to this literature by showing how conflict in the home country affects the timing of return.[2]

Within the scope of conflict-related studies, prior research has largely examined the effects of past exposure to violence, highlighting its negative consequences on various life outcomes (for example, see Chamarbagwala and Morán (2011), Shemyakina (2011), León (2012), Rodriguez and Sanchez (2012), Verwimp and Van Bavel (2014), Akbulut-Yuksel (2014)). However, Becker et al. (2020) and Aksoy et al. (2024) show that experiences of forced migration can lead to increased investment in education. We contribute to this literature using rich panel data that allows us to link return migration, changes in return intentions, and refugees' investments in the host country with detailed local conflict data and expectations about the outcome of the war.

Finally, our research also relates to the broader body of work examining the determinants of refugees' labor market integration.[3] Previous studies have shown that immigrant networks (Edin, Fredriksson and Åslund, 2003; Damm, 2009; Beaman, 2012), language training (Arendt et al., 2022), job search assistance (Battisti, Giesing and Laurentsyeva, 2019), and positive attitudes of natives (Aksoy, Poutvaara and Schikora, 2023) have a positive influence on refugee integration success. Our study adds to this literature by examining how ongoing conflicts in refugees' home municipalities impact their integration efforts in host countries.[4]

Our context differs from other refugee scenarios where return is either forced through deportation or made impossible by persecution. Moreover, Ukrainian refugees have autonomy in

[2]Previous literature has highlighted the importance of return migration for the development and reconstruction of the country of origin in terms of innovation (Choudhury, 2016) and entrepreneurship (Massey and Parrado, 1998; Demurger and Xu, 2011; Krasniqi and Williams, 2019). Return migration and contacts with the diaspora can also foster trade (Parsons and Vézina, 2018; Bahar et al., 2022), investment (Mayda et al., 2022), and political change (Chauvet and Mercier, 2014; Barsbai et al., 2017). Hence, we provide evidence on the return intentions of Ukrainian refugees, which is crucial for policymakers both in Ukraine and in refugees' destination countries.

[3]For a comprehensive review, see Strang and Ager (2010) and Becker and Ferrara (2019).

[4]Two recent studies are relevant to our paper. Zaiour (2023) documents that drug-related violence in Mexico increases naturalization rates of Mexican immigrants in the US, but does not affect labor market outcomes and human capital investment decisions. Bassetto and Freitas Monteiro (2024) finds that terrorism in the home country reduces return intentions and increases employment.

making decisions about their integration, unlike other refugee groups who often encounter temporary work restrictions or are required to participate in mandatory integration and language courses. Consequently, there is more to learn from the decisions made by Ukrainian refugees compared to those made by refugees who face legal restrictions or are prevented from making choices.

The paper is organized as follows: Section 3.2 describes the data, followed by Section 3.3, which provides descriptive statistics. Section 3.4 presents the main empirical specification. Section 3.5 discusses the main results of conflict on return, return intentions, and a range of integration outcomes. Section 3.6 explores how expectations about the conflict at large shape return intentions. Section 3.7 aims to better understand why so many Ukrainians want to return, compared to other refugees, through the lens of migration aspirations among Ukrainians in Ukraine before and after the Russian invasion. Section 3.8 concludes the paper.

3.2 Data

3.2.1 Survey of Ukrainian Refugees in Europe

We collaborated with the survey company Verian (formerly Kantar Public) to conduct a six-wave online panel survey of Ukrainian refugees across Europe. For the first wave (hereafter: baseline) survey, respondents aged 18 and over were recruited via Facebook ads, and for subsequent waves, contact was made via email. The baseline was conducted between 14 June 2022, and 22 December 2022. On average, respondents completed the survey 194 days after leaving Ukraine. The survey was completed by 11,783 respondents with Ukrainian citizenship, of whom 6,299 agreed to participate in future waves.[5] Figure 3.A.1a shows the distribution of Ukrainian refugees across European countries and Figure 3.A.1b shows the sampling rate across European destinations, dividing the number of baseline respondents by the number of Ukrainians registered for temporary protection in December 2022.[6] All major host countries have a sampling rate of at least 1 in 1000 refugees. Those who agreed to be recontacted were asked by email to complete five follow-up surveys between September 2022 and November 2023. The follow-up emails explicitly asked respondents who returned to Ukraine to complete the survey. Participants received a 3 Euro voucher to encourage participation and minimize attrition rates in each survey wave. Table 3.A.1 details the specific times and number of observations for each wave and Figure 3.A.2 graphically shows the distribution of interviews over time.[7] A total of 18,202 interviews were completed, with 2,674 individuals participating in at least two interviews that are at least 30 days apart.

The first survey wave includes questions on migrants' demographic characteristics, past and present employment status, their current living situation, and intentions to return. Specifically, we explore the intention to return through the following question: *What are your plans regarding returning to Ukraine?* Response options include: (i) *I intend to go back very soon*; (ii) *I intend to go back at some point later when I feel it is safe to return*; (iii) *I do not intend to go back and plan to settle outside Ukraine*; (iv) *Do not know yet*; (v) *Prefer not to answer*.

Furthermore, we ask respondents to indicate where in Ukraine they lived before leaving Ukraine. Specifically, we ask for the region in a drop-down menu and municipality (*hromada*

[5]In all analyses, we exclude the small proportion of respondents (101 in the baseline survey) who do not hold Ukrainian citizenship.

[6]Due to the timing of our recruitment period, our sample is not representative of Ukrainians leaving in 2023. Ukrainians who applied for temporary protection in 2022 accounted for 80 percent of all registrations by the end of 2023, according to Eurostat (2023).

[7]Table 3.A.2 shows the participation frequency of each respondent across the survey waves.

in Ukrainian) in a write-in field.[8] To match respondents to local conflict measures (see below), we parse the fill-in field for municipality of origin and match 82 percent of Wave 1 respondents to a unique municipality of origin. Figure 3.A.3 shows where our participants are from. The largest sampling rates (as a share of the 2021 population) can be found in high-conflict regions in the east and in the south, as well as in Lviv and Kyiv.[9]

In the five follow-up waves, respondents were also asked about their current location, their expectations about the war, and a range of integration outcomes. Importantly, we ask respondents about their main activity (e.g., working, studying, or unemployed), the number of Ukrainian and local friends in the destination country, whether they are taking a language course, two questions on host country language skills (speaking and reading), and subjective integration. We combine the questions on language skills and subjective integration in principal component analyses (PCAs). For more information on the Verian survey and its detailed questions, see Appendix 3.A.1. We construct two primary samples. First, we consider the full baseline survey sample. Second, we construct a sample of long differences between each individual's last response and their response in the baseline survey (hereafter: long differences sample). The average number of days between the interviews in this sample is 268, with a minimum of 30 (by construction) and a maximum of 506.

3.2.2 Conflict data

ACLED and UCDP

To obtain measures of local conflict intensity, we use the Armed Conflict Location Event Data Project (ACLED) (Raleigh et al., 2010) and the Uppsala Conflict Data Program's Georeferenced Event Dataset (UCDP-GED) 23.1 (Sundberg and Melander, 2013) databases.

ACLED and UCDP automatically collect news reports of conflict data that are human-coded using standardized methods and, if possible, geocode the event. ACLED includes the primary actor, the type of conflict and the number of reported fatalities, among others. UCDP-GED is also an event-level dataset, but with the strict inclusion criterion that at least one death should have been recorded.[10] In many cases, death tolls are estimates and may vary between UCDP and ACLED for the same event. Death tolls may not be known, or may be measured with error. Furthermore, the events may also be included in less severe instances (especially for ACLED), but on average may provide a reasonable summary measure of conflict. Although this introduces some measurement error, Ukrainian refugees may be no better informed than what ACLED and UCDP can infer from news reports. Because of these concerns, we use both the number of events as well as the number of deaths from both ACLED and UCDP in the following analysis. Between February 24, 2022 and November 7, 2023 (the last day of the sixth wave of our survey) ACLED recorded 85,298 events and 61,446 fatalities, while UCDP recorded 11,099 events and 157,015 fatalities. The latter includes statistical corrections for the number of Ukraine-wide casualties, which we disregard in our analysis. In the following analysis, we only use events that are exactly geocoded or geocoded at the municipality level. This drops less than 1 percent of the events in ACLED and 15 percent of the events in UCDP.

To calculate a measure of conflict intensity in a municipality between two dates, we calculate the number of events and deaths per 30 days and take the $log(x+1)$ transformation. As this gives us four different measures (events and deaths in ACLED and UCDP) measuring the intensity of

[8] The regions include 24 *oblasts*, the Autonomous Republic of Crimea, the city of Kyiv and the city of Sebastopol. As of 2022, Ukraine has 27 regions, 137 districts and 1469 municipalities.

[9] As there is no representative data on the exact origin of refugees within Ukraine, we cannot exactly assess how representative our sample is in terms of origins.

[10] Raleigh and Kishi (2019) compare different conflict datasets and show that in the case of the 2018 conflict in Donbass, Ukraine, ACLED and UCDP give more plausible results than automated conflict datasets. Therefore we do not use these datasets. They also find that ACLED captures more events that only appeared in non-English speaking media than UCDP, which is an advantage of ACLED in the current context.

local conflict, which are strongly correlated,[11] we combine them through a principal component analysis and obtain the standardized first principal component $Conflict_{mt_1t_2}$.[12] We illustrate the distribution of our measure of conflict at the municipal level in Figure 3.A.4. This Figure shows the conflict intensity between the first and last interview in our long differences sample. Conflict during this period is concentrated in a band along the front line, as well as along the border with Russia, and in Kyiv, Dnipro, and other bigger cities where Russian missile and drone strikes have brought devastation.

Institute for the Study of War

We construct a daily dataset of the location of the frontline using the maps created by the Institute for the Study of War (ISW). ISW's maps visualize the state of the war based on publicly available information sourced from news outlets, social media and satellite imagery. Importantly, these maps include a line that approximates the front line of the conflict. We categorize a district (*Raion* in Ukrainian) as either *under Ukrainian control*, *on the frontline*, or *occupied*. For subsequent analysis we calculate the change in the frontline status between the two interview dates for each respondent. As we are particularly interested in the effects of the liberation of one's district of origin, we calculate whether a district has been continuously under full Ukrainian control, whether one's district has been fully liberated, or whether it has been continuously on the frontline or occupied by Russia between two survey dates. We do this at the district level, as the proximity of conflict matters to the perceived threat of conflict proximity. Most of the variation occurred up to September 2022 in Kharkiv region and up to November 2022 in Kherson region. Figure 3.A.5 shows the changes in occupation status between the first and last interview in our long differences sample. On the long differences sample, 8 percent of individuals originate from districts that were liberated, whereas 18 percent of individuals originate from districts that were continuously occupied or on the frontline.

3.2.3 Other Data

To further explore the mechanisms underlying our findings, we use data from the Gallup World Poll and IKDIF/Razumkov. These datasets consist of survey responses collected in Ukraine both before and after the outbreak of the war. They provide insights into how the intentions and beliefs of Ukrainians who remained in the country have evolved. A more detailed description of these data can be found in the Appendices 3.A.2 and 3.A.4.

3.2.4 Selection and attrition

Our online survey leverages Facebook Ads for recruitment, providing an advantage with its wide reach, precise targeting, and cost-effectiveness. This platform allows us to quickly access diverse and hard-to-reach populations affected by conflict, facilitating timely and efficient data collection. The anonymity provided by online engagement enhances participant safety and encourages candid responses. Despite challenges such as digital access and sample representativeness, the use of Facebook Ads for surveys allows for rapid survey rollout and data collection, establishing it as a powerful method for conducting research in complex areas such as refugee migration and conflict. More than 15 million Ukrainians used Facebook on a monthly basis in early 2022 (Datareportal, 2022), reaching more than 41 percent of the population over the age of 13.

To examine the representativeness of our sample, we compare the observable characteristics of our sample with administrative data from Eurostat on Ukrainians who received a Temporary Protection Status (TPS) (Eurostat, 2023). Table 3.A.3 shows how the baseline survey sample and

[11] Pairwise correlations range from 0.69 to 0.77.

[12] On the long differences sample, the eigenvector of the first principal component are 0.47 for ACLED events, 0.50 for ACLED deaths, 0.55 for UCDP events, and 0.47 for UCDP deaths.

the long differences sample differ from Ukrainians who applied for Temporary Protection Status. Overall, our sample matches the characteristics of temporary protection beneficiaries reasonably well. Women and middle-aged respondents are more likely to respond to our survey, while refugees in Czechia are less likely to do so.

Second, response in the follow-up waves could be nonrandom, biasing our aggregate statistics and estimated treatment effects. To understand what determines response, we estimate logit regressions of follow-up response on initial return intentions and measures of conflict. Table 3.B.1 presents the main results of this analysis. Columns 1 to 3 show that weaker return intentions predict responding to more waves, ever responding to a follow-up wave, and being part of the long difference sample. Those who plan to settle outside Ukraine are 28 percent more likely and those who plan to return very soon are 28 percent less likely to be in the long difference sample than those who plan to return when safe. Since more concrete return intentions predict a lower probability of a follow-up response, these results suggest that we are underestimating return rates due to selective attrition. In addition, column 4 shows that our three main measures of conflict intensity do not strongly predict follow-up response. A one standard deviation increase predicts only a 12% higher probability of being part of the long differences sample. We discuss how selective attrition might affect our results in Section 3.4.

To better represent the Ukrainian refugee population in descriptive statistics and regressions, we use two types of weights. To make the long differences sample more representative of our first wave population, we weight with inverse probability weights based on the predicted probabilities from column 3 of Table 3.B.1. To make the sample more representative of the whole Ukrainian refugee population in Europe, we construct population weights based on the probability of observing a respondent in each sex-age-host country bin using the EU temporary protection status data.[13]

3.3 Descriptive statistics

We weight the baseline wave of the Verian survey with population weights and compare the sample to the nationally representative 2019-2021 Gallup World Poll surveys in Ukraine in Table 3.A.4. Since men aged 18-60 are generally not allowed to leave the country, we show characteristics for men and women separately in Table 3.A.4. There are some notable differences between migrants and the general population. Individuals with a partner are more likely to have migrated, as are women (but not men) with children. In addition, those with tertiary education and those living in urban areas are more likely to have migrated.

Table 3.A.5 shows descriptive statistics of the values of all covariates used in the long differences estimation sample, as well as the baseline sample without missing covariates. 19 percent of baseline survey respondents originate from districts on the frontline and 6 percent originate from districts behind the frontline. 8 percent of respondents in the long differences sample originate from territories that were liberated after June 2022, 18 percent from territories still under occupation, and 3 percent originated from territories already occupied before the large-scale Russian invasion. The average number of days in the destination country is 268, 6 percent left before the February 24, 2022. The average population of one's municipality is very large, because of the strong propensity to leave from large cities, such as Kyiv.

[13] For several of these bins, we have zero respondents in our survey and no data to weight. These are males 18-34 in Iceland, Luxembourg and Malta, 35-64 in Denmark and Iceland, and 65+ in Denmark, Estonia, Luxembourg, Norway and Cyprus. For a further 42 respondents, we do not have information on the exact country of destination and discard them. For Hungary, Moldova and the United Kingdom, we do not have detailed information on the number of refugees by detailed bin, so we weight these observations only for the whole country.

3.3.1 Return intentions

Most Ukrainian refugees intend to return soon or when it is safe. Table 3.A.5 shows that during the baseline wave 7 percent of respondents planned to return soon, 58 percent when it is safe, 24 percent do not know, and only 8 percent planned to settle outside Ukraine. This is in stark contrast to the return intentions of other refugee groups. Appendix Figure 3.E.1 shows that the intention to stay in Germany of refugees from different countries is above 90 percent for all refugee groups in the first, second to fifth and sixth to tenth year after arrival. Weighting the Verian survey responses with population weights changes these figures only minimally (rightmost column of Panel A of Table **??**).

Figure 3.C.1 and 3.C.2 illustrate the correlates of first-wave return intentions among Ukrainian refugees. The most important predictors of intentions to return soon are having a partner left behind in Ukraine, and, not surprisingly, coming from an area not on or behind the frontline. Moreover, respondents living in Eastern and Southern European countries are more likely to plan to return soon than those living in Germany and the rest of Western Europe (including Northern Europe). Plans to settle outside Ukraine are highest among those from places behind the frontline, men, those without a partner in Ukraine, and those who speak English. Those living in Eastern European countries, except Poland, are less likely to plan to settle outside Ukraine than those living in other countries. Although those from districts behind the frontline are more likely to plan to settle outside Ukraine, conflict intensity does not affect plans to settle outside Ukraine. This may be due to the counteracting forces of the causal effect of conflict and selection. Since those from places with higher conflict intensity are more likely to be forcibly displaced, individuals from such places may be less likely to plan to settle outside Ukraine. Our identification strategy allows us to account for these selection effects.

	(1) Unweighted	(2) Inverse probability weighted (IPW)	(3) IPW and population weighted
Panel A: Baseline return intentions			
Baseline return intention			
Return soon	5.4	7.1	6.5
Return when safe	56.8	56.9	56.0
Do not know	25.6	24.6	25.0
Settle outside Ukraine	9.9	8.2	8.9
Prefer not to answer	2.3	3.1	3.5
Panel B: Share returned to Ukraine			
Baseline return intention			
Return soon	33.1	30.9	27.3
Return when safe	9.2	8.5	7.6
Do not know	2.7	2.6	2.2
Settle outside Ukraine	0	0	0
Prefer not to answer	3.8	2.2	1.6
Panel C: Outcomes			
Outcome			
Returned to Ukraine	7.8	7.8	6.7
Returned to home municipality	6.0	5.8	4.7
Moved to third country	4.4	4.4	4.8
Started planning to settle outside Ukraine	7.5	6.7	7.5
Observations	2,301	2,301	2,296

Two natural questions arise. First, are intentions to return predictive of actual return? In the long difference sample, 7.8 percent of individuals were living in Ukraine at the follow-up wave.[14]

[14] Although we ask where people are living, one might be concerned that many of them will leave Ukraine again soon. Of the 97 individuals who responded to at least three waves and returned before the last wave they responded to, 11 left Ukraine again over an average period of 168 days. Although this is a small sample, the rate of individuals leaving Ukraine again is only 7 percent per 100 days. In addition, in Wave 3 (conducted in January 2023) and Wave 4 (conducted in April 2023), we asked whether people had temporarily returned to Ukraine since their arrival. In January 2023, 14 percent had done so, and in April 2023, 30 percent had done so. This shows that the cost of returning

Panel B of Table **??** shows that respondents with stronger intentions to return are more likely to have returned and that the levels of intentions to return are clearly ordered in terms of propensity to return. Of those planning to return soon, 33 percent have returned and none of those planning to settle outside Ukraine have actually returned. Weighting reduces the return rates only slightly, due to the low sampling rate of high-income destinations such as the Netherlands, Scandinavia, Switzerland, and the United Kingdom. Second, are these strong return intentions persistent? To answer this question, we turn to within-person changes in the first-differences sample. Figure 3.C.3 shows that most changes occur between adjacent levels of return intentions. Because of the large proportion of people who report plans to return when safe, most returnees come from this level. To quantify how quickly return intentions change over time, we nonparametrically plot the evolution of return intentions on the full sample after netting out individual fixed effects in Figure 3.3.1. We find that most levels follow nearly linear trends over time. The share of individuals planning to settle outside Ukraine increased over time (1.6 percentage points per 100 days), as did the share of individuals returning to Ukraine (2.7 percentage points per 100 days). However, the number of individuals who said they would return when it was safe to do so decreased sharply over time (4.7 percentage points per 100 days).

FIGURE 3.3.1: Within-individual return intentions and return over time since arrival

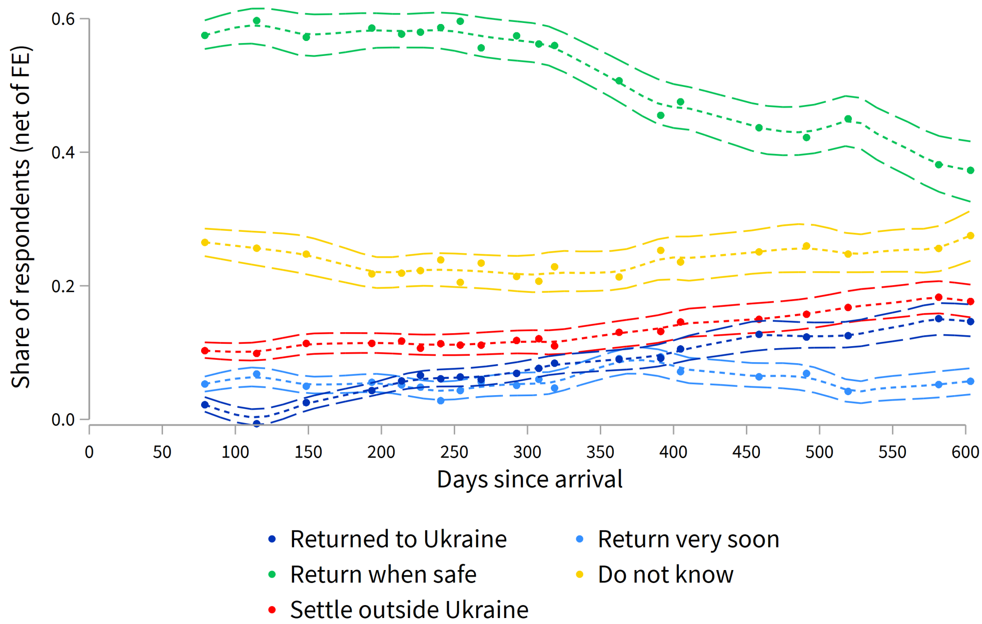

Notes: Binned scatterplot with non-parametric trend for levels of return intentions over time, net of individual fixed effects, with 90 percent confidence interval. For each level of return intentions, we perform the following procedure. First, we assign all observations to 20 equally sized bins over the number of days since arrival in the destination country of residence in the baseline survey. We residualize the outcome by regressing it on individual fixed effects and the number of days since arrival in the first destination country. We perform this procedure for 100 bootstrap samples to obtain smoothed 90 percent confidence intervals. We draw markers for (i) the mean for each of the 20 equally sized bins, (ii) a predicted mean for each bin of the number of days since arrival, (iii) a 90 percent confidence interval around the predicted mean. N = 8,752.

In the following analysis, we are interested in three main aspects of mobility and changes in return intentions. The first measure is whether a respondent returned to Ukraine. Second,

is low. Interestingly, it is not strongly correlated with distance. Among those in Poland, only 33 percent have returned temporarily.

we are interested in whether people returned to their home municipality.[15] Third, we are interested in people who report in the latest wave that they plan to settle outside Ukraine and did not do so in the baseline survey. Less than 5 percent of respondents moved to a country other than Ukraine. Panel C of Table **??** shows the rates of these four outcomes on the long differences sample, unweighted and weighted to reflect the baseline population and the full population under temporary protection. Although those with stronger intentions to return are more likely to respond to follow-up waves, we do not underestimate return rates. The main reason for this is that in the regressions presented in Table 3.B.1, covariates associated with return are also predictive of responding to follow-up waves. After weighting with inverse probability and population weights, 6.7 percent returned to Ukraine, 4.7 percent returned to their home municipality, 4.8 percent moved to another country, and 7.5 percent started planning to settle outside Ukraine. At the same time, 2.9 percent no longer planned to settle outside Ukraine (not shown in the table).

3.3.2 Integration

Refugees often have lower initial levels of integration than other migrant groups (Fasani, Frattini and Minale, 2022). However, in the baseline survey 36 percent of Ukrainian refugees had already started working. Ukrainian refugees under temporary protection could start working directly, which benefits subsequent integration (Fasani, Frattini and Minale, 2021). Moreover, 25 percent of baseline respondents speak at least some of their host country language already. An additional reason for strong initial employment could be the welcoming attitudes of Europeans to Ukrainian refugees as previous research has shown that refugees are more likely to find employment when local attitudes towards immigrants are more positive (Aksoy, Poutvaara and Schikora, 2023). In a similar vein to Figure 3.3.1, we show the non-parametric levels of integration over time since arrival in Figure 3.3.2. The rate of Ukrainians subsequently entering the labor market is 4.0 percentage points per 100 days. However, this hides strong country-level heterogeneity. In Appendix Figure 3.C.4 we show destination country groups. Employment rates in Germany, which pursues a language-first policy, are considerably lower than in all other country groups. Over time, the increase in employment is strongest in the rest of Western Europe (+3.2pp/100 days) and Southern Europe (+2.9pp/100 days), whereas the slope is smaller and not significant at a 5 percent level for Germany (+1.3pp/100 days) and Eastern European countries (+0.9pp/100 days).[16]

Regarding other measures of integration, about 20 percent of respondents were mainly in training in 200 days after arrival, which increased by 2.0 percentage points every 100 days. Being in training as one's main activity is an important measure of demand for host-country specific skills. An important aspect of host-country specific human capital are language skills. More than 60 percent of Ukrainians had enrolled in a language course 200 days after arrival, which increased gradually thereafter (+2.5 pp/100 days).

[15] If one returned, in two follow-up waves (3 and 6) we asked where in Ukraine one returned. In this sample, 80 percent of those who returned did so within the same municipality.

[16] Nevertheless, many Ukrainians work in jobs below their skill level. In wave six, we ask respondents whether their current job matches their qualifications. 38 percent of respondents aged 25 – 59 indicate that they have found a job at their qualification level (see Table 3.C.1). In Germany this is 46 percent.

FIGURE 3.3.2: Integration over time since arrival in the destination, net of controls

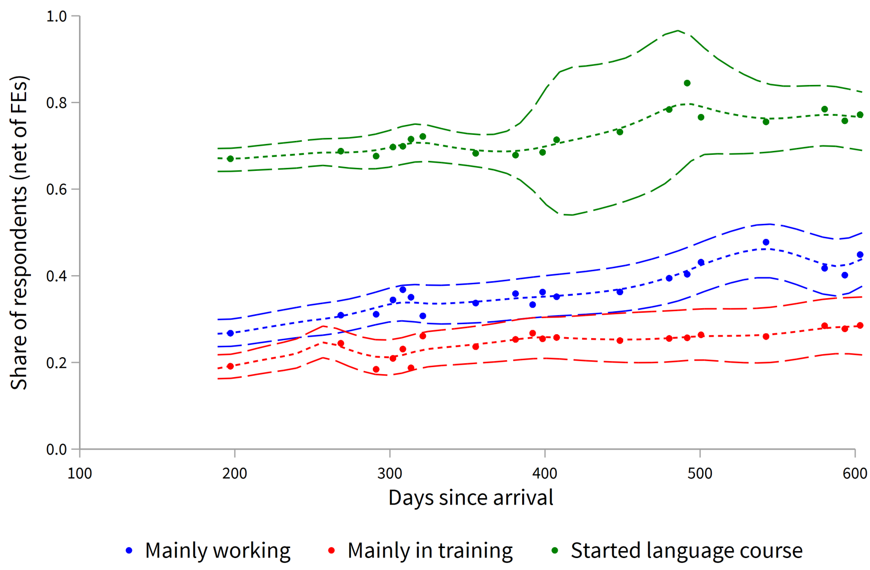

Notes: Binned scatterplot with non-parametric trend of levels of integration outcomes over time, net of individual fixed effects, with 90 percent confidence intervals. We restrict the sample to all respondents aged 25 – 59. N = 3,837 for mainly working and mainly in training and N = 1,875 for started a language course. See Figure 3.3.1 for details about the construction of the non-parametric plots.

3.4 Empirical strategy

To examine the causal impact of local conflict on individuals' return (intentions) and integration outcomes, we regress changes in outcomes between two interview dates on changes in conflict intensity measured in the period between those interviews. Our method offers a significant advantage over cross-sectional analysis as it eliminates the effects of unobserved individual heterogeneity. In particular, by focusing on changes over time within the same individuals, our approach effectively isolates the direct effects of conflict from other confounding factors.

In our primary analysis, we focus on the long difference sample introduced in Section 3.2. We index individuals with i, their municipality (district) of origin by $m(d)$, the start of the full-scale war (February 24, 2022) with t_0, the time they left Ukraine with t_l, and the first and second wave interview dates by t_1 and t_2.

We analyze individual-level *changes* $Y_{im(d)t_1t_2}$ in outcomes between the baseline interview and the last available interview. In the analysis of migration, the outcomes are (i) whether one returned to Ukraine, (ii) whether one returned to one's home municipality, (iii) whether one moved to a country other than Ukraine, and (iv) whether one started planning to settle outside Ukraine. The results for moving to a country other than Ukraine are presented in the online appendix. When analyzing economic integration, we consider whether one is mainly in work and whether one is mainly in some form of training. When analyzing other dimensions of integration, we consider changes in having started a language course and in subjective integration. We estimate the following regression equation:

$$Y_{im(d)t_1t_2} = \alpha t + \beta_1 RegainControl_{dt_1t_2} + \beta_2 RemainOccupied_{dt_1t_2} + \beta_4 Conflict_{mt_1t_2} + \\ \beta_5 Conflict_{mt_0t_1} + \gamma' \mathbf{ReturnInt}_{imt_1} + \delta' \mathbf{X}_{it_1} + \theta_h + \phi_{t_1} + \psi_{t_l} + \epsilon_i \quad (3.1)$$

t is the time elapsed between the two interviews. $RegainControl_{dt_1t_2}$ is a binary indicator for whether or not a district has been fully liberated by Ukraine between the two survey dates, $RemainOccupied_{dt_1t_2}$ is a binary indicator if a district has been continuously occupied between the two interview dates. The reference category for these two mutually exclusive variables is Ukraine fully controlling the district during both interviews. $Conflict_{mt_1t_2}$ is the local conflict intensity per 30 days in one's municipality of origin between the two survey waves. We also control for conflict intensity *before* the first wave interview, $Conflict_{mt_0t_1}$. Controlling for this is important as conflict before the first wave and between survey waves is positively correlated ($\rho = 0.54$). Without this control, the variable $Conflict_{mt_1t_2}$ would likely also capture selection according to initial conflict. To further relax the identifying assumptions, we include the levels of first wave return intentions, $\mathbf{ReturnInt}_{it_1}$. This controls for two factors. First, it accounts for initial factors influencing return intentions by absorbing any factor that determines the first wave return intentions. Second, it accounts for the situation where outcomes, such as the decision to settle outside Ukraine, are inherently zero for those who made such plans in the baseline survey. In robustness checks, we exclude these controls and find that our results remain similar.

We also include host country fixed effects, (θ_h), which account for any unobserved, time-invariant characteristics specific to each host country that might influence the results. In addition, we include fixed effects for the week of departure from Ukraine, (ψ_{t_l}), which control for temporal factors related to the specific week in which individuals left Ukraine (such as the intensity of the conflict at the time of departure or the availability of resources for refugees).

We incorporate the following individual-level baseline covariates $\mathbf{X}_{it_1}$ to account for differences in changes in return intentions driven by demographic factors. First of all, as most refugees are women and men could only leave under special circumstances or by paying a bribe (and thus may face stigma upon return), we are interested in whether there are gender differences in changes in return intentions. Secondly, as many refugees are women whose partners are still in Ukraine, we study whether the presence of a partner affects changes in return intentions. In a similar vein, we study the effect of children remaining in Ukraine. Third, many women are accompanied by children (88 percent of women with underage children answered in the baseline that children were the primary motive to leave Ukraine, see Table 3.C.2) and we study whether the presence of children in the household discourages return. Finally, we are interested in whether return migration may be selective in terms of education. We study the effect of both (i) a formal tertiary education (sample mean: 70 percent) and (ii) whether one indicated to speak English in the baseline survey (sample mean: 46 percent).[17] Given that the conflict-related variables are determined at either the municipality (*hromada* in Ukrainian) or district (*raion* in Ukrainian) level, we address spatial correlation by clustering standard errors at the district level.

Our long difference sample of Ukrainian refugees may exhibit selection bias for two reasons. First, the sample of individuals who respond to a Facebook ad for a survey may be different from the general population of Ukrainian refugees across Europe. To alleviate concerns that our results may be driven by this particular sample, we re-weight our regressions in a series of robustness checks using the weights introduced in Section 3.2.4. Second, those who choose to respond to a request for a follow-up survey may differ from baseline respondents too. First, selective attrition on the outcome could attenuate the point estimates. If returnees are less likely to answer follow-up waves, we underestimate return rates. This is somewhat alleviated by inverse probability weighting on covariates and return intentions predicting return. Second, unobserved factors could simultaneously drive attrition and return intentions. If these factors are correlated with our conflict measures, they could bias our treatment effect estimates. An example of such a concern

[17] Additional factors included in X_{it_1} consist of 7 age categories (18-24; 25-34; 35-44; 45-54; 55-59; 60-64; 65 and older), a binary indicator for partnership status, a binary indicator for originating from an urban area, whether one's home municipality was occupied before February 24, 2022, whether one left Ukraine before February 24, 2022, whether one completed the baseline survey in Russian, the population of one's home municipality, and the squared-term of the population. Controlling for English skills helps to account for the possibility that a considerable share of Ukrainians who have formally tertiary education may not be able to apply their education abroad, due to missing language skills.

FIGURE 3.5.1: The effect of conflict and predictors of changes in return (intentions)

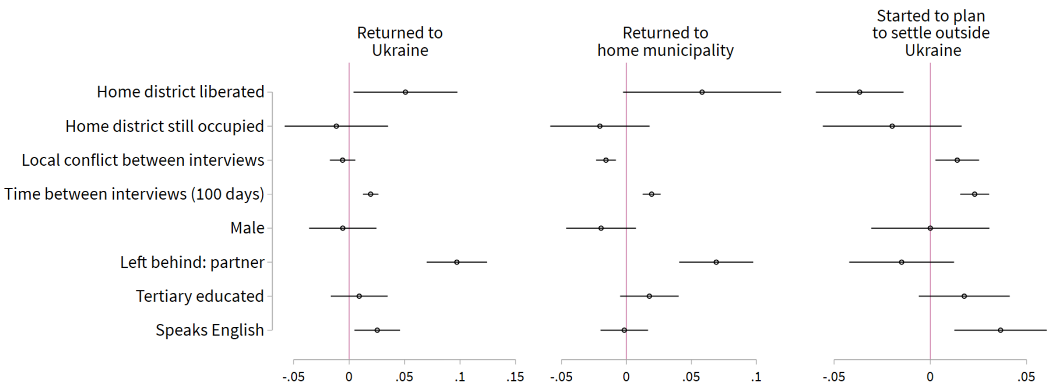

Notes: This figure shows coefficient plots of three multivariate OLS regressions as introduced in Equation 3.1. The outcomes (from left to right) are returned to Ukraine, returned to home municipality, and started to plan to settle outside Ukraine on conflict-related variables and personal characteristics. 95 percent confidence intervals are based on standard errors clustered on the district level. Each regression includes a wide set of control variables and fixed effects as outlined in equation 1. "Home district liberated" and "Home district still occupied" are binary indicators for full liberation of one's home district and whether one's district is at least partially occupied during both survey waves as discussed in Online Appendix 3.A.5. The reference category are districts that have been continuously under Ukrainian control. "Local conflict between interviews" is the standardized first PCA of conflict intensity as discussed in Section 3.2.2. Baseline controls are initial levels of return intentions, age bins (18-24; 25-34; 35-44; 45-54; 55-59; 60-64; 65 and older), the number of days elapsed between the two waves, the population of one's home municipality, population squared and binary indicators for sex, partnership status, tertiary education, speaking English, originating from an urban area in Ukraine, being accompanied by children, having a partner left in Ukraine, having children left in Ukraine, continuing one's Ukrainian job remotely, having left Ukraine before February 24, 2022, originating from a territory that was occupied by Russia or allied forces before February 24, 2022, and answering the survey in Russian. For simplicity of exposition, not all control variables are shown in the figure. N = 2,301 (column 1 and 3); N = 1,433 (column 2).

could be that Ukrainian localities with higher conflict intensity have worse infrastructure (due to Russian attacks). If refugees return to these regions, they may be less likely to respond in the follow-up waves. This would lead to a downward bias in the coefficient of conflict in a regression of return on conflict intensity. However, if such selective attrition occurs, one would expect the conflict to affect response rates for all respondents. This is not what we find. Column 4 of Table 3.B.1 shows that conflict intensity and occupation status do not strongly affect the propensity to be in the long differences sample.

3.5 The causal effect of local conflict on return, return intentions and integration

3.5.1 Return and return intentions

In this section, we focus on the impact of variation in conflict intensity at the local level. We analyze the effect of whether the refugee's home district is liberated, remains occupied, or is on the frontline, and the effect of the conflict in the refugee's home municipality. We examine three outcomes: whether respondents returned to their home municipalities, whether they returned to Ukraine, and whether they started to plan to settle outside Ukraine.

Figure 3.5.1 presents the results from equation (1). We find that the liberation of one's home district significantly increases the likelihood of individuals returning to Ukraine, while simultaneously reducing the propensity to make new plans to settle outside Ukraine. The similarity in the point estimates for returning to Ukraine in general and returning specifically to one's home municipality suggests that most of the increase in returns to Ukraine is due to people returning to their home municipalities after the liberation of their district. Conversely, continued occupation does not have a statistically significant impact on any of the outcomes.

Turning to the effect of conflict intensity, we find that more intense conflict in one's home municipality reduces return to one's home municipality, but not to Ukraine in general. A one standard deviation higher conflict intensity reduces return to one's home municipality by 1.8 percentage points, but return to Ukraine altogether by only 0.8 percentage points ($p \geq 0.10$).[18] Furthermore, more intense conflict in the home municipality makes it more likely that refugees start planning to settle outside Ukraine.

We also examined additional predictors of return in Figure 3.5.1. Having a partner in Ukraine increases the likelihood of returning by 9.7 percentage points. Contrary to expectations, tertiary-educated immigrants are not less likely to return. Surprisingly, proficiency in English increases the likelihood of returning. At the same time, English speakers are also more likely to consider settling outside of Ukraine for the first time. These findings suggest that, if anything, return migrants are not negatively selected from the available sample of migrants.

To alleviate concerns about non-random attrition and selection as discussed in Section 3.4, we weight regressions with inverse probability weights and population weights in Figure 3.D.1. We find that point estimates are quantitatively similar to those in the main results. However, after population weighting standard errors are considerably larger, which is driven by the large variation in the weights.

What factors are most influential in determining return and changing return intentions? We can assess this using our estimates and the individual-level variation in the regressors depicted in Figure 3.5.1. Having a partner remaining in Ukraine contributes 2.6 percentage points and liberation contributes 0.4 percentage points to the total return rate of 7.8 percentage points. Thus, the effect of having a partner in Ukraine accounts for 33 percent of the total returns during the sample period. This aligns with the stated reasons for returning to Ukraine: 43 percent of all returnees indicated they returned to reunite with their spouse or other relatives, as shown in Table 3.C.3. Although local conflicts do not have a statistically significant impact on the overall likelihood of returning to Ukraine, they do influence decisions to settle outside Ukraine. Local conflicts account for 1.7 percentage points of the respondents newly planning to settle abroad, which represents 22 percent of the total proportion of individuals who began planning to settle outside Ukraine.

Additionally, the choice of host country strongly predicts return intentions. Countries hosting Ukrainian refugees vary significantly in terms of income levels, labor market conditions, and the generosity of welfare benefits. Our sample is too small to analyze the effects of individual destination countries. However, when we group destination countries together, results in Figure 3.C.5 reveal that respondents in Germany were 5 percentage points less likely to return to Ukraine than those in Poland and the rest of Eastern Europe, and they were more than 6 percentage points more likely to begin planning to settle outside Ukraine. Figure 3.C.2 shows that return intentions in the baseline survey were also weakest in Western European countries. These results suggest that refugees who are less willing to plan a return to Ukraine are more likely to have chosen to relocate to Western European countries. Due to endogenous sorting, the small sample size for individual countries, and the correlation of country characteristics, we do not further examine this aspect of the analysis.

We next show how the results evolve as we progressively incorporate conflict-related factors and additional controls, as shown in Table 3.5.1. All columns include socio-demographic controls. The first column in each panel includes, in addition to socio-demographic controls and initial country of refuge fixed effects, only occupation and frontline status of the home district, the second column local conflict intensity between the two waves, the third column shows these together and the fourth column includes these together and adds controls for the week of the

[18] This is not driven by sample composition. Figure 3.C.6 and Table 3.C.4 show that on the sample where return location is elicited, the estimate on return to home municipality is statistically significant, whereas return to Ukraine generally is not.

TABLE 3.5.1: The effect of conflict on return intentions

Panel A: **Returned to Ukraine**	(1)	(2)	(3)	(4)
Home district liberated	0.042*		0.051**	0.046
	(0.022)		(0.024)	(0.028)
Home district still occupied	-0.023		-0.012	-0.015
	(0.019)		(0.023)	(0.023)
Local conflict between interviews		-0.005	-0.006	-0.008
		(0.006)	(0.006)	(0.006)
Observations	2306	2301	2301	2299
R^2	0.14	0.14	0.14	0.16
Average dependent variable	0.078	0.078	0.078	0.078
Panel B: **Returned to home municipality in Ukraine**	(1)	(2)	(3)	(4)
Home district liberated	0.046		0.058*	0.047
	(0.033)		(0.031)	(0.036)
Home district still occupied	-0.040**		-0.020	-0.023
	(0.015)		(0.019)	(0.019)
Local conflict between interviews		-0.016***	-0.016***	-0.018***
		(0.005)	(0.004)	(0.005)
Observations	1436	1433	1433	1432
R^2	0.13	0.13	0.14	0.15
Average dependent variable	0.061	0.061	0.061	0.061
Panel C: **Started to plan to settle outside Ukraine**	(1)	(2)	(3)	(4)
Home district liberated	-0.029**		-0.037***	-0.027*
	(0.012)		(0.011)	(0.015)
Home district still occupied	-0.008		-0.020	-0.020
	(0.022)		(0.018)	(0.018)
Local conflict between interviews		0.009	0.014**	0.016***
		(0.006)	(0.006)	(0.006)
Observations	2306	2301	2301	2299
R^2	0.13	0.13	0.13	0.14
Average dependent variable	0.075	0.075	0.075	0.075
Baseline controls	✓	✓	✓	✓
Destination country FE	✓	✓	✓	✓
Week of interview FE				✓

Notes: This Table shows regression results of equation 3.1 for three different outcomes: a) whether someone has returned to Ukraine, b) whether someone has returned to his or her home municipality in Ukraine, and c) whether someone no longer plans to settle outside Ukraine. Standard errors, corrected for clustering at the district level, are shown in parentheses. Columns 1-3 differ only in the regressors for which point estimates are shown. Column 4 adds fixed effects for the week of the first and the last interview. For details on the full specification of column 3, see notes to Figure 3.5.1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

initial and the final interview fixed effects. The latter partial out variation due to survey timing. Throughout the columns, point estimates exhibit a notable degree of stability, with a few exceptions.

As the family situations of Ukrainians starkly differ, their reactions to positive and negative shocks in their locality of origin may also vary. Figure 3.C.7 displays a regression model from Column 3 of Table 3.5.1, augmented with an interaction between three measures of conflict and an individual-level characteristic. The results indicate that the effects of regaining control in the home district, as well as local conflict intensity, are primarily driven by those who have a partner remaining behind and by those who are tertiary educated. The influence of local conflict intensity on the likelihood of settling outside Ukraine is strongest among people with a tertiary education, possibly because they have the best labor market prospects abroad.

We also analyzed predictors of relocating to a third country, as shown in Figure 3.C.8. Individuals whose home district has been liberated are less likely to move to a third country, possibly due to an increased likelihood of returning to Ukraine. Furthermore, higher local conflict intensity correlates with a greater probability of moving to a third country, aligning with the likelihood of planning to settle abroad.

3.5.2 Integration outcomes

Theory suggests that refugees who do not intend to return invest more in acquiring host-country-specific human capital, such as language skills, and integrating into the local labor market (Chiswick and Miller, 1994). Figure 3.5.2 presents three sets of results on the relationship between return intentions and subsequent integration outcomes in terms of employment, training, subjective integration, and participation in language courses. The first set shows how initial return intentions and time since arrival predict integration outcomes without accounting for individual fixed effects. The second set is similar but includes interactions between time since arrival and initial return intentions. The third set further incorporates individual fixed effects. In this third set, initial return intentions do not serve as explanatory variables, as they are absorbed by individual fixed effects; instead, the focus shifts to how changes in integration outcomes between the first and the last wave vary based on initial return intentions, once individual idiosyncratic factors are controlled for.

FIGURE 3.5.2: The relation between initial return intentions and subsequent integration outcomes

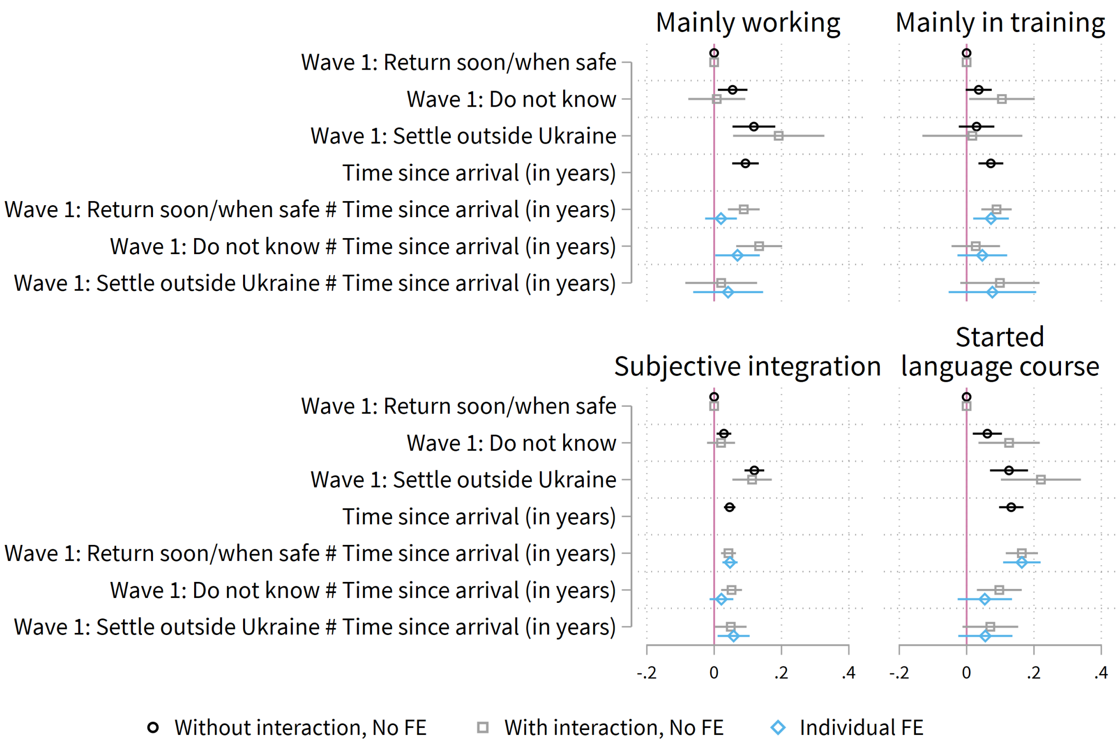

Notes: This figure shows regression results of levels of integration outcomes in wave 2 – 6 on (i) the levels of initial return intentions and the time since arrival in the first destination country and (ii) additional interactions of initial return intentions and time since arrival and (iii) individual fixed effects. We show 95 percent confidence intervals based on standard errors clustered on the district level. N = 5,200 (upper left), N = 2,765 (upper right), N = 2,856 (lower right), and N = 1,718 (lower left).

Integration outcomes reveal a distinct pattern: individuals initially planning to settle outside Ukraine are the most integrated, while those intending to return are the least integrated. This suggests that the motivation to learn new skills relevant to the host country is strongest among those planning to settle permanently abroad. These differences are most pronounced shortly after leaving Ukraine. For subjective measures of integration, effects related to initial intentions to return and time since arrival show little variation across analyses, indicating that initial differences tend to persist. In terms of language acquisition, those planning to return soon are initially less likely to start a language course but reach similar levels after one year. The pattern for employment is more complex, with analyses not adjusted for individual characteristics suggesting a 'catch-up' in employment rates between those intending to return and those undecided. However, this catch-up effect becomes less apparent after adjusting for individual characteristics.

FIGURE 3.5.3: Conflict and different integration outcomes

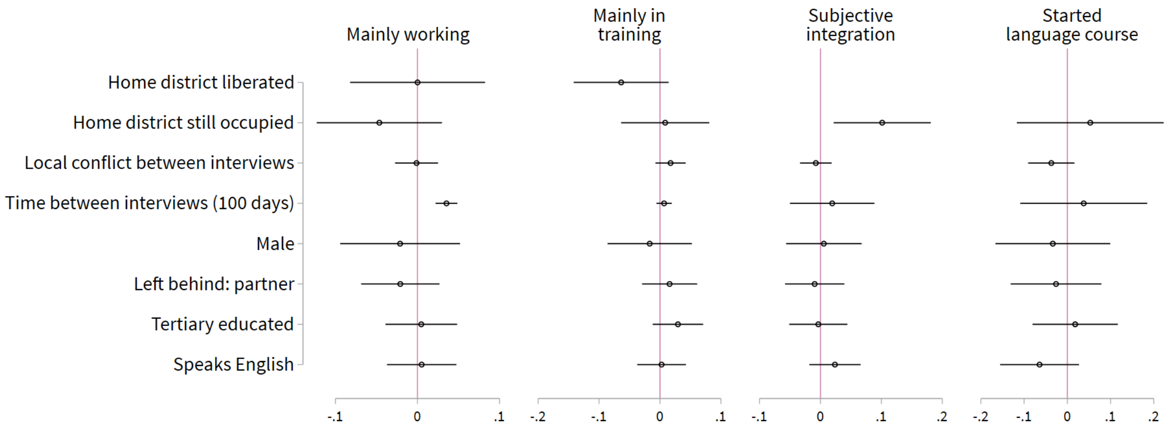

Notes: This figure shows coefficient plots of four multivariate OLS regressions. 95 percent confidence intervals are based on standard errors clustered on the district level. We restrict the sample to all respondents 25 – 59. The outcomes in the first two columns are in levels on the long differences sample, and control for initial levels of started working or not in wave 1. N = 1,966 for both. The last two columns are changes on the sample of long differences between the earliest response in wave 2 and 3 and the response in wave 6. N = 503 and N = 544, respectively. The latter two do not include estimates for "home district liberated" as no district was liberated during the sample period. All other control variables are identical to those in Figure 3.5.1. For the coefficients on the conflict-related variables, see Table 3.C.5.

The effect of conflict in one's home municipality on integration outcomes is unclear, a priori. On one hand, it might encourage investment in integration by reducing return intentions. On the other hand, more intense conflict could also lead to stress and trauma, which may negatively affect labor market outcomes, subjective integration, and investment in language skills. Figure 3.5.3 displays the regression coefficients for the model introduced in equation 3.1 across four key measures of changes in economic, subjective, and linguistic integration. Additionally, Figure 3.C.9 shows further measures of integration outcomes between the first response in Waves 2 and 3, and the response in Wave 6.[19] Because most frontline changes occurred before Waves 2 and 3, no districts were liberated between these survey waves and Wave 6; hence, this regressor is absent in columns 3 and 4 of Figure 3.5.3. It is important to note that all reported results pertain only to respondents who did not return to Ukraine.

Our results suggest that the conflict variables have no significant effect on whether refugees are employed. The liberation of one's home district appears to make refugees less likely to participate in any kind of training, which aligns with a higher likelihood of return, reducing incentives to invest further in integration in the host country (Chiswick and Miller, 1994; Cortes, 2004; Adda, Dustmann and Görlach, 2022). Conversely, if one's home district remains occupied for the duration of our surveys, refugees report a positive change in their subjective integration. This

[19]The other waves did not include all integration questions.

can be attributed to the lower return intentions among this group, which encourages investment in integration.

Conflict intensity in the home municipality does not appear to systematically affect integration outcomes. Individuals from regions with higher conflict intensity are slightly less likely to have started a language course, although this result is only statistically significant at the 10 percent level. When weighting the regressions with inverse probability and population weights, as shown in Figure 3.D.2, the effects of conflict intensity on subjective integration and language course participation appear slightly stronger.

3.5.3 Robustness

To further establish the robustness of our results, we examine how different specifications and approaches to treatment construction affect the results. First, we show that allowing for spatial correlation in the error structure does not substantially change the results in Section 3.5.1 (Figure 3.D.3). In addition, we show several alternative specifications in Figure 3.D.4. Excluding return intentions or previous conflict, or including region fixed effects and the shortest distance to Russia or the frontline during the second interview do not change the results substantially. An exception is the inclusion of region fixed effects, which changes the estimated effect of "Home district still occupied" on "Started to plan to settle outside Ukraine". The effects of conflict on integration outcomes also remain largely unchanged across specifications, as reported in Figure 3.D.5.

In addition, we show that the results are largely robust to alternative ways of constructing the treatment. In Figures 3.D.6 and 3.D.7 we use the four measures of conflict underlying the first principal component used as our primary measure of conflict, both in logs (as used in the PCA) and linearly. Although not always statistically significant, the results for conflict intensity in logs always have the same magnitude as in the PCA. The linear results give much more weight to places with very high conflict. With the exception of return to Ukraine, we find similar results for the linear specification. Despite the fact that the effect is negative and significant for return to home, the effect on return overall is positive for all four measures (and statistically significant for one). In Figure 3.D.8, we test whether the results change when we use conflict not only in the home municipality, but also in municipalities within a r kilometer radius. We find that the effect of conflict intensity becomes insignificant when using radii of 100 kilometers, confirming that we are capturing the effects of local conflict rather than a larger scale.

To confirm that the effect of conflict is not only driven by respondents from areas on or behind the frontline, we exclude the respondents from areas on the frontline or occupied in the first wave in Figure 3.D.9. We find that the effect of local conflict on returning to one's home municipality is even stronger among respondents from districts not on or behind the frontline since the first survey. Additionally, we exclude the 6 percent of respondents who left Ukraine before 2022 and the 7 percent of respondents who came from areas already occupied by Russia before 2022 in Figure 3.D.9, as these respondents may be very different from the rest of the refugee population. We find that the results hardly change, even if we exclude all three groups at the same time. Figure 3.D.10 breaks down the effects of different types of events as recorded in ACLED. The results point in a similar direction; airstrikes, which are able to target locations far behind the frontline, most strongly reduce return.

3.6 Beyond local conflict: the role of expectations about the war

Expectations about the outcome of the war can play an important role in return plans and in investments in host-country specific human capital. To quantify these effects, we collected panel data on individuals' expectations about the outcome of the war. Initially, most Ukrainians were very optimistic about Ukraine's chances of winning the war without ceding any territory. From September 2022 to January 2023, 71 percent expected Ukraine to win and liberate all occupied territories by the end of 2024. In October and November 2023, this decreased to only 35 percent. Figure 3.6.1 shows how expectations about the outcome of the war by the end of 2024 changed between the survey waves from September 2022 to January 2023 and the survey wave in October and November 2023. Ukrainian refugees became considerably more pessimistic about a quick victory, and more likely to expect the war to still be ongoing at the end of 2024. Despite the increase in pessimism, the propensity to return to Ukraine or make plans to do so — either immediately or when conditions are deemed safe — decreased only slightly on the same sample, from 66 percent to 54 percent. At the same time, the proportion of respondents who have decided to settle outside Ukraine has increased from 13 percent to 20 percent.

Figure 3.6.2 documents that expectations about the outcome of the war have become less optimistic over time across survey waves in a close-to-linear fashion. The share of Ukrainian refugees who plan to return or have already returned has also declined over time, but at a much lower rate. While the shares planning to return and expecting Ukraine to win by the end of 2024 without losing any territory were about the same at the beginning, by November 2023 the share expecting Ukraine to win by the end of 2024 without losing any territory was considerably lower.

FIGURE 3.6.1: Sankey diagram of changes in expectations about the outcome of the war until the end of 2024

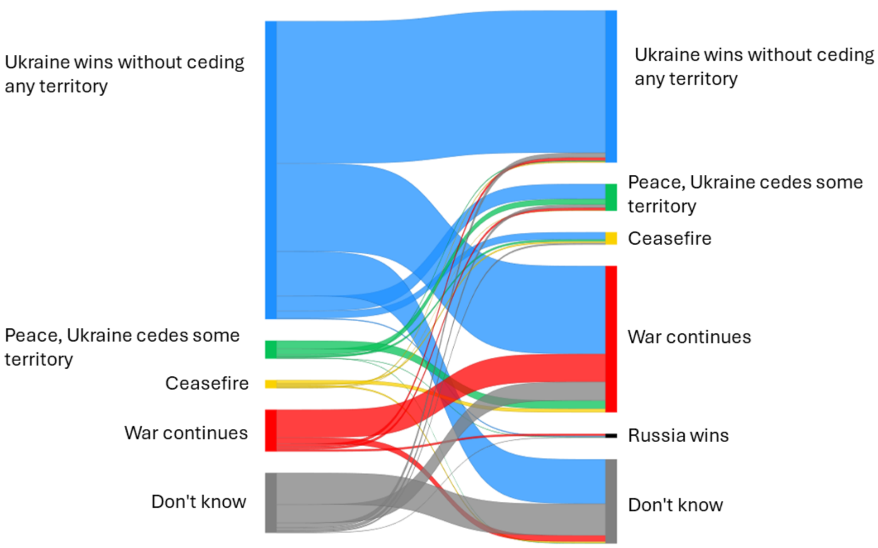

Notes: This figure shows individual-level changes in expectations about the outcome of the war by the end of 2024. The sample consists of differences between the first interview in the second and third wave (September 2022 – January 2023) and the interview in the sixth wave (October – November 2023). The average number of days between the interviews is 285 days. N = 834.

Return intentions are strongly correlated with the expected outcome of the war, as shown in Figure 3.C.10. Only 9 percent of respondents who expect Ukraine to win by the end of 2024 plan to settle outside Ukraine. In contrast, 21 percent of those expecting territorial concessions or a ceasefire, and 26 percent of those who anticipate the war lasting until 2025 or longer, plan to settle outside Ukraine.

To test the strength of the relation between changes in expectations and our main outcomes of interest, we regress returning to Ukraine, planning to settle outside Ukraine, and being employed on changes in expectations about the war. As the share of respondents expecting fighting to end by the end of 2024 with ceasefire or Ukraine ceding any territory is small throughout our survey period, we pool the changes in Panel A of Figure 3.C.10 in three categories: whether

FIGURE 3.6.2: The percentage of people expecting Ukraine to win by the end of 2024 gradually decreases over time

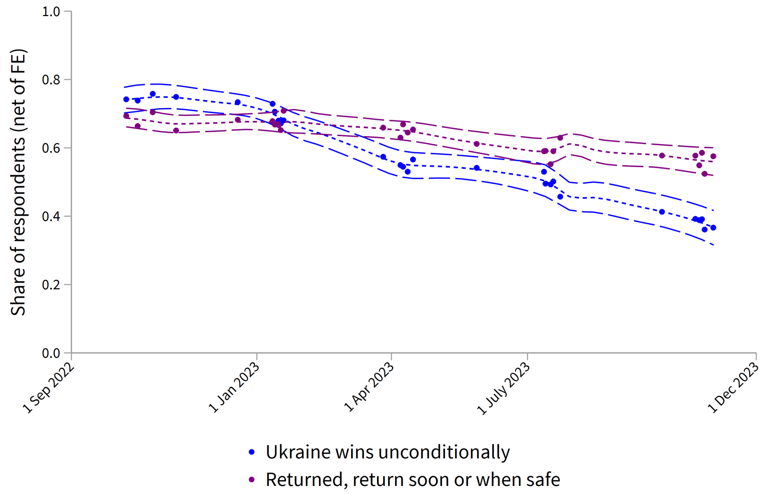

Notes: Binned scatterplot with non-parametric trend of the share of respondents expecting Ukraine to win the war by the end of 2024 over time and the total share of respondents who returned, plan to return soon or when safe, net of individual fixed effects, with 90 percent confidence interval. The binned scatterplot is based on 20 bins. For an explanation of the construction of this Figure, see notes to Figure 3.3.1. Based on waves 2 – 6. N = 5,669.

one always thought that Ukraine would win and liberate all occupied territories by the end of 2024 or newly thinks so in the last wave (the reference category), whether the respondent no longer expects that Ukraine would win and liberate all occupied territories by the end of 2024, and whether the respondent never thought that Ukraine would win and liberate all occupied territories by the end of 2024. We show the coefficients from the regression analysis in Table 3.6.1. We find that negatively updating war expectations or always being pessimistic increases plans to settle outside Ukraine by almost 5 percentage points. This is a sizeable effect, as the sample mean is 7.4 percent. Surprisingly, negatively updating expectations or always having pessimistic expectations does not correlate to returning to Ukraine. This could either indicate that there is no effect of expectations on return, or that any such effect is offset by returnees negatively updating their expectations. However, Table 3.C.6 shows that return to Ukraine in a previous wave does not predict more pessimistic expectations about the outcome of the war, after controlling for prior expectations. We do not find statistically significant effects on employment.

Expecting Ukraine no longer to win by the end of 2024 contributes to 1.4 percentage points of the increase in respondents planning to settle outside Ukraine. Furthermore, those who have consistently been pessimistic about Ukraine's victory by the end of 2024 are more likely to have started planning to settle outside the country than those who have been consistently optimistic. This difference accounts for about 2 percentage points in the share of respondents who are newly planning to settle permanently outside Ukraine. These findings indicate that changes in expectations about the war's outcome explain twice as much of the shift toward settling outside Ukraine (46 percent) as the intensity of local conflict does (22 percent; see Section 3.5.1).

TABLE 3.6.1: The relation between changes in expectation and changes in return intentions

	(1) Returned to Ukraine	(2) Started to plan to settle outside Ukraine	(3) Found work
Does not think anymore Ukraine would win	0.009	0.047***	0.027
	(0.016)	(0.016)	(0.020)
Never thought Ukraine would win	-0.009	0.048***	0.008
	(0.013)	(0.016)	(0.019)
Time between interviews (100 days)	0.023***	0.010*	0.023***
	(0.006)	(0.005)	(0.007)
Observations	1668	1668	1668
R^2	0.087	0.078	0.009
Average dependent variable	0.070	0.074	0.118

Notes: Regression results of changes in return intentions on changes in expectations about the outcome of the war. The reference group is "Always thought or newly thinks Ukraine would win". For the full set of control variables, see notes to Figure 3.5.1. The sample is composed of long differences between the first and the last interview among waves 2 – 6, on the sample of individuals who answered at least two follow-up waves 30 days apart. The outcomes in the first two columns are identical to those in Figure 3.5.1, whereas the third column is a binary indicator for whether the individual is working in the later wave, but not in the earlier. The mean number of days between survey waves is 242. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.7 What could explain strong return intentions? Evidence from Ukrainians in Ukraine

Since refugees were only surveyed after leaving Ukraine, analyzing their responses alone is insufficient to assess the relative importance of the mechanisms discussed in the introduction—increased confidence in the government and military, optimism about the future, and strong national identity. To further evaluate the mechanisms behind high return intentions, we analyze representative surveys of Ukrainians in Ukraine. We utilize Gallup World Polls and surveys conducted by the Ilko Kucheriv Democratic Initiatives Foundation (IKDIF) in collaboration with the Razumkov Center, both before and during the full-scale war in 2022 and 2023.

Figure 3.7.1 illustrates the competing forces that influence Ukrainians' decisions to emigrate. A traditional perspective, represented by a red arrow, suggests an increased propensity to emigrate and a decreased willingness to return due to the costs of conflict exposure. In contrast, blue arrows represent factors that counteract this trend: strengthened national identity, increased trust in government and military capabilities, and greater optimism. These mechanisms are likely to influence both the decision to emigrate and the decision of refugees to return.

Although the figure focuses on Ukraine, the interplay of national identity, trust in government and military, and optimism could similarly affect responses to conflict in various contexts. In some cases, the influences of these factors—national identity, confidence in the government and military, and optimism about the future—might be reversed. For example, in the context of a civil war, all these factors could encourage emigration and discourage returning.

3.7.1 The full-scale war reduced desire to emigrate from Ukraine

The Gallup World Poll (GWP) is an annual survey conducted across more than 150 countries that provides nationally representative data. Each year, the GWP conducts repeated cross-sectional surveys, interviewing around 1,000 individuals per country on a variety of topics. These surveys

FIGURE 3.7.1: The effect of the Russian invasion on the desire to live in Ukraine is ambiguous.

gather information on topics such as migration aspirations, views and attitudes towards the government, and the socio-demographic characteristics of respondents. The advantage of the Gallup World Polls is their ability to compare changes in attitudes in Ukraine with those in other countries, including Russia and other Eastern European countries.

Figure 3.7.2A illustrates the share of individuals desiring to emigrate permanently over time as recorded in the Gallup World Polls for Ukraine, Russia, and five other country groups. Between 2015 and 2021, the share of the population with a desire to emigrate was higher in Ukraine than in any other European country group. Following the large-scale Russian military offensive, there was a notable shift in the emigration intentions among Ukrainians: the share of Ukrainians wanting to emigrate permanently dropped from 35.3 percent in July 2021 to 9.5 percent in September 2022, before rising slightly to 12.9 percent by August 2023.

In 2022 and 2023, a lower share of respondents expressed a desire to emigrate from Ukraine than from Western and Northern European countries. Figure 3.E.3 illustrates that the decrease in the desire to emigrate from Ukraine from 2021 to 2022 is the largest year-to-year change ever recorded in the Gallup World Poll (GWP). Importantly, this decrease cannot be attributed to selective outmigration. Figure 3.E.4 demonstrates that the decrease remains similar even after accounting for the observable characteristics of the respondents.

Furthermore, Appendix 3.A.3 discusses four scenarios of selective outmigration based on outmigration rates and responses to the Verian survey. Figure 3.E.5 indicates that in all four scenarios, only a small part of the decline in Ukrainians' desire to emigrate can be explained by the selective out-migration of those with a stronger desire to emigrate. Specifically, even under the assumption that all emigrants from Ukraine had a preference to live abroad, the decrease in emigration desire between 2021 and 2022 remains substantial at 19 percentage points.

3.7.2 The role of confidence in government and military, optimism, and national identity

The finding that Ukrainians' desire to emigrate permanently decreased significantly during the war starkly contrasts with theories of international migration that consider conflict a major push

FIGURE 3.7.2: Desire to emigrate, optimism, and confidence in government and military

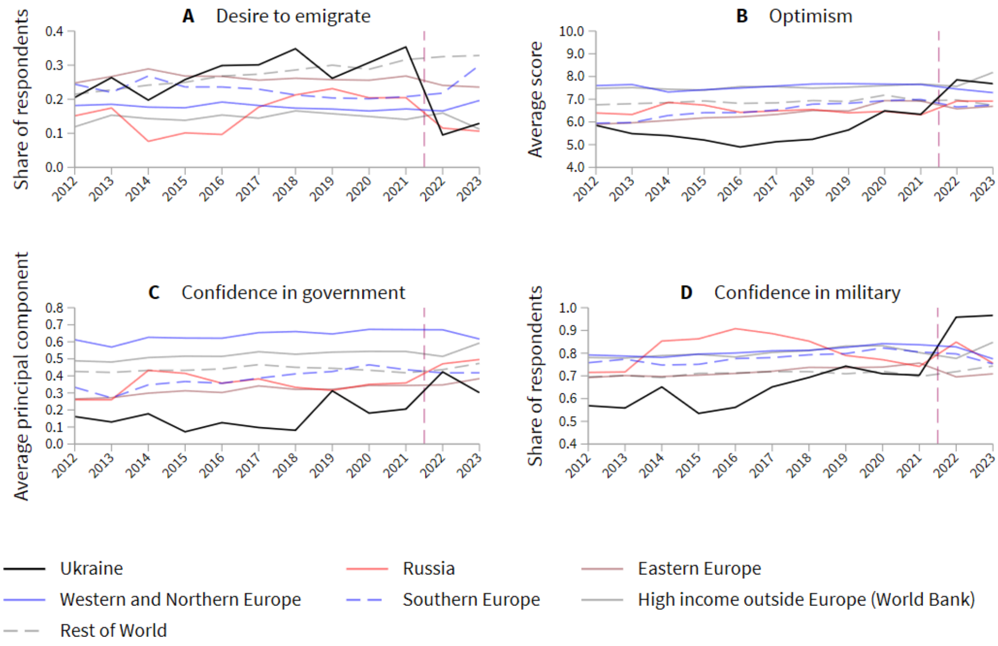

Notes: (A) illustrates the share of respondents desiring to emigrate. (B) shows the average optimism score which measures how good the respondent expects their life to be in absolute terms in five years (ranging from 0 for the worst possible to 10 for the best possible life). Importantly, this is asked in absolute terms, not relative to the current situation. (C) shows the measure of confidence in the government as constructed in (Guriev, Melnikov and Zhuravskaya, 2021). The separate components of this measure are shown in Figure 3.E.2. (D) shows the share of respondents who have confidence in the military. We exclude respondents from regions in Ukraine in all years that were partially or fully occupied by Russia during the 2022 and 2023 survey interviews.

factor for emigration (Massey et al., 1993; Bohra-Mishra and Massey, 2011; Adhikari, 2013). Concurrently, Panels B-D of Figure 3.7.2 reveal that Ukrainians have grown more optimistic about life in Ukraine and more confident in the government and military. The year-on-year increases in optimism and confidence in both the government and the military are also exceptionally large, as shown in Figure 3.E.3.

An Oaxaca-Blinder decomposition, presented in Table 3.E.1, suggests that these increases account for much of the change in the desire to emigrate between 2021 and 2022. Increases in confidence in government, confidence in the military, and overall optimism explain 41 percent of this gap, while other covariates account for only 5 percent, leaving 54 percent unexplained. Interestingly, these patterns remain similar when comparing 2021 to 2023, although the share explained by these three factors drops to 32 percent. A notable difference is that the increase in confidence in government explains half as much in 2023 as in 2022, aligning with literature that suggests rally-around-the-flag effects typically only temporarily boost government confidence (Mueller, 1970; Dinesen and Jæger, 2013).

The unexplained gap might be partially due to a stronger national identity. To explore this possibility, we utilized additional survey data from the Razumkov Center, which inquires about respondents' pride in being Ukrainian and their migration plans. These data reveal that 55 percent of Ukrainians were very proud to be Ukrainian in August 2022, and 50 percent felt the same in August 2023, an increase from 27 percent in 2021 (see Figure 3.E.6). Consistent with findings from the Gallup World Poll, most respondents express a desire to build their future in Ukraine, as shown in Figure 3.E.7. National identity is strongly correlated with plans to build a future in

Ukraine, as indicated in Figure 3.E.8. Assuming the relationship between national identity and plans to build a future in Ukraine remains constant, the observed increase in national identity could explain 22 percent of the decline in the desire to emigrate.

3.8 Conclusion

We analyzed the return intentions and integration outcomes of Ukrainian refugees in a panel survey across Europe. In the baseline survey in 2022, the majority of respondents planned to return to Ukraine soon or when it is safe, and 10 percent planned to settle permanently outside Ukraine. Return intentions are remarkably stable and strong predictors of actual return. Among all respondents, the realized return rate was 2.7 percentage points and the net increase in plans to settle outside Ukraine was 1.6 percentage points over 100 days. High local conflict intensity between survey waves deters refugees from returning to their home municipality, but not to Ukraine as a whole. The liberation of one's home district increases the probability of returning to Ukraine by 5.1 percentage points. Those refugees who planned to settle outside Ukraine integrate faster, but subsequent conflict intensity in the home municipality and the liberation of the home district do not change most integration outcomes. However, there is suggestive evidence that the liberation of one's home district reduces the likelihood of having training as one's main activity. This is consistent with higher return intentions reducing incentives to invest in integration in the host country.

What explains the remarkably high intention to return among Ukrainian refugees? At the individual level, family ties are an important reason for return. Most men between the ages of 18 and 60 are not legally allowed to leave Ukraine. Our back-of-the-envelope calculation suggests that having a partner left behind can explain one third of realized returns, which is consistent with self-reported main reasons for return. At the individual level, becoming more pessimistic is associated with a 4.7 percentage point higher probability of planning to settle outside Ukraine. Changes in war expectations are not associated with changes in employment. At the population level, data from the Gallup World Poll show that the desire to emigrate permanently among Ukrainians living in Ukraine has also declined sharply. The share of Ukrainians desiring to emigrate permanently dropped from 35 percent in 2021 to 9 percent in 2022. Selective emigration of refugees can explain only a small part of the observed decline. Among Ukrainians who remain in Ukraine, increased confidence in the government, confidence in the military, optimism, and a stronger national identity play an important role in explaining the observed decline in the desire to emigrate.

Our findings provide insights into the causal effects of conflict on refugee return and integration outcomes, offering guidance for policymakers in refugee-hosting countries. The data show that Ukrainian refugees exhibit a significantly higher desire to return compared to other refugee groups previously studied in Germany, regardless of their length of stay. For instance, only 7 percent of Syrian refugees registered in the Middle East and North Africa have voluntarily returned to Syria by March 2023 (UNHCR, 2023c), and only 14 percent of Syrian refugees in Germany express a desire to return to Syria if it becomes as safe as before the civil war (Al Husein and Wagner, 2023). A critical difference is that Ukrainians were fleeing a democratic country facing external aggression, while most refugees flee internal conflicts or persecution by their own governments. This distinction has vital implications for host countries as they develop plans to support voluntary returns and formulate integration policies for future conflict scenarios. It underscores the importance of enhancing government legitimacy, fostering refugees' trust in their government, and promoting national identity to attract refugees back to their home countries.

Despite the potential victory in the war against Russia, Ukraine faces considerable challenges.

The Ukrainian population was declining even before the Russian invasion, with deaths outnumbering births annually since 1991. Moreover, confidence in the judiciary remains low and corruption is pervasive, factors that could deter returns. The critical challenge for Ukraine is to leverage the common purpose fostered by the war to drive broader institutional and cultural changes. By addressing these push factors, Ukraine can make returning more appealing and stabilize its demographic trends.

Appendix

3.A Detailed Description of Data

3.A.1 Verian Survey "Voice of Ukraine"

The survey includes a wide range of background variables relating to demographics, employment status, and municipality of origin. Importantly, to elicit return intentions we ask individuals the following question on return intentions in every wave:

Return intentions *What are your plans regarding returning back to Ukraine?* With the following answer options:

- I intend to go back very soon
- I intend to go back at some point later when I feel it is safe to return
- I do not intend to go back and plan to settle outside Ukraine
- Do not know yet
- Prefer not to answer

In addition to aforementioned question and various demographic variables, we use several other questions directly in the main text:

Started working *Did you start working in the country you are currently residing in, yes or no?*

Current location and return *In which country are you currently located?* [drop-down menu]

Respondents answer this question from a list of countries. In the second wave, this list also includes Ukraine, which enables us to identify those who have returned to Ukraine.

In the follow-up waves, we ask several additional questions:

Expectations about the outcome of the war *What do you find the most likely outcome of the war by the end of 2024?* With the following answer options:

- Ukraine wins and Russia withdraws from all territory it currently occupies
- Ukraine cedes some territory to Russia as part of peace agreement
- There is ceasefire
- Russia wins and annexes big parts of Ukraine
- The war continues
- Do not know
- Prefer not to answer

Expectations about the duration of the war *When do you expect the war in Ukraine to end?* With the following answer options:

- Within 3 months
- In 4 to 6 months
- In 7 to 12 months
- In 1-2 years
- I expect the war to continue more than 2 years
- Do not know
- Prefer not to answer

Work-related integration (wave 1 and 6) *Did you start working in the country you are currently residing in?* With the following answer options:

- Yes
- No

Work-related integration (wave 2) *Which of these descriptions best apply to what you have been doing for the last four weeks?* With the following *non-exclusive* answer options:

- In paid work – working remotely in Ukraine (employee, self-employed, working for your family business)
- In paid work – working in the current country of residence (employee, self employed, working for your family business)
- In any kind of schooling or training (including language courses)
- Unemployed and actively looking for a job
- Unemployed and not actively looking for a job
- Doing unpaid housework, looking after children or other persons

Work-related integration (wave 3-6) *Have you started working?* With the following *exclusive* options:

- In paid work – working remotely in Ukraine (employee, self-employed, working for your family business)
- In paid work – working in the current country of residence (employee, self employed, working for your family business)
- In any kind of schooling or training (including language courses)
- Unemployed and actively looking for a job
- Unemployed and not actively looking for a job
- Doing unpaid housework, looking after children or other persons

Language course participation *Have you started a course to learn the language of your host country?* With the following answer options:

- Yes

- No

We use the following two questions on self-assessed language skills in a principal components analysis:
Linguistic integration I *Please evaluate your own language skills in your current country of residence. I can understand the main points in simple newspaper articles on familiar subjects when reading in the local language.* With the following answer options (recoded to):

- Very well (4)
- Well (3)
- Moderately well (2)
- Not well (1)
- Not well at all (0)

Linguistic integration II *Please evaluate your own language skills in your current country of residence. In a conversation, I can speak in the local language about familiar topics and express personal opinions.* With the following answer options (recoded to):

- Very well (4)
- Well (3)
- Moderately well (2)
- Not well (1)
- Not well at all (0)

We use the following three questions on subjective integration in a principal components analysis:
Subjective integration I *How often do you feel like an outsider in the current country of residence?* With the following answer options (recoded to):

- Never (4)
- Rarely (3)
- Sometimes (2)
- Often (1)
- Always (0)

Subjective integration II *How well integrated do you feel in the city/town you currently live in?* With the following answer options (recoded to):

- Not at all integrated (0)
- Very little integrated (1)
- Moderately integrated (2)
- Integrated (3)
- Well integrated (4)

Subjective integration III *Do you feel welcome in the city/town you currently live in?* With the following answer options (recoded to):

- Never (0)
- Rarely (1)
- Sometimes (2)
- Often (3)
- Always (4)

Social integration with Ukrainians *How many Ukrainian friends/family members do you have in the city you currently live in?* With the following answer options:

- Nobody
- 1-2
- 3-5
- 6-10
- 11-15
- 16-20
- More then 20

Social integration with locals *How many friends among* ***Nationality of country of residence*** *do you have in the city you currently live in?* With the following answer options:

- Nobody
- 1-2
- 3-5
- 6-10
- 11-15
- 16-20
- More then 20

Data cleaning and processing

To determine an individual's place of residence before they evacuated during the war, the baseline wave of the survey asks: (i) which region they lived in before February 24, 2022, and (ii) the specific locality through a write-in field. Eighteen percent of respondents did not answer the latter question. To match individuals with the municipality (hromada) of their residence before the war, we utilize geospatial data on Ukraine's administrative divisions as of 2020 from the United Nations Office for the Coordination of Humanitarian Affairs (UN OCHA) (United Nations, 2023). The average municipality has 30,800 inhabitants, and the median has 13,200. Larger cities comprise a single municipality. The spatial files encompass all 1,469 municipalities (hromada), nested in 137 districts (raions) across 27 regions and six macro-regions. These localities generally align with administrative divisions, but for 550 individuals, the localities had to be manually matched to municipalities within the specified regions. Localities were classified into municipalities using the Ukrainian government website `https://gromada.info/`. However, since not all region-municipality pairs are unique, we were unable to assign a unique municipality to 12 respondents, and thus classified their municipality as missing.

TABLE 3.A.1: Survey waves, number of respondents and timing

	Number of responses	First month	Last month
Wave			
1	11,783 (6,299 panellists)	June 2022	December 2022
2	1,005	September 2022	December 2022
3	1,610	January 2023	January 2023
4	1,411	April 2023	April 2023
5	1,218	July 2023	July 2023
6	1,175	October 2023	November 2023
Total	18,202	June 2022	November 2023

Notes: Number of respondents by wave and first and last interviews month per wave.

TABLE 3.A.2: Number of waves per respondent

	Number of respondents
Number of waves	
1	9,067
2	1,048
3	586
4	385
5	441
6	256
Total	11,783

TABLE 3.A.3: Demographic characteristics of the baseline and long differences sample compared to Temporary Protection beneficiaries

	Dataset		
	Baseline	Long differences	TPS (Eurostat)
Female	0.88	0.88	0.78
18 - 34	0.26	0.27	0.38
35 - 64	0.65	0.68	0.53
65 and older	0.08	0.06	0.08
Czechia	0.05	0.04	0.10
Germany	0.23	0.25	0.18
Italy	0.05	0.05	0.03
Poland	0.28	0.27	0.36
Spain	0.04	0.04	0.04
Other	0.33	0.33	0.29
N	11,783	2,674	4,377,305

TABLE 3.A.4: Demographic characteristics of the baseline sample and the Ukrainian population before 2022

	Male		Female	
	Verian	GWP	Verian	GWP
	mean	mean	mean	mean
Age 16-24	0.14	0.12	0.06	0.10
Age 25-34	0.26	0.24	0.30	0.14
Age 35-44	0.18	0.23	0.21	0.15
Age 45-54	0.13	0.17	0.19	0.15
Age 55-59	0.04	0.06	0.08	0.09
Age 60-65	0.16	0.06	0.08	0.12
Age 65+	0.09	0.13	0.08	0.24
With partner	0.67	0.60	0.54	0.49
With children under 18	0.31	0.37	0.56	0.33
Tertiary educated	0.57	0.19	0.66	0.16
From urban settlement	0.67	0.45	0.70	0.40
Observations	1359	1260	9996	1776

Notes: Descriptive statistics for the baseline sample (weighted with population weights) compared to the Gallup World Polls (weighted with survey weights) in Ukraine 2019 – 2021, for men and women separately.

TABLE 3.A.5: Demographic characteristics of the baseline and long differences estimation samples

	Baseline		Long differences	
	Mean	S.D.	Mean	S.D.
On the frontline	0.19	0.39		
Behind the frontline	0.06	0.24		
Home district liberated			0.08	0.27
Home district still occupied			0.18	0.39
Local conflict between interviews			0.00	1.00
Days since arrival	185	71		
Time between interviews (100 days)			2.68	1.39
Male	0.11	0.31	0.11	0.31
Left behind: partner	0.25	0.43	0.27	0.44
Left behind: children	0.17	0.38	0.17	0.37
With children under 18	0.57	0.50	0.61	0.49
Tertiary educated	0.66	0.47	0.71	0.45
Speaks English	0.39	0.49	0.47	0.50
Spoke destination-country language upon arrival	0.25	0.19	0.25	0.19
Age 16-24	0.05	0.21	0.04	0.20
Age 25-34	0.20	0.40	0.22	0.41
Age 35-44	0.27	0.45	0.32	0.47
Age 45-54	0.22	0.41	0.23	0.42
Age 55-59	0.08	0.27	0.06	0.23
Age 60-65	0.10	0.30	0.09	0.28
Age 65+	0.08	0.27	0.05	0.21
Married or partner	0.56	0.50	0.57	0.49
From an urban settlement in Ukraine	0.76	0.43	0.77	0.42
Continued job in Ukraine remotely	0.15	0.36	0.17	0.38
Left before the 24th of February	0.06	0.24	0.06	0.24
Survey language: Russian	0.10	0.29	0.07	0.26
District occupied before 2022	0.03	0.16	0.03	0.18
Population municipality [1000s]	1059.38	1070.95	1078.35	1083.12
Wave 1: Return soon	0.07	0.25	0.05	0.23
Wave 1: Return when safe	0.58	0.49	0.57	0.50
Wave 1: Settle outside Ukraine	0.08	0.28	0.10	0.30
Wave 1: Do not know	0.24	0.43	0.26	0.44
Wave 1: Prefer no answer	0.03	0.16	0.02	0.15
Observations	9052		2301	

Notes: Descriptive statistics of all regressors in the baseline and long differences estimation sample for whom information on the municipality of origin is non-missing. Local conflict between interviews is standardized and therefore has mean zero and a standard deviation of one.

FIGURE 3.A.1: Number and sampling rate of Ukrainian refugees

(A) Number of Ukrainian refugees across Europe

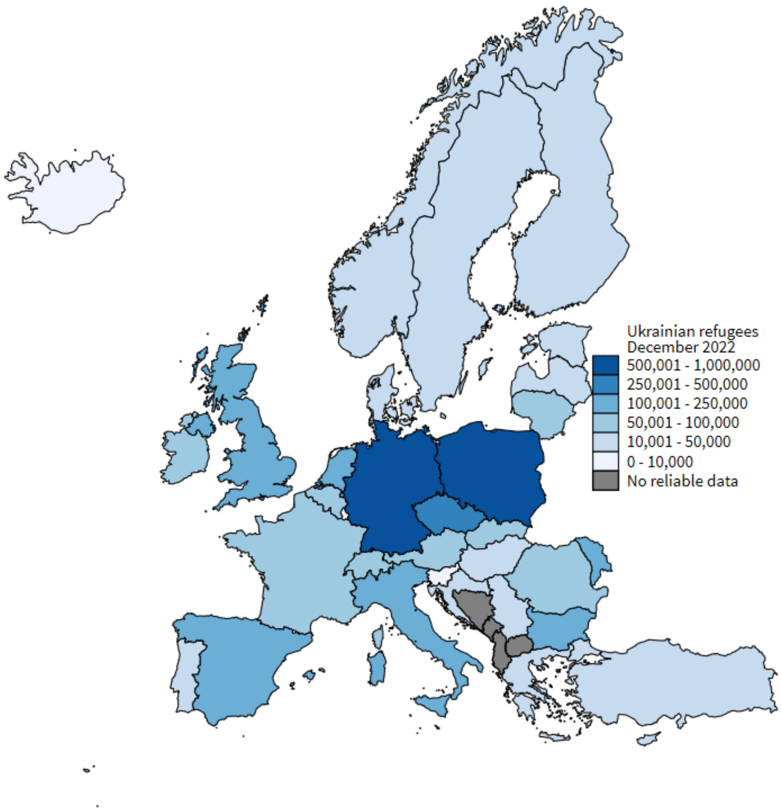

(B) Sampling rate of Ukrainian refugees across Europe

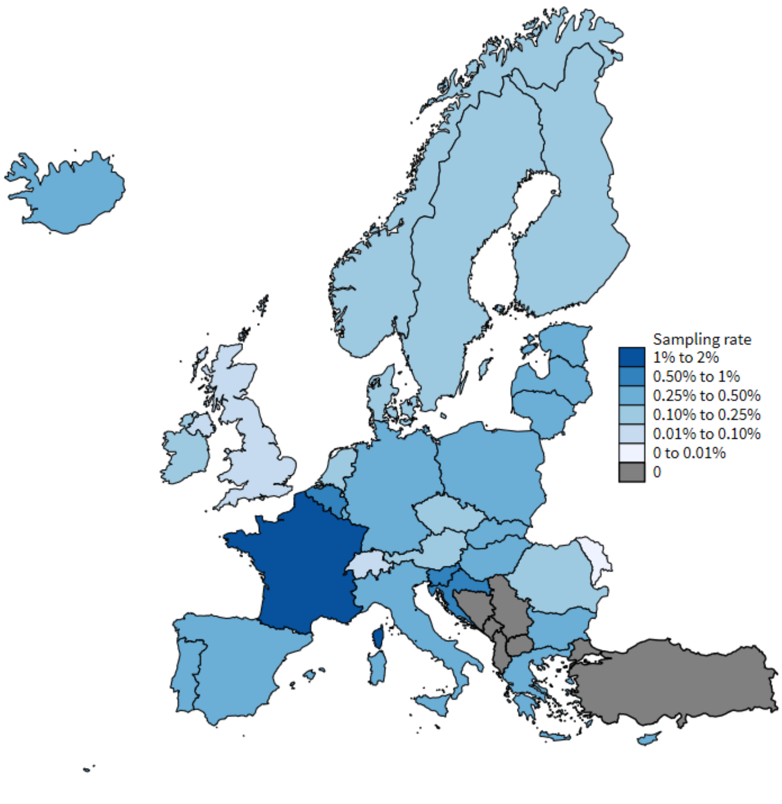

Notes: Panel (a) shows the number of Ukrainian refugees who are beneficiaries of Temporary Protection Status by December 2022, by host country. Data from the Eurostat table *migr_asytpsm*. Panel (b) shows the sampling rate of Ukrainian refugees from across European countries. Obtained by dividing the total number of respondents in the baseline wave by initial destination country by the total beneficiaries in December 2022 from the Eurostat table *migr_asytpsm*.

FIGURE 3.A.2: Distribution of the dates of interview

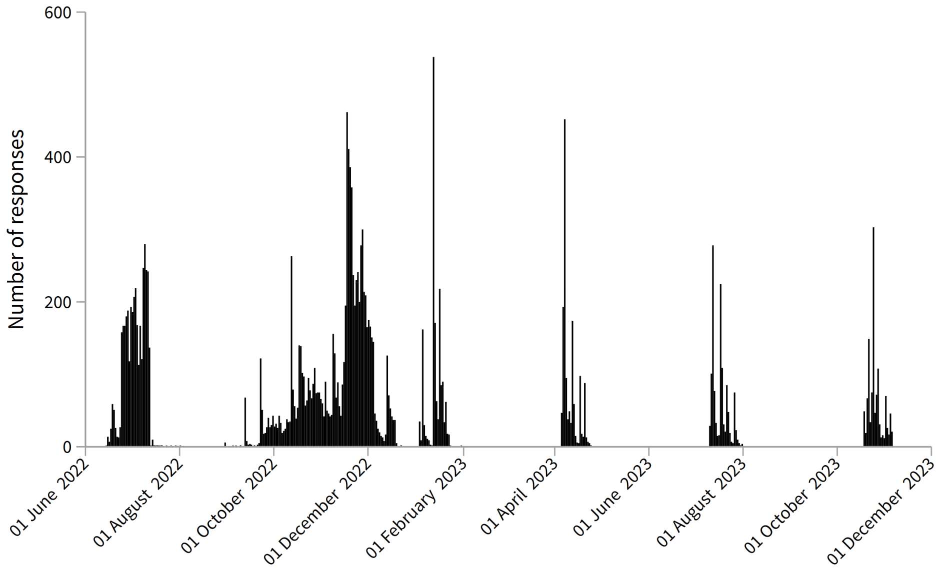

Notes: Temporal distribution of interviews of all six waves. Every bin represents one day. N = 18,202.

FIGURE 3.A.3: Origin municipalities of respondents

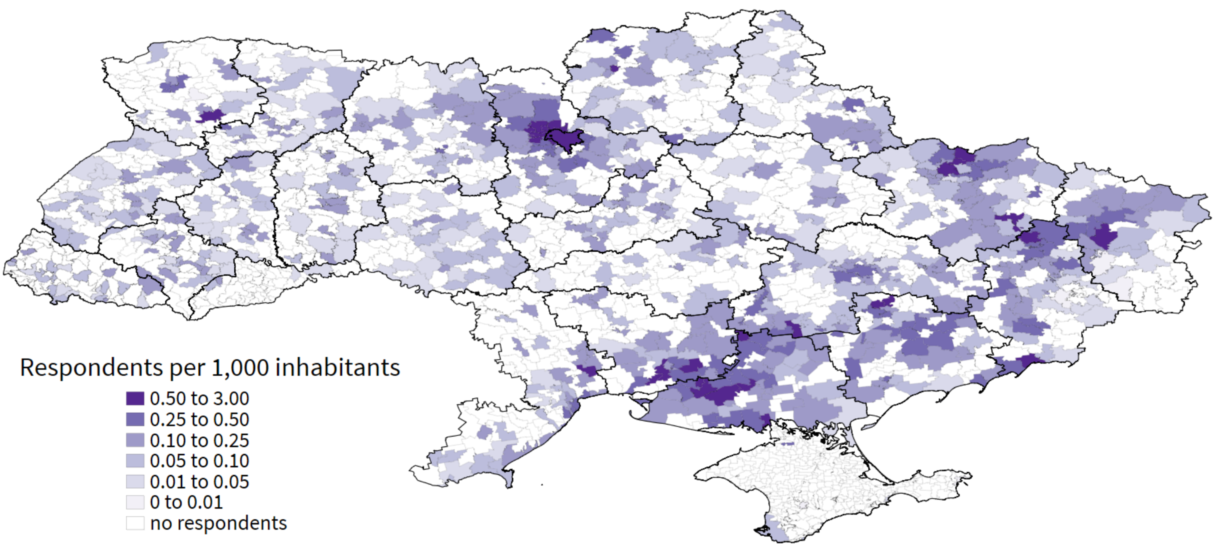

Notes: Distribution of respondents by municipality of origin in Ukraine. Excludes those respondents for whom no home municipality could be uniquely determined. N = 9,655.

3.A.2 Gallup World Poll

We use the Gallup World Polls (GWP) conducted from 2012 to 2023 to gather data on Ukrainians' emigration aspirations, perceptions, and confidence in the Ukrainian government and other institutions. The GWP annually surveys a nationally representative sample of approximately 1,000 individuals in a majority of countries worldwide. Interviews in Ukraine after Russia's full-scale attack were conducted in September 2022, and in July and August 2023. These surveys allow us to compare Ukrainian refugees surveyed in the Verian survey with Ukrainians who remained in Ukraine.

Additionally, the presence of numerous pre-war waves and the ability to compare these with other countries provide valuable insights. For our analysis, we restrict the sample to the 142 countries (including Ukraine) that were surveyed by Gallup both in 2012 and in 2022 after February 24, and/or in 2023. We categorize all other countries, except Ukraine and Russia, into five mutually exclusive groups. As GWP does not conduct surveys in every country every year, we interpolate missing values at the country level before aggregating them into country groups. For Ukraine, the only missing data of interest was the question regarding the desire to emigrate in 2020. Specifically, we investigate this desire using the following question:
Desire to emigrate *Ideally, if you had the opportunity, would you like to move permanently to another country, or would you prefer to continue living in this country?* With the following answer options:

- Move permanently
- Prefer to stay
- Don't know
- Refused to answer

We recode the desire to emigrate to 1 if the respondent answers "yes", to 0 if the respondent answers "no" or "don't know," and to missing if the respondent answers "prefer not to answer". In addition to the aforementioned question and various demographic variables, we use several other questions from GWP:
Optimism *Please imagine a ladder, with steps numbered from 0 at the bottom to 10 at the top. The top of the ladder represents the best possible life for you and the bottom of the ladder represents the worst possible life for you. Just your best guess, on which step do you think you will stand in the future, say about five years from now?*
Confidence in Government We follow Guriev, Melnikov and Zhuravskaya (2021) and perform a Principal Components Analysis (PCA) and obtain the first principal component of the following four questions with yes, no, do-not-know, and refuse-to-answer as answer options:

- *In this country, do you have confidence in each of the following, or not? How about national government?*
- *In this country, do you have confidence in each of the following, or not? How about judicial system and courts?*
- *In this country, do you have confidence in each of the following, or not? How about honesty of elections?*
- *Is corruption widespread throughout the government in this country, or not?*

Confidence in Military *In this country, do you have confidence in each of the following, or not? How about the military?* with yes, no, do-not-know, and refuse-to-answer as answer options.

Furthermore, we use the following question on corruption in business, which serves as a robustness test in the decomposition analysis in S3.1.2:
Corruption in Businesses *Is corruption widespread within businesses located in this country, or not?* with yes, no, do-not-know, and refuse-to-answer as answer options.

Data cleaning and processing

In further analysis, we code the desire to emigrate as 1 if the answer is "Move permanently", 0 if the answer is "Prefer to stay" or "Don't know" and we omit the individual when "Refused to answer" is chosen. In the full sample described in the next paragraph, 1.6 percent of individuals indicate "Don't know" and less than 0.2 percent of individuals indicate "Refused to answer". We adhere to the same procedure for all other variables with such an answer option structure.

For our analysis of changes by country groups in the desire to emigrate in Figure 3.7.2, we limit the sample to the 142 countries that were visited by Gallup in 2012 and in 2022 after February 24 and/or in 2023.

As GWP does not visit every country each year and the question on the desire to emigrate is not included in all country-years, we linearly interpolate the share of respondents who would like to emigrate on the country-year level for the 302 missing observations (out of 1,563) in Figure 3.7.2. To obtain yearly averages for each country group, we take the unweighted mean of the (interpolated) country-level averages. For Ukraine, only the 2020 values are interpolated because the question on the desire to emigrate was not asked.

The method of contacting respondents changed over the years in Ukraine. Until 2019 surveys were conducted face-to-face, in 2020 and 2021 by landline or mobile phone, and in 2022 only by mobile phone. In 2019, 90.4 percent of respondents either had a landline connection or a mobile phone and in 2021 99.2 percent of respondents indicated that they used a mobile phone for making phone calls. This suggests that a mobile phone-based sampling approach is able to reach a closely comparable sample of respondents as in 2021. As in 2022 respondents were contacted via mobile phone, also Internally Displaced Persons (IDPs) are included, although we have no way of identifying them. In all years, respondents could answer the survey in either Ukrainian or Russian.

In all analyses, we weight observations by nationally representative weights supplied by Gallup to calculate statistics as representative as possible. Gallup's weights variable reflects the inverse probability of selection, calculated using respondents' information and (among others) national demographics, number of phone connections per household and the number of household members.

As some explaining factors used in the Oaxaca-Blinder (OB) decomposition in Table 3.E.1 have missing responses (e.g., because of answering "don't know" or "refused to answer" on some items), the sample is limited to those respondents without missing responses for the respective questions.

3.A.3 Selective Migration

The large drop in desire to emigrate between 2021 and 2022 could be driven by selective out-migration of Ukrainians on observable and unobservable factors that (directly) affect migration intentions.

On observables To illustrate how the desire to emigrate, optimism, confidence in government and confidence in the military in Ukraine would have altered in 2022 if the composition in terms of age-by-gender and education-by-gender would not have changed, we residualize the outcome. For each of the four outcomes, we regress the outcome on the covariates, obtain the residuals from that regression and plot the residuals over time in Figure 3.E.4. The Figure looks qualitatively similar to Figure 3.7.2.

On unobservables To understand what part of the drop could be explained by out-migration selected on unobservables, we perform a back of the envelope calculation based on the observed migration intentions in Ukraine in GWP, observed return intentions in the Verian surveys as well as UNHCR data on population movements. As the Gallup World Poll was fielded in early

September and participation was restricted to those residing in Ukraine at the time of survey, we take the information available on refugee populations on the midpoint of the interviews on 05 September 2022.

We proxy the size of the refugee populations on 05 September 2022 by the gross number of 2.4 million of border crossings to Russia and Belarus (there is no information about movements from Russia and Belarus into Ukraine) and by the 4.2 million net border crossings from Ukraine to the rest of Europe from UNHCR (UNHCR, 2023). We have no information on the return intentions of 2.4 million Ukrainians who crossed the border to Russia and Belarus.

We assume that the share of minors in both refugee populations is 37 percent, in line with the share of minors among those who were granted Temporary Protection Status by 31 August 2022 (Eurostat, 2023). The pre-war adult population of Ukraine was 33.9 million of whom an estimated 12.3 percent left the country before the GWP was fielded (4.5 percent to Russia and Belarus; 7.8 percent to the rest of Europe).

Using these numbers, we can adjust the numbers in Figure 3.7.2 for potentially selective out-migration by making various assumptions about the counterfactual desire to emigrate of the refugee population based on return intentions in the Verian survey. In the following, we analyze the following four cases:

- **Case 1** We assume that the survey is representative of the adult refugee population (including those who crossed the border with Russia and Belarus) and that only those refugees who want to settle outside Ukraine are those who would have otherwise desired to emigrate.
- **Case 2** We assume that the survey is representative of the adult refugee population (including those who crossed the border with Russia and Belarus) and that those who want to settle outside Ukraine and those who do not know where to live are those who would have otherwise desired to emigrate.
- **Case 3** We assume that the survey is representative of the adult refugee population only in the countries it covers and that those who want to settle outside Ukraine and those who do not know where to live are those who would have otherwise desired to emigrate. Furthermore, we assume that no individuals who crossed the border with Russia and Belarus plan to settle outside of Ukraine.
- **Case 4** We assume that the Verian survey is representative of the adult refugee population only in countries it covers and that those who want to settle outside Ukraine and those who do not know where to live are those who would have otherwise desired to emigrate. Furthermore, we assume that all individuals who crossed the border with Russia and Belarus plan to settle outside Ukraine.

Cases 3 and 4 represent polar opposites on desire to emigrate among Ukrainians who fled or were forcibly displaced to Russia and Belarus. Figure 3.E.5 demonstrates how the change in desire to emigrate between 2021 and 2022 would have looked for these four scenarios. We find that the observed drop of 25.8 percentage points increased to 26.0 pp in Case 1, and decreased to 22.8 pp in Case 2, 24.4 pp in Case 3 and 20.0 pp in Case 4. Even in the very conservative case 4, the drop in return intentions would still be in the 99th percentile of year-year changes shown in Fig. 3.E.5.

3.A.4 IKDIF/Razumkov Center Survey

To probe the strength of Ukrainians' national identity over time, we draw on publicly available data from a survey conducted by the Ilko Kucheriv Democratic Initiatives Foundation (IKDIF) together with the Razumkov Center in August 2022 and 2023 and in several earlier years.

The IKDIF/Razumkov Center surveyed individuals in person in August 2022 and 2023. In 2022, 2,024 individuals were interviewed, of whom 54 percent were female. In 2023, 2019 individuals were interviewed, of whom 55 percent were female. Both surveys included questions on national identity and on future plans. The former question has been asked in previous surveys since 2002, interviewing about 2,000 individuals per year.

National Identity *To what extent are you proud or not proud to be a citizen of Ukraine?* With the following answer options:

- Very proud
- Rather proud
- Hard to answer
- Rather not proud
- Not proud at all

Future plans *Would you like to build your future life in Ukraine?* With the following answer options:

- Yes, definitely
- Rather yes
- Hard to answer
- Rather not
- Definitely not

3.A.5 ISW: Frontline Data

To capture whether an individual's home region is under Ukrainian control, contested by fighting, or occupied by Russian forces, we construct a daily dataset of the position of the frontline. To construct the dataset, we draw on the (almost) daily updated maps of the war in Ukraine provided by the Institute for the Study of War (ISW) between June 2022 and November 2023 (Institute for the Study of War and AEI's Critical Threats Project, 2023). Since the start of the war, ISW has been providing reports with maps visualizing the state of the war based on publicly available information sourced from news outlets, social media, and satellite imagery. Importantly, these maps include a line approximately indicating the frontline of the conflict. The constructed dataset is on the district level (average size of 4,406 km2) rather than the municipality level (average size of 342 km2). This makes it possible to realistically capture meaningful changes in the position of the frontline with respect to the locality of origin. As municipalities are relatively small, a municipality may be liberated but an adjacent municipality could still be on the frontline. By using the district as the level of analysis, we are better able to capture whether localities' status changes from the zone of conflict to being firmly under Ukrainian control. For instance, upon the withdrawal of Russian forces and advancements achieved by the Ukrainian military, several districts in the Kharkiv region were liberated.

We proceed by classifying districts in one of the following three categories, treating districts that are divided by a large watercourse, such as the Dnipro River, as if each side were a separate district:

1. The district is marked as "Under Ukrainian Control" if the full district is under control of the Ukrainian government.

FIGURE 3.A.4: Conflict intensity on the municipality level between the first and last interview days

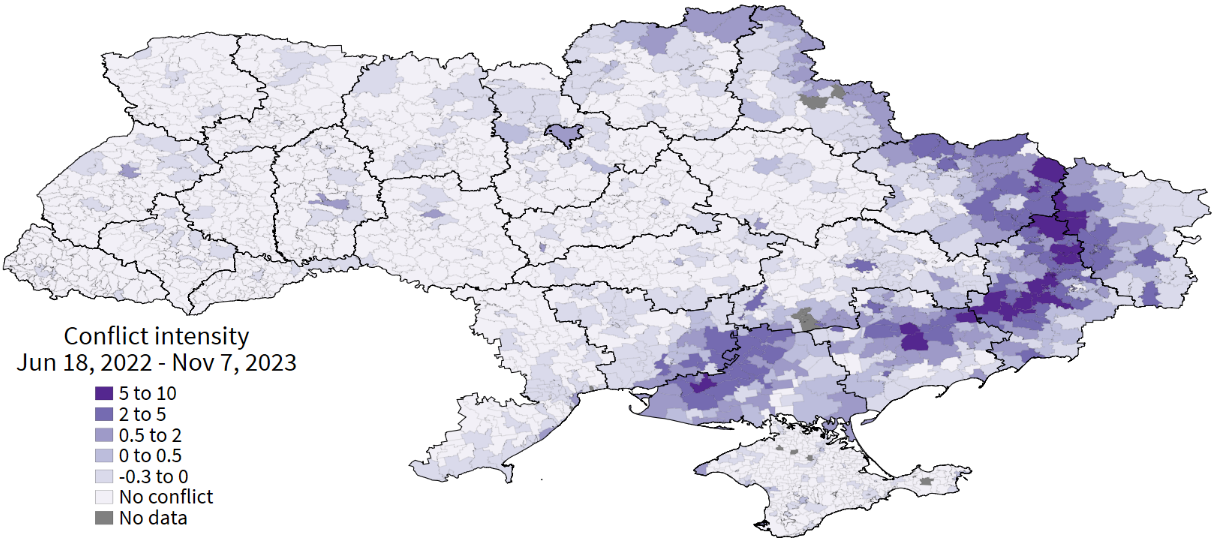

Notes: This figure shows the first principal component of conflict intensity across municipalities of origin in Ukraine.

2. The district is marked as “On the Frontline" if any area inside the district is occupied by Russian forces and any area is under control of the Ukrainian government.

3. The district is marked as “Occupied" if the full district is occupied by Russian forces.

FIGURE 3.A.5: Change in occupation states on the district level between the first and last interview days

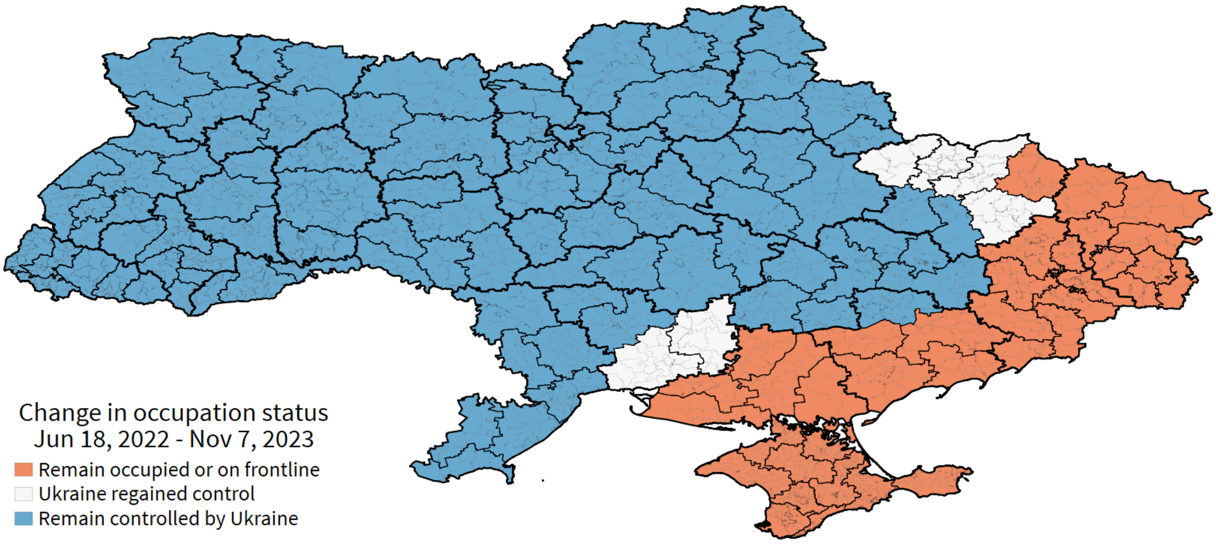

Notes: This figure shows changes in occupation status between the first baseline survey (June 18, 2022) and the last wave 6 interview (November 7, 2023) across districts of origin in Ukraine.

3.B Representativeness

TABLE 3.B.1: Predictors of follow-up response

	(1) Number of follow-up waves	(2) Responded in at least one follow-up survey	(3) Long differences	(4) Long differences + conflict
Return very soon	0.757***	0.735***	0.718***	0.775***
	(0.059)	(0.061)	(0.063)	(0.069)
Return when safe (reference)	.	.	.	.
Do not know	1.079	1.126*	1.068	1.114
	(0.063)	(0.077)	(0.061)	(0.075)
Settle outside Ukraine	1.283***	1.242**	1.276**	1.269**
	(0.101)	(0.117)	(0.127)	(0.124)
Prefer not to answer	0.827**	0.747***	0.789**	0.943
	(0.072)	(0.082)	(0.094)	(0.110)
Home district liberated				0.970
				(0.042)
Home district still occupied				1.050
				(0.101)
Local conflict intensity				1.124**
				(0.056)
Observations	10884	10884	10884	8458
Pseudo R^2	0.04	0.03	0.04	0.03
Average dependent variable	0.559	0.237	0.204	0.263
Model:	Poisson	Logit	Logit	Logit

Notes: OLS regressions of measures of follow-up survey response. Standard errors clustered at the district level are shown in parentheses. The table shows relative risk ratios from Poisson and odds ratios from logistic regressions for three different outcomes: the number of follow-up waves a respondent answers (column 1), ever responding to a follow-up survey (column 2), and the respondent being in the long differences sample (columns 3 and 4). The reference category for initial return intentions is the intention to return when safe. Baseline controls are age bins (18-24; 25-34; 35-44; 45-54; 55-59; 60-64; 65 and older), and binary indicators for sex, partnership status, tertiary education, speaking English, originating from an urban area in Ukraine, accompanied by children, having a partner left in Ukraine, having children left in Ukraine, continuing one's Ukrainian job remotely, having left Ukraine before February 24, 2022 and answering the survey in Russian. Only individuals without missing covariate values are included in the analysis. Column 4 imputes local conflict intensity for attriters by drawing a date from the empirical distribution of follow-up response dates, and calculates measures of occupation status and conflict in the same way as for respondents. As information on the home municipality is missing for 2,426 individuals, we drop these. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

3.C Additional Results on Ukrainian Refugees

TABLE 3.C.1: Share of refugees working and in a job according to qualifications by destination country group

	Number of working-age individuals	Share working	Share of those working in a job according to qualifications
Country group			
Germany	290	0.26	0.46
Rest of Western Europe	185	0.41	0.34
Poland	233	0.52	0.33
Rest of Eastern Europe	193	0.53	0.44
Southern Europe	112	0.43	0.36
Total	1,013	0.42	0.38

Notes: Share of respondents working and working in a job according to qualifications by destination country group in wave six, for respondents aged 25 – 59. Only respondents who are working are asked about whether their job corresponds to their qualifications.

TABLE 3.C.2: Stated reasons for leaving Ukraine, by gender

	Female	Male
Under direct attack	0.19	0.19
Fear for own life	0.27	0.18
Reason of leaving: fear for children's live	0.60	0.31
Among those accompanied by minor children	0.89	0.72
Life is disturbed	0.21	0.26
Fear of fighting in military	0.00	0.04
Fear of chemical or nuclear attack	0.08	0.03
Uncertainty about the future	0.07	0.12
Fear of forced displacement by authorities	0.01	0.02
Taking the opportunity	0.01	0.06
Other	0.05	0.17
Prefer not to answer	0.00	0.01
Observations	10353	1430

Notes: Reason for leaving asked from all wave 1 responses. Respondents could choose up to two reasons.

TABLE 3.C.3: Stated reasons for returning to Ukraine

	Share of respondents
Reuniting with spouse or relatives	0.43
Willingness to continue education (your or your children/ grandchildren)	0.12
Professional, work-related matters	0.17
Willingness to look after a house or other property	0.03
Homesickness	0.32
Feeling that the situation in Ukraine has stabilized	0.05
Lack of funds for living abroad	0.18
Inability to find job	0.12
Insufficient support from local authorities	0.05
Difficulty adjusting to life abroad	0.09
Feeling that I am not welcome in previous country	0.03
To join the armed forces	0.01
Other reasons	0.05
Prefer not to answer	0.01
Observations	103

Notes: Reason for return stated by returnees on the long difference sample. Respondents could choose up to two reasons.

TABLE 3.C.4: The effect of conflict on return intentions on the restricted sample (where place of return in Ukraine is elicited)

Panel A: **Returned to Ukraine**	(1)	(2)	(3)	(4)
Home district liberated	0.067***		0.079***	0.069***
	(0.019)		(0.019)	(0.023)
Home district still occupied	-0.028		-0.010	-0.015
	(0.023)		(0.031)	(0.032)
Local conflict between interviews		-0.008	-0.011	-0.012
		(0.007)	(0.008)	(0.009)
Observations	1436	1433	1433	1432
R^2	0.15	0.15	0.15	0.16
Average dependent variable	0.074	0.074	0.074	0.074
Panel B: **Returned to home municipality in Ukraine**	(1)	(2)	(3)	(4)
Home district liberated	0.046		0.058*	0.047
	(0.033)		(0.031)	(0.036)
Home district still occupied	-0.040**		-0.020	-0.023
	(0.015)		(0.019)	(0.019)
Local conflict between interviews		-0.016***	-0.016***	-0.018***
		(0.005)	(0.004)	(0.005)
Observations	1436	1433	1433	1432
R^2	0.13	0.13	0.14	0.15
Average dependent variable	0.061	0.061	0.061	0.061
Panel C: **Started to plan to settle outside Ukraine**	(1)	(2)	(3)	(4)
Home district liberated	-0.045***		-0.042**	-0.033
	(0.017)		(0.020)	(0.022)
Home district still occupied	-0.017		-0.017	-0.019
	(0.033)		(0.034)	(0.032)
Local conflict between interviews		0.007	0.012	0.015
		(0.010)	(0.011)	(0.010)
Observations	1436	1433	1433	1432
R^2	0.14	0.14	0.14	0.16
Average dependent variable	0.090	0.090	0.090	0.090
Baseline controls	✓	✓	✓	✓
Destination country FE	✓	✓	✓	✓
Week of interview FE				✓

Notes: This table shows results from OLS regressions for three different outcomes: a) whether someone has returned to Ukraine, b) whether someone has returned to his or her home municipality in Ukraine, c) whether someone no longer plans to settle outside Ukraine on the sample where place of return in Ukraine is elicited. See notes to Figure 3.5.1 and Table 3.5.1 for details on the specification. Standard errors, corrected for clustering at the district level, are shown in parentheses. N = * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 3.C.5: The effect of conflict on integration outcomes

	(1) Mainly working	(2) Mainly in training	(3) Subjective integration	(4) Started language course
Time between interviews (100 days)	0.032***	0.007	0.019	0.038
	(0.006)	(0.006)	(0.036)	(0.094)
Home district still occupied	-0.025	0.001	0.101**	0.053
	(0.040)	(0.027)	(0.042)	(0.075)
Local conflict between interviews	-0.006	0.012	-0.007	-0.037*
	(0.010)	(0.011)	(0.012)	(0.019)
Observations	2272	2272	503	544
R^2	0.36	0.26	0.16	0.22
Average dependent variable	0.325	0.224	0.024	0.061
Baseline controls	✓	✓	✓	✓
Destination country FE	✓	✓	✓	✓

Notes: This table shows regressions for four different outcomes on the long differences samples: a) Work as a main activity , b) Training as a main activity, c) Changes in the first PCA of subjective integration and d) changes in whether someone started a language course. See notes to Figure 3.5.3 for details on the specification and a coefficient plot including demographic covariates. Standard errors, corrected for clustering at the district level, are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 3.C.6: The effect of return on expectations

	Expecting Ukraine to win
	(1)
Returned during prior wave	0.035
	(0.032)
Observations	3605
R^2	0.25
Average dependent variable	0.052

Notes: This table shows results from OLS regressions of a binary indicator for expecting Ukraine to win the war until the end of 2024 on an indicator of return during the previous survey wave. We control for the levels of war expectations in the prior wave. As this sample requires information about prior return status and about expectations, dependent variable values are from waves 3 – 6 and independent variable values are from waves 2 – 5. Standard errors, corrected for clustering at the respondent level, are shown in parentheses. N = * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE 3.C.1: Predictors of baseline levels of return intentions

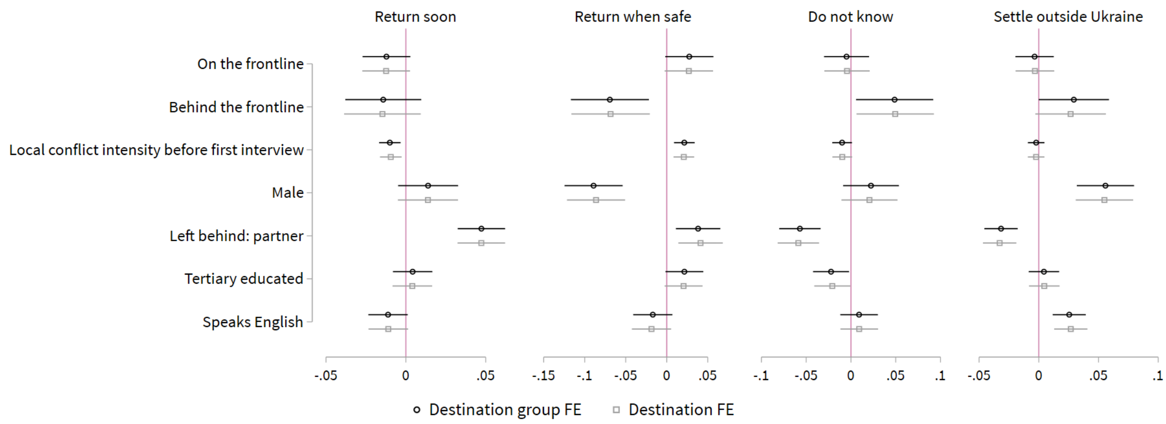

Notes: This figure shows coefficient plots of four multivariate OLS regressions using levels of return intentions as outcome variables during wave 1. We include the same control variables as in Figure 3.5.1. N = 9,041.

FIGURE 3.C.2: The role of destination countries in baseline return intentions

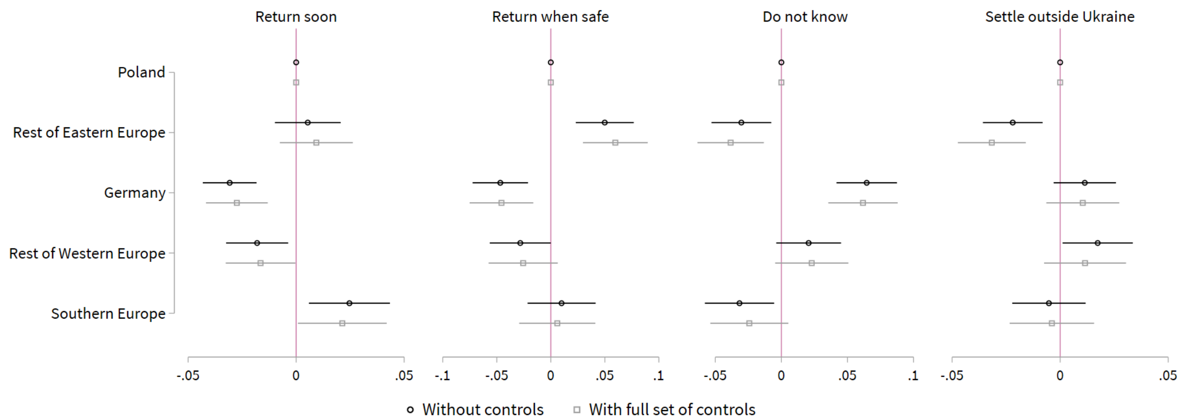

Notes: This figure shows coefficient plots of two sets of four multivariate OLS regressions of levels of return intentions as in Figure 3.C.1, replacing destination country fixed effects with country group indicators. We show the results for models without and with all controls in Figure 3.5.1. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. We include the same control variables as in Figure 3.5.1. N = 9,041.

FIGURE 3.C.3: Most refugees plan to return and return intentions are predictive of actual return

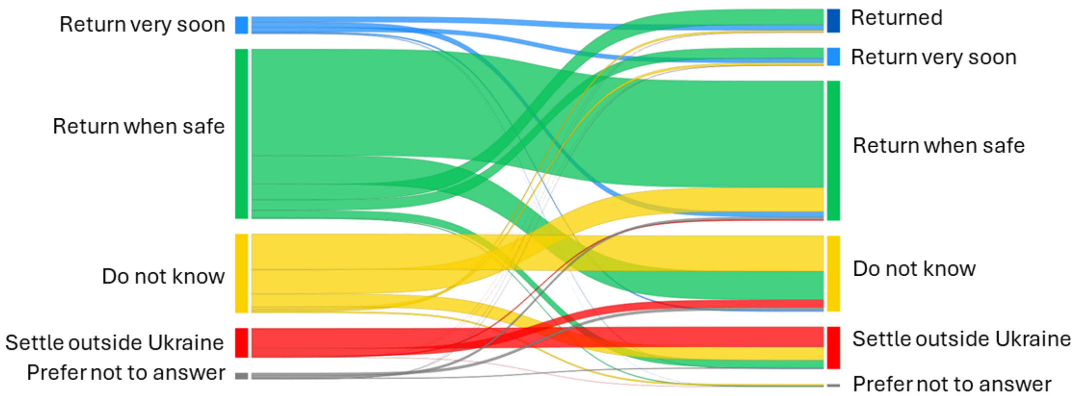

Notes: Sankey diagram of changes in return intentions between the first wave and the last wave recorded on the individual level. For an explanation of this sample, see Section 3.2. N = 2,674.

FIGURE 3.C.4: Employment by destination country

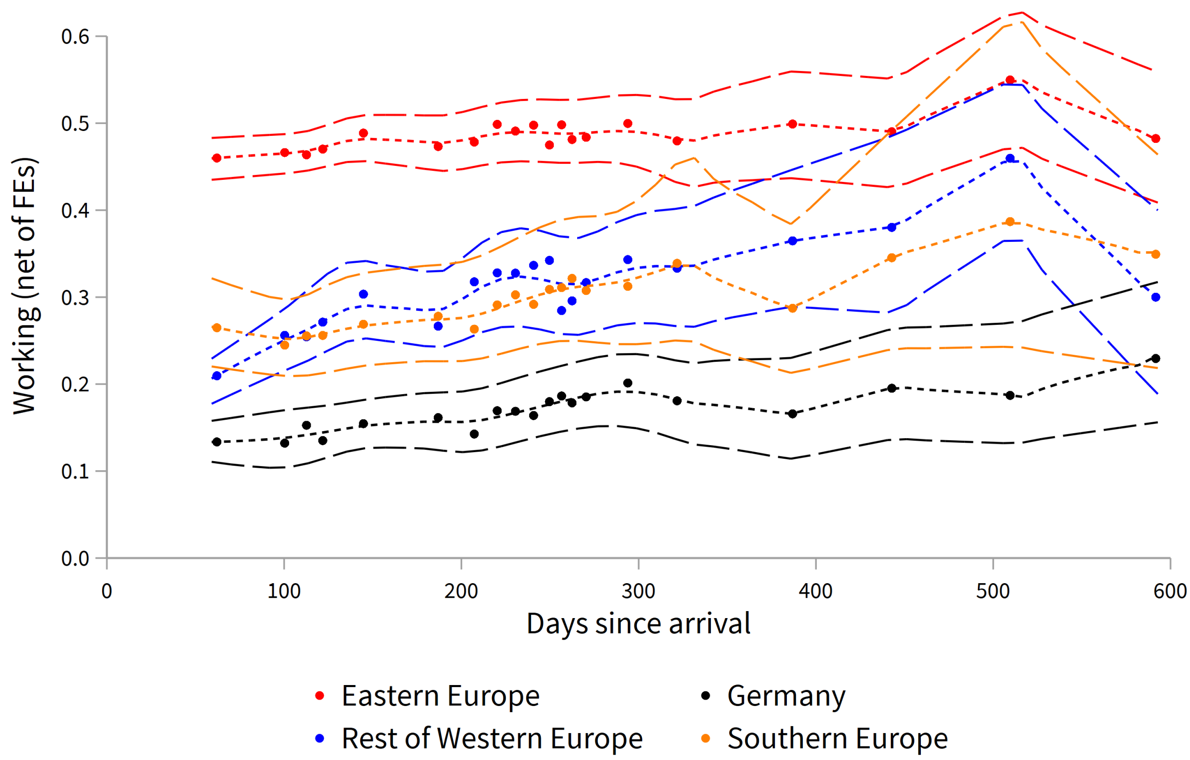

Notes: Binned scatterplot with non-parametric trend of levels of employment over time, net of individual fixed effects, with 90 percent confidence interval. To have sufficient statistical power to show patterns by country group, we combine the measures of any work and work as main activity as shown in Figure 3.3.2. To do so, we partial out level differences between any work (wave 1) and work as main activity (wave 2 to 6) through the inclusion of binary indicators when residualizing the raw outcome variables. For details on the procedure, see the notes to Figure 3.3.1. N = 3,165 (Eastern Europe), N = 1,708 (Germany), N = 742 (Rest of Western Europe), 1,318 (Southern Europe).

FIGURE 3.C.5: The role of destination countries in changes in return intentions

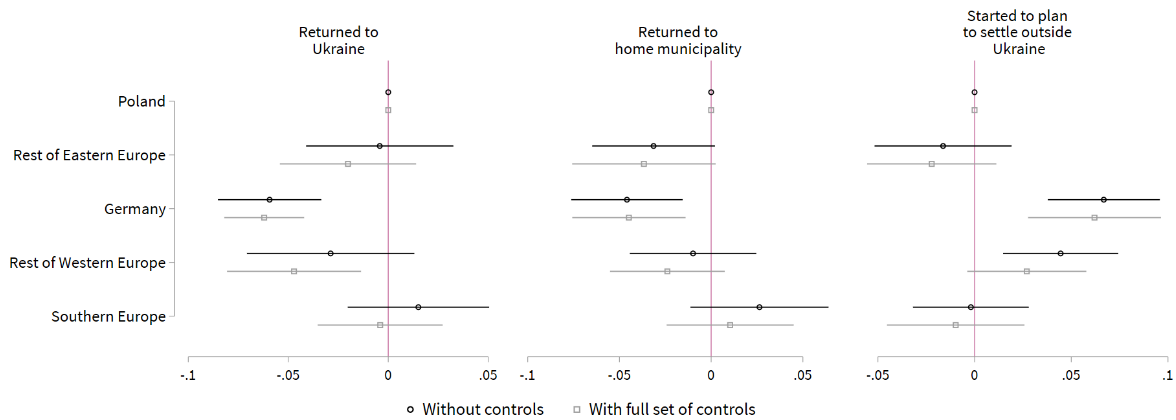

Notes: This figure shows coefficient plots of two sets of four multivariate OLS regressions as introduced in Equation 3.1, replacing destination country fixed effects with country group indicators. We show the results for models without and with all controls from Figure 3.5.1. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 2,301 (column 1 and 3), N = 1,433 (column 2).

FIGURE 3.C.6: Predictors of changes in return (intentions) on the restricted sample

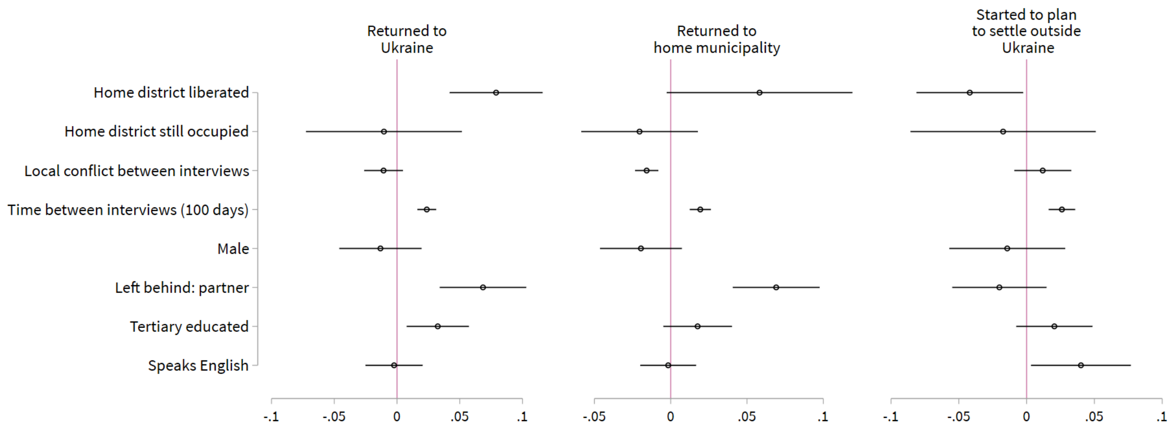

Notes: This figure shows coefficient plots of three multivariate OLS regressions as introduced in Equation 3.1, on the sample where place of return in Ukraine is elicited. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 1,433.

FIGURE 3.C.7: Heterogeneity in the effects of conflict

Notes: This Figure shows coefficient plots of multivariate OLS regressions for four dimensions of heterogeneity. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the specification, controls and fixed effects, see Figure 3.5.1. N = 1,433 (column 1); 2,301 (column 2, 3 and 4)

FIGURE 3.C.8: Predictors of moving to a third country

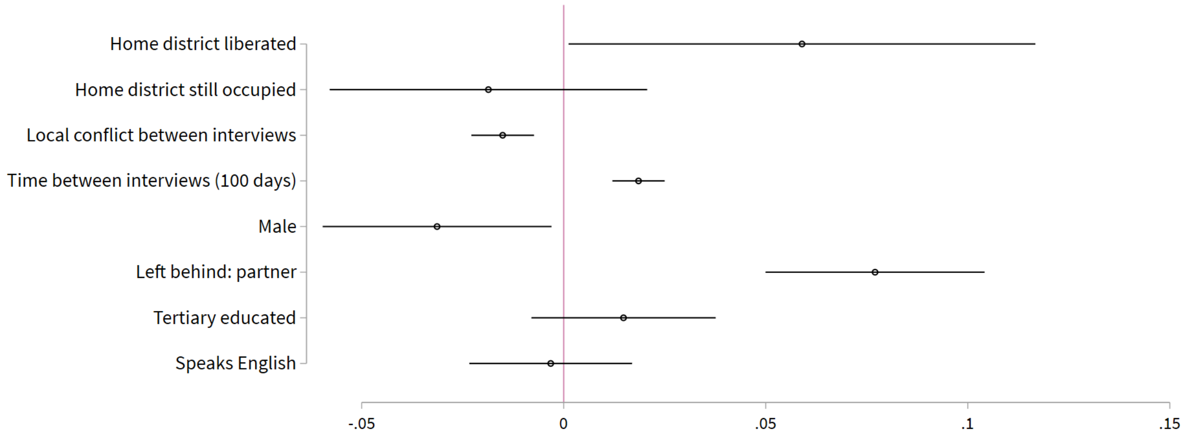

Notes: This Figure shows coefficient plots of a multivariate OLS regression of a binary indicator for moving to a different country than the initial destination country of residence or Ukraine. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the specification, controls and fixed effects, see Figure 3.5.1. N = 2,301.

FIGURE 3.C.9: Conflict and additional integration outcomes

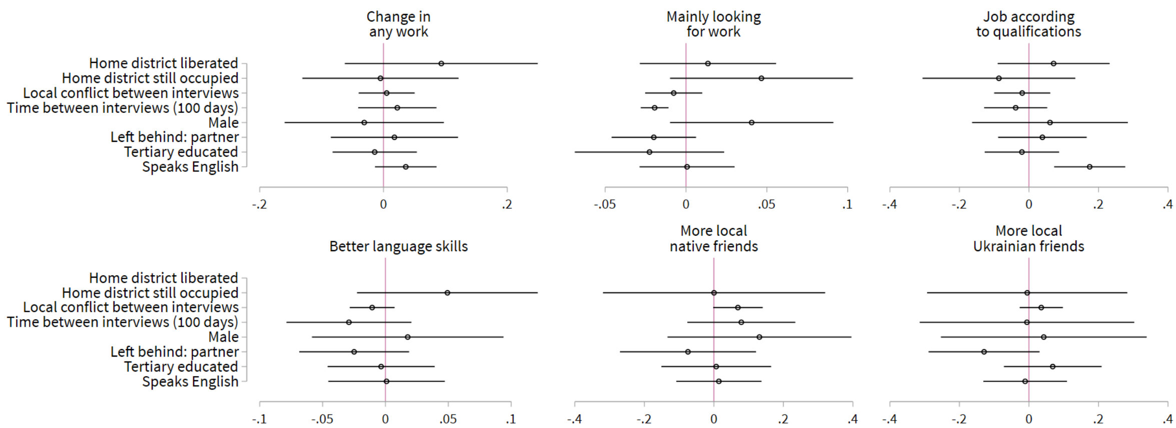

Notes: This figure shows coefficient plots of six multivariate OLS regressions. 95 percent confidence intervals are based on standard errors clustered on the district level. We restrict the sample to all respondents aged 25 – 59. "Any work" is the change in doing any work between the first and last survey of the long differences sample and can take values -1, 0 and +1 (N = 2,301). "Mainly looking for work" is a binary indicator for looking for work as a main activity in the last survey, after controlling for initial work status (N = 1,966). "Job according to qualifications" is a binary indicator for whether the respondent deems their job to fit their qualifications in wave 6 (N = 456), which is only asked to those working at the time of the survey. "Language skills" is the change in the first PCA of language skills between surveys (N = 550), "Local native friends" and "Local Ukrainian friends" are changes in the number of friends between the earliest response in wave 2 and 3 and the response in wave 6 (N = 546). The change takes value -1 if one reports less friends, 0 an equal number, and 1 if one reports more friends among natives and Ukrainians, respectively. The lower three subfigures do not include estimates for "Home district liberated" as no district was liberated during the sample period. All other control variables are identical to those in Figure 3.5.1.

FIGURE 3.C.10: Return intentions by expectations about the war.

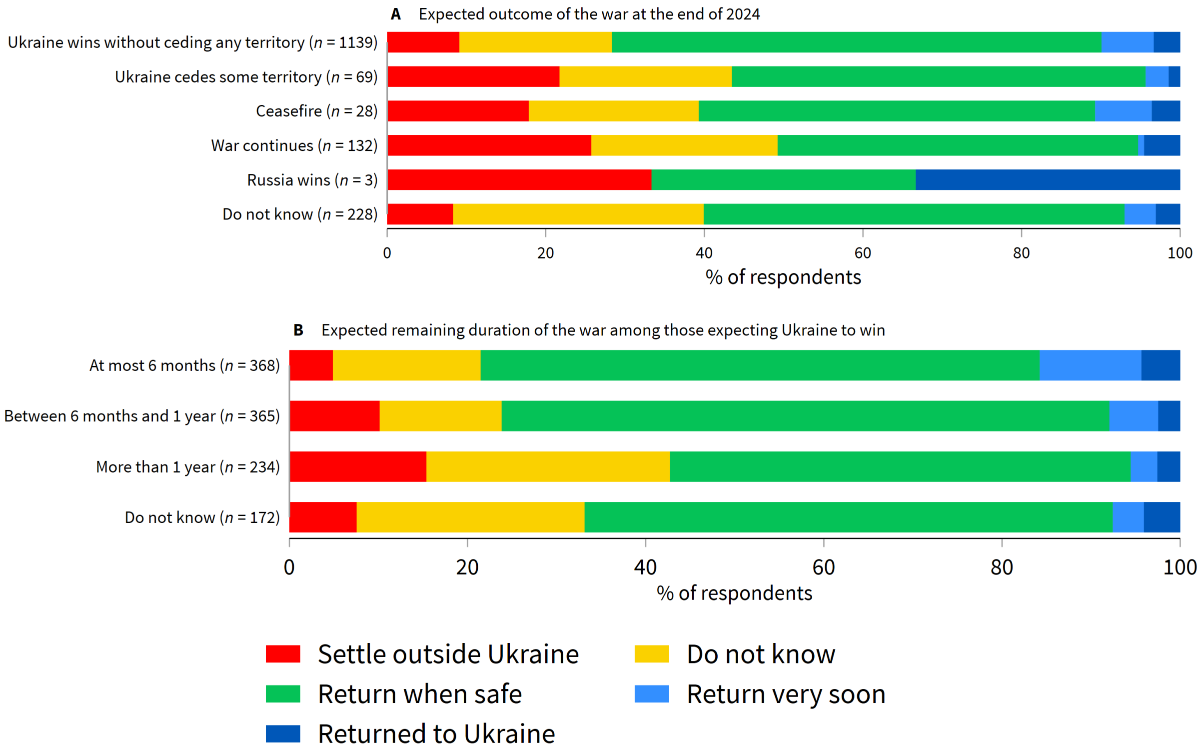

Notes: (A) shows the distribution of return intentions by expectations about the outcome of the war elicited in the second and third wave. (B) shows the distribution of return intentions by expectations about the duration of the war, contingent on expecting Ukraine to win the war without ceding territory in the third wave.

3.D Robustness

FIGURE 3.D.1: Robustness test: weighting

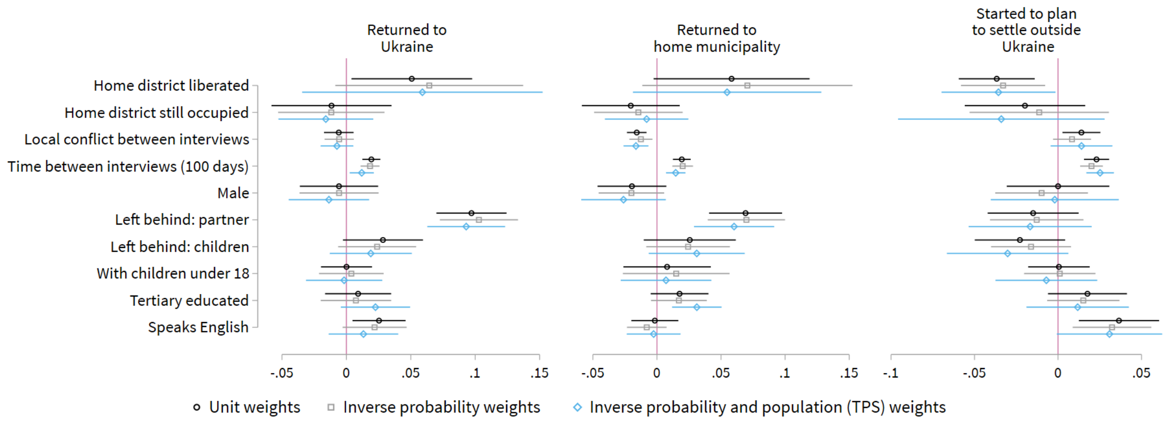

Notes: This figure shows coefficient plots of three sets of three multivariate OLS regressions as introduced in Equation 3.1, weighting the regressions. The black markers indicate unweighted regression results as in Figure 3.5.1, grey markers weighted with inverse probability weights obtained from the logistic regression in Column 3 of Table 3.B.1, and blue markers with both inverse probability weights as well as population weights as discussed in Section 3.2.4. N = 2,301 for the unweighted and IPW weighted and 2,296 for the population weighted regressions. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1.

FIGURE 3.D.2: Robustness test for integration outcomes: weighting

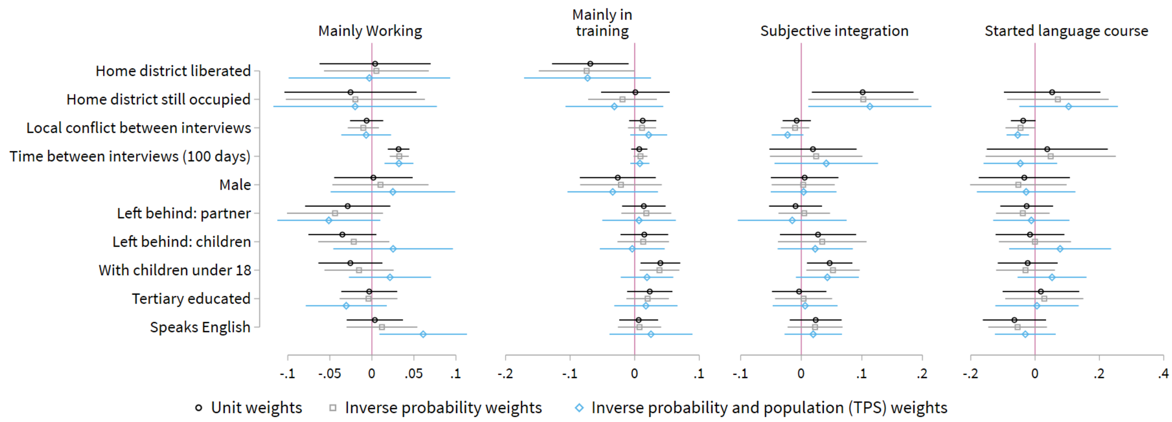

Notes: This figure shows coefficient plots of three sets of four multivariate OLS regressions as introduced in Figure 3.5.3, weighting the regressions differently. The black markers indicate unweighted regression results as in Figure 3.5.3, grey markers weighted with inverse probability weights obtained from the logistic regression in Column 3 of Table 3.B.1, and blue markers with both inverse probability weights as well as population weights as shown in Table 3.A.3. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. Number of observations for the unweighted and IPW weighted regressions from left to right: N = 1,966, N = 1,966, N = 503, N = 544, and for population weighted regressions: N = 1,962, N = 1,966, N = 499, N = 540.

FIGURE 3.D.3: Robustness test: spatially clustered standard errors

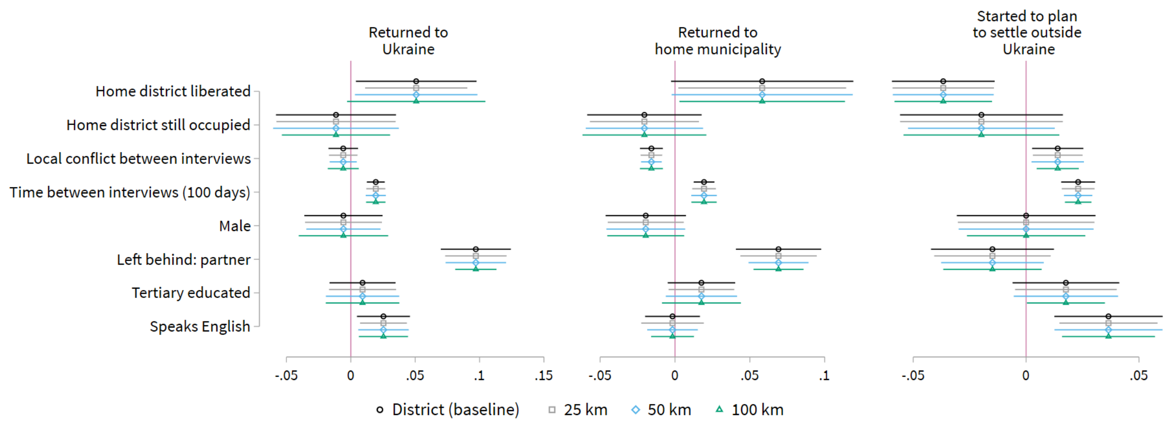

Notes: This figure shows coefficient plots of four sets of three multivariate OLS regressions as introduced in Figure 3.5.1. The Figure shows 95 percent confidence intervals constructed from standard errors allowing for arbitrary correlation at the district level (as in Figure 3.5.1), 25, 50 and 100 kilometers. Allowing for clustering at more than 100 km renders standard errors unreasonably small in some instances, suggesting a low number of effective clusters, and is therefore omitted. This is unsurprising as Ukraine's surface area is equal to about five circles with a radius of 200km. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 2,301 (column 1 and 3); N = 1,433 (column 2).

FIGURE 3.D.4: Robustness test: additional specifications

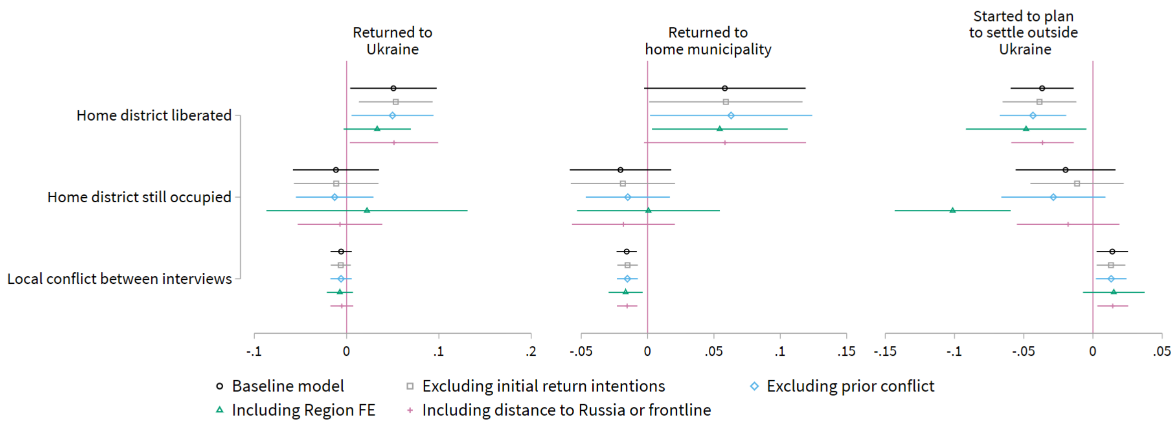

Notes: This figure shows coefficient plots of five sets of three multivariate OLS regressions as introduced in Figure 3.5.1, showing the coefficients on the conflict-related variables for four additional specifications. The first shows the baseline model, the second excludes the levels of initial return intentions, the third excludes the measure of prior conflict $Conflict_{imt_0t_1}$, the fourth includes region fixed effects and the fifth includes a measure of distance to the frontline or Russia (minimum distance) during the latest interview, for those areas under Ukrainian control. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 2,301 (column 1 and 3); N = 1,433 (column 2).

FIGURE 3.D.5: Robustness test for integration outcomes: additional specifications

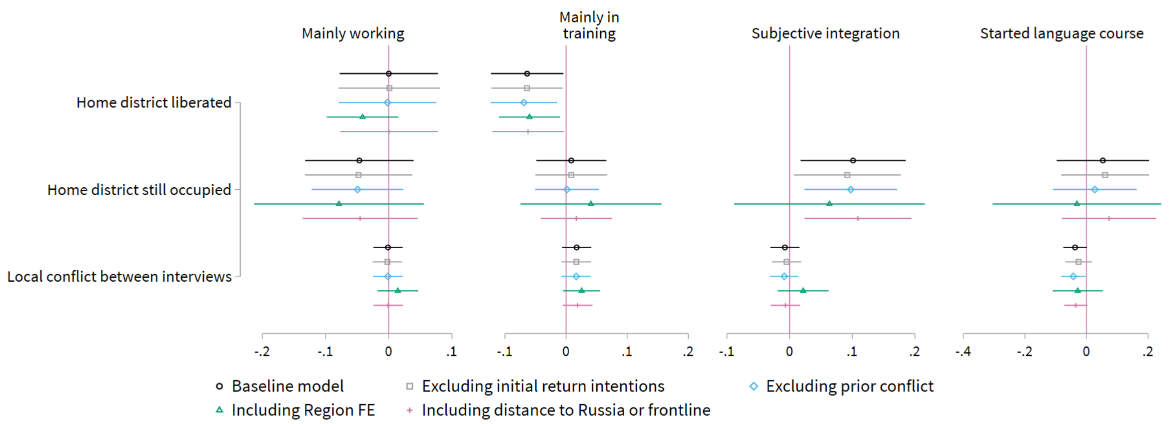

Notes: This figure shows coefficient plots of five sets of four multivariate OLS regressions as introduced in Figure 3.5.3, additionally showing the coefficients on the conflict-related variables for four additional specifications. The first shows the baseline model, the second excludes the levels of initial return intentions, the third excludes the measure of prior conflict $Conflict_{imt_0t_1}$, the fourth includes region fixed effects and the fifth includes a measure of the minimum distance to the frontline or Russia during the second wave interview, for those areas under Ukrainian control. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. From left to right: N = 1,966, N = 1,966, N = 503, N = 544.

FIGURE 3.D.6: Robustness test: independent conflict measures, logarithmic

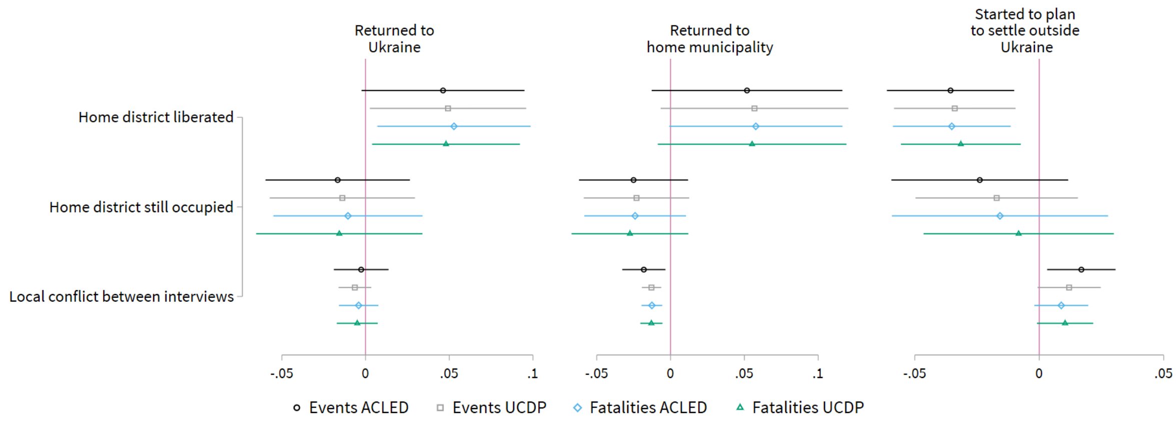

Notes: This figure shows coefficient plots of four sets of three multivariate OLS regressions as introduced in Figure 3.5.1, replacing the PCA-based measure of conflict intensity by its four individual constituents, one at the time. These are: the log (plus one) number of events in ACLED, events in UCDP, fatalities in ACLED, and fatalities in UCDP per 30 days between interviews. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 2,301 (column 1 and 3); N = 1,433 (column 2).

FIGURE 3.D.7: Robustness test: independent conflict measures, linear

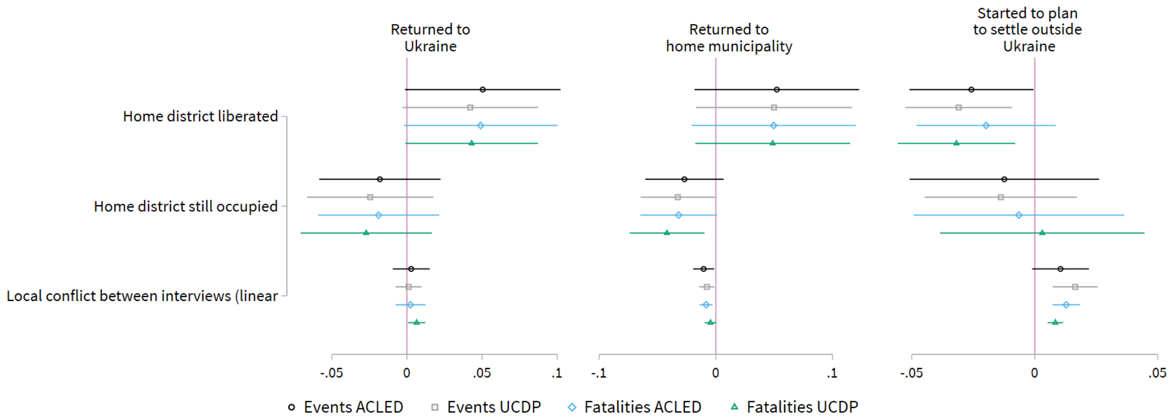

Notes: This figure shows coefficient plots of four sets of three multivariate OLS regressions as introduced in Figure 3.5.1, replacing the PCA-based measure of conflict intensity by the linear measures of four individual constituents. These are: the number of events in ACLED, events in UCDP, fatalities in ACLED, and fatalities in UCDP per 30 days between interviews. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 2,301 (column 1 and 3); N = 1,433 (column 2).

FIGURE 3.D.8: Robustness test: radius of conflict

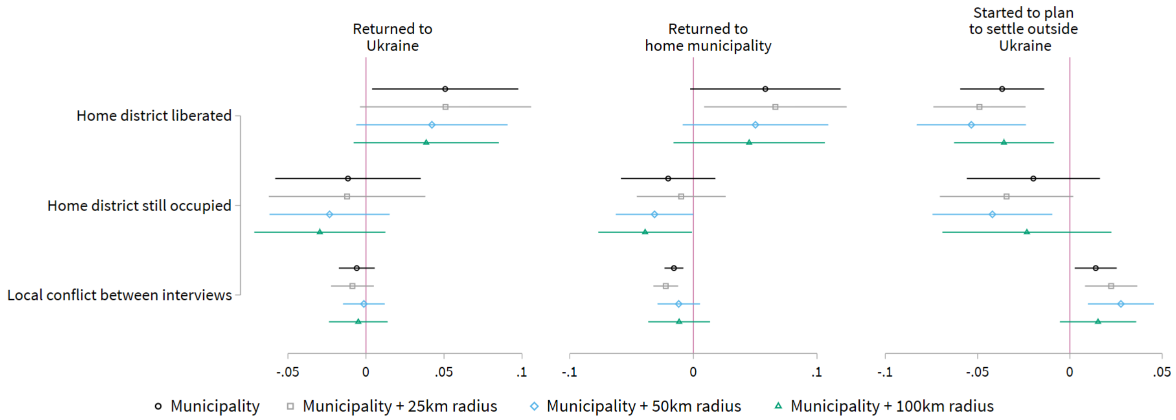

Notes: This figure shows coefficient plots of four sets of three multivariate OLS regressions as introduced in Figure 3.5.1, replacing the measure of conflict intensity with an analogous measure also including conflict in a radius around one's home municipality. We show results for conflict within municipality and within the municipality and 25, 50, and 100 kilometer radii around the municipality. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 2,301 (column 1 and 3); N = 1,433 (column 2).

FIGURE 3.D.9: Robustness test: various sample restrictions

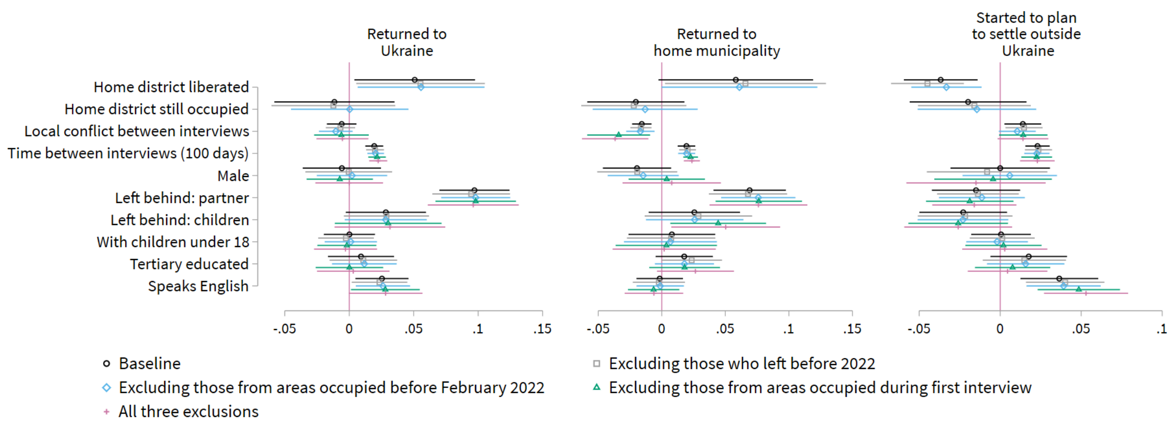

Notes: This figure shows coefficient plots of five sets of three multivariate OLS regressions as introduced in Equation 3.1. The black markers indicate estimates on the full sample as in Figure 3.5.1, grey markers indicate regressions without those who left Ukraine already before 2022 (146 individuals in column 1 and 3; 97 individuals in column 2), blue markers indicate regressions without those who are from areas that were already occupied before February 24, 2022 (77; 58), green markers indicate regressions without those who are from areas occupied during the first interview (610; 366) and pink markers indicate results omitting all three groups (732; 446). The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1.

FIGURE 3.D.10: Robustness test: ACLED casualty event types

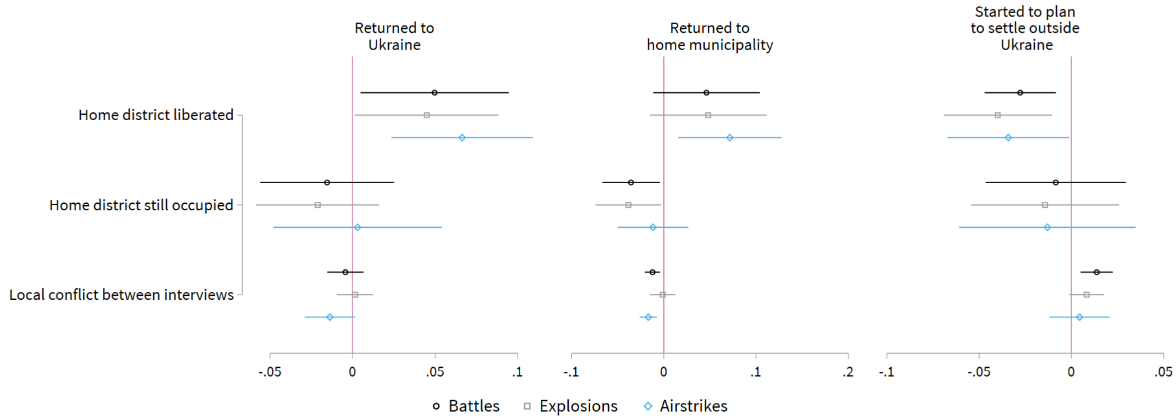

Notes: This figure shows coefficient plots of three sets of three multivariate OLS regressions as introduced in Figure 3.5.1, replacing the measure of conflict intensity with the log of the number of fatalities by event types as classified by ACLED. The figure shows 95 percent confidence intervals constructed from standard errors clustered at the district level. For details on the controls and fixed effects, see notes to Figure 3.5.1. N = 2,301 (column 1 and 3); N = 1,433 (column 2).

3.E Additional Results from Other Surveys

TABLE 3.E.1: Oaxaca-Blinder decomposition of the gap in desire to emigrate

Panel A: Levels and explained and unexplained changes in desire to emigrate				
Year	2022		2023	
	Value	S.E.	Value	S.E.
2021	.383***	.019	.383***	.019
2022	.111***	.012		
2023	.		.154***	.015
Difference	.272***	.023	.229***	.024
Explained	.125***	.019	.099***	.016
Unexplained	.147***	.028	.129***	.027
Panel B: Explaining factors				
	Explained %	S.E.	Explained %	S.E.
Female	-3.1	2.7	-3.5	3.1
Age	1.7	1.7	1.4	2.0
Female × age	4.2	3.2	4.1	3.2
Children	0.1	0.4	0.1	0.3
Tertiary education	0.0	0.3	0.1	0.3
Optimism	12.7***	3.6	10.8***	3.2
Confidence in Government	14.2**	6.0	7.6**	3.3
Confidence in Military	13.7**	4.3	13.2**	4.1
Corruption in Business	2.6	1.6	0.7	0.8
Number of observations	1,329		1,243	

FIGURE 3.E.1: Return intentions of refugees in Germany by time since arrival and origin country

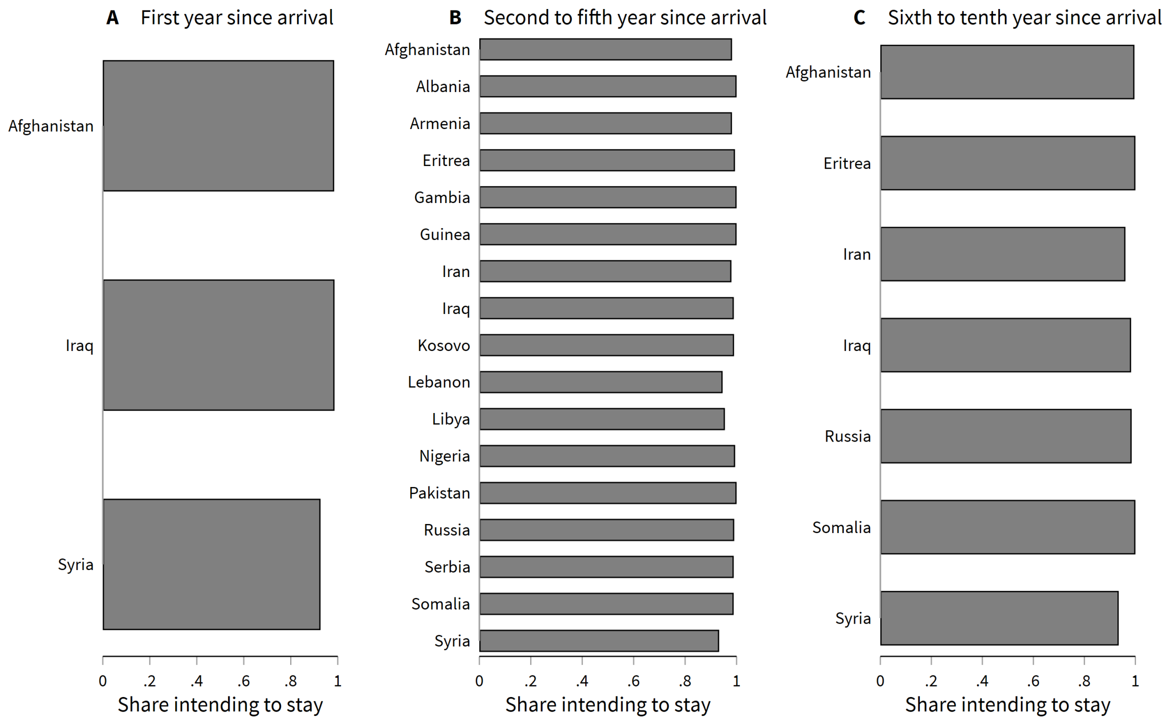

Notes: This Figure shows the share of respondents answering positively to the question: "Do you plan to stay in Germany?", by country and three spans of time since arrival, from the GSOEP 1984-2020. A) N = 1041; B) N = 15,835; C) N = 2,423. Only countries of birth with more than 50 respondents are shown.

FIGURE 3.E.2: Components of the first principal component of confidence in the government in Ukraine and across country groups between 2012 and 2023.

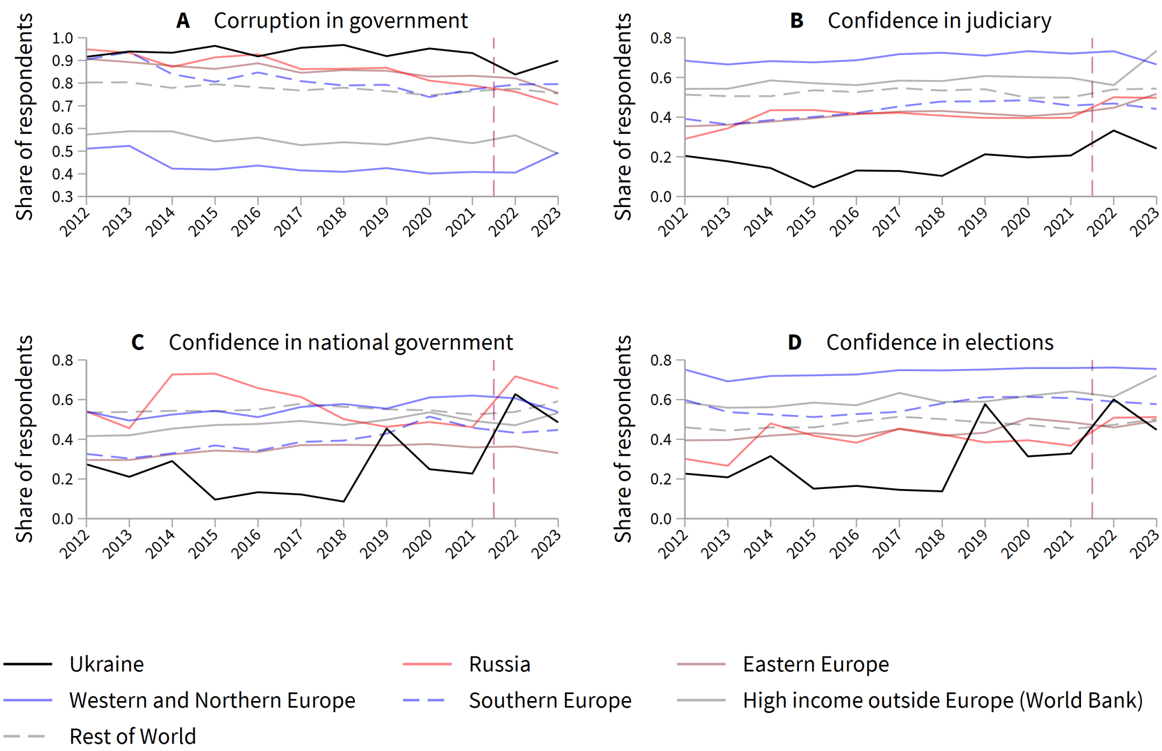

Notes: Changes in the four components of confidence in government. See notes to Figure 3.7.2 for construction of the sample.

FIGURE 3.E.3: Year-on-year changes in the desire to emigrate, optimism, and confidence in the government and military

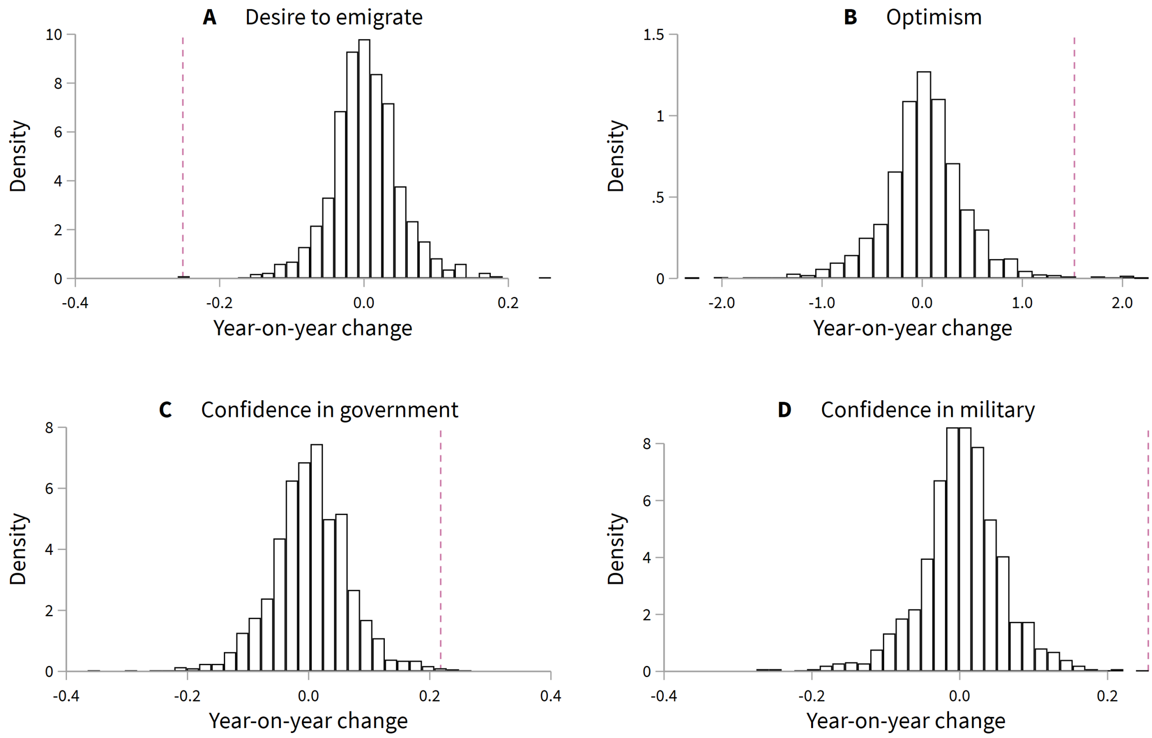

Notes: This figure is based on the sample of all country-years for which the country was surveyed by GWP as well as the year before and the respective question(s) are included, between 2006 and 2023. We omit observations with less than 500 valid observations in either the focal year or the year before. The dashed vertical line refers to the change in Ukraine between 2021 and 2022. A) N = 1,286 B) N = 1,620; C) N = 1,344; D) N = 1,401.

FIGURE 3.E.4: Desire to emigrate, optimism, confidence in government and military in Ukraine and country groups, controlling for demographic factors

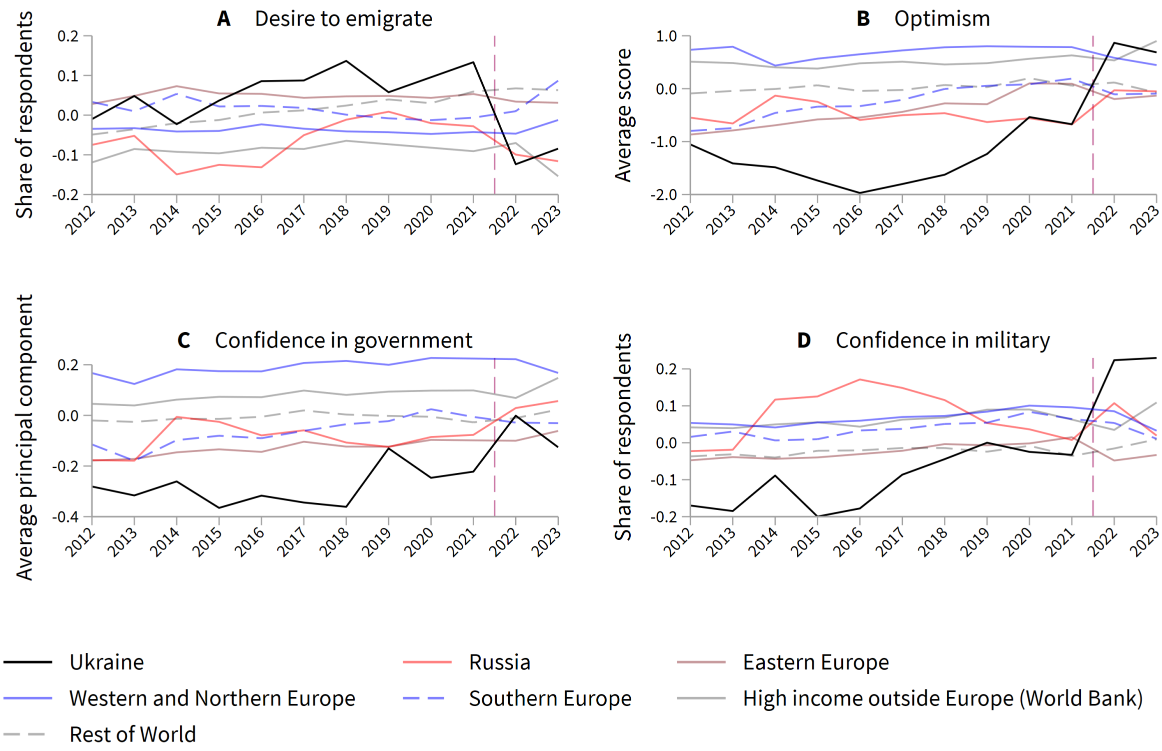

Notes: Desire to emigrate, optimism, confidence in government and military in Ukraine, Russia and five country groups, controlling for demographic factors. See notes to Figure 3.7.2 for information about the underlying sample and questions. Here, we show the residuals of each of the four outcomes after controlling for demographic factors. We obtain the residuals of the respective variables after regressing each of them on (1) age, age squared, dummies for secondary and tertiary educational attainment and (2) a binary indicator for being female, and all interactions of (1) with (2).

FIGURE 3.E.5: Desire of Ukrainians to live outside Ukraine under four scenarios.

Notes: This figure shows how Ukrainians' desire to emigrate would have changed from 2021 to 2022 under four different scenarios about refugees' desire to live in Ukraine, discussed in Section 3.A.3. We take the values for desire to emigrate for Ukrainians in Ukraine from GWP and use results from the full baseline Verian sample to evaluate the counterfactual desire to emigrate of Ukrainian refugees in case they would still be in Ukraine. The number of Ukrainians in Ukraine, in European countries covered by Verian and in Russia and Belarus are based on data from Eurostat from 31 August 2022. See text in Section 3.A.3 for a detailed discussion of the four cases. N = 1,024 (2020), N = 974 (2021), N = 991 (2022).

FIGURE 3.E.6: National pride has increased over time and skyrocketed in 2022

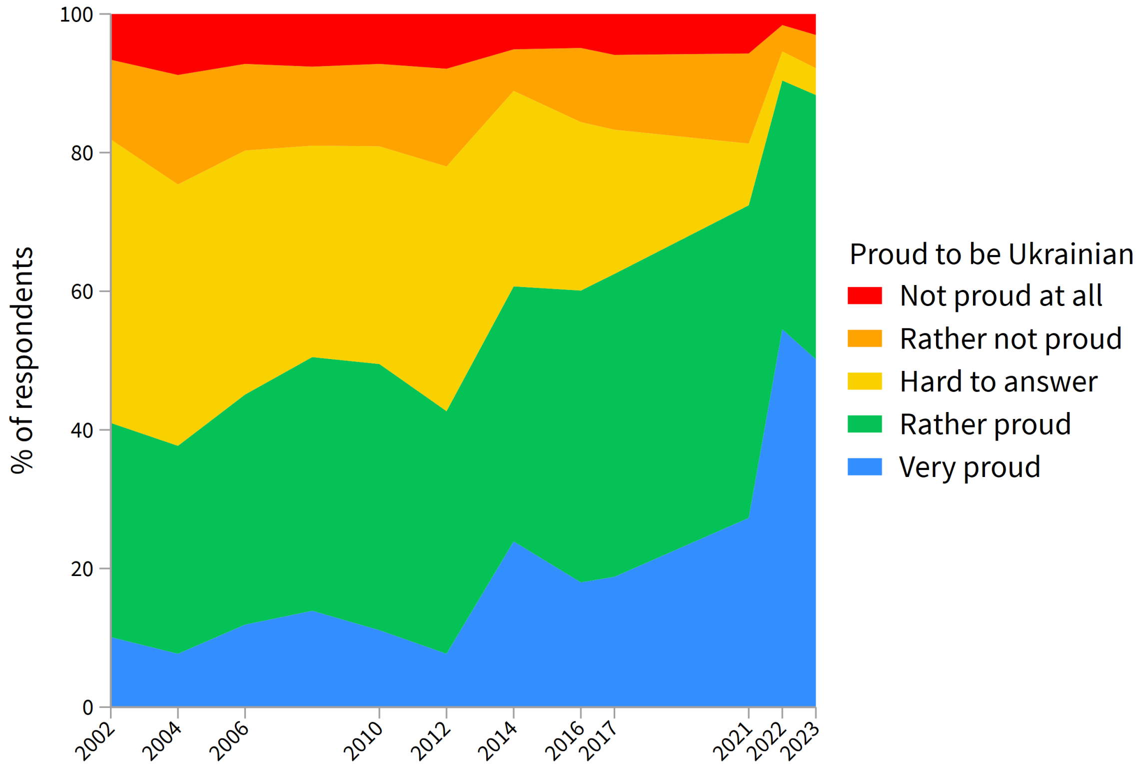

Notes: Proud to be Ukrainian between 2002 and 2023. Interviews were conducted in person by the Ilko Kucheriv Democratic Initiatives Foundation in cooperation with the Razumkov Centre (9). Interviews in 2022 and 2023 were conducted in August. In the Zaporizhzhya, Mykolaiv, Kharkiv regions the survey was conducted only in the territories controlled by the Ukrainian government and where there are no combat actions. All years in which surveys were conducted are indicated on the x-axis. About 2,000 respondents are interviewed in every year.

FIGURE 3.E.7: Most Ukrainian in Ukraine plan to build a future in Ukraine

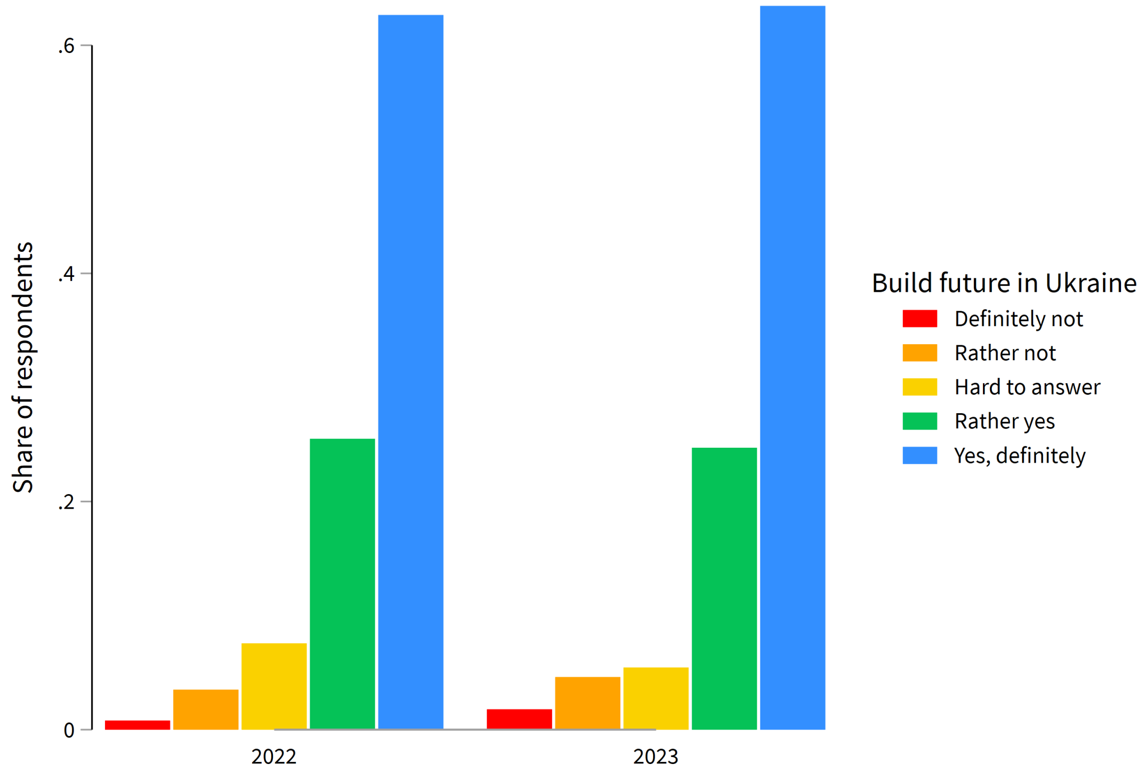

Notes: Distribution of plans to build a future in Ukraine. See the text and notes to Figure 3.E.6 about the survey, of which this figure only uses the 2022 and 2023 data on plans to build a future. N = 2,024 (2022) and N = 2,002 (2023).

FIGURE 3.E.8: Relation between national pride and plans to build a future in Ukraine

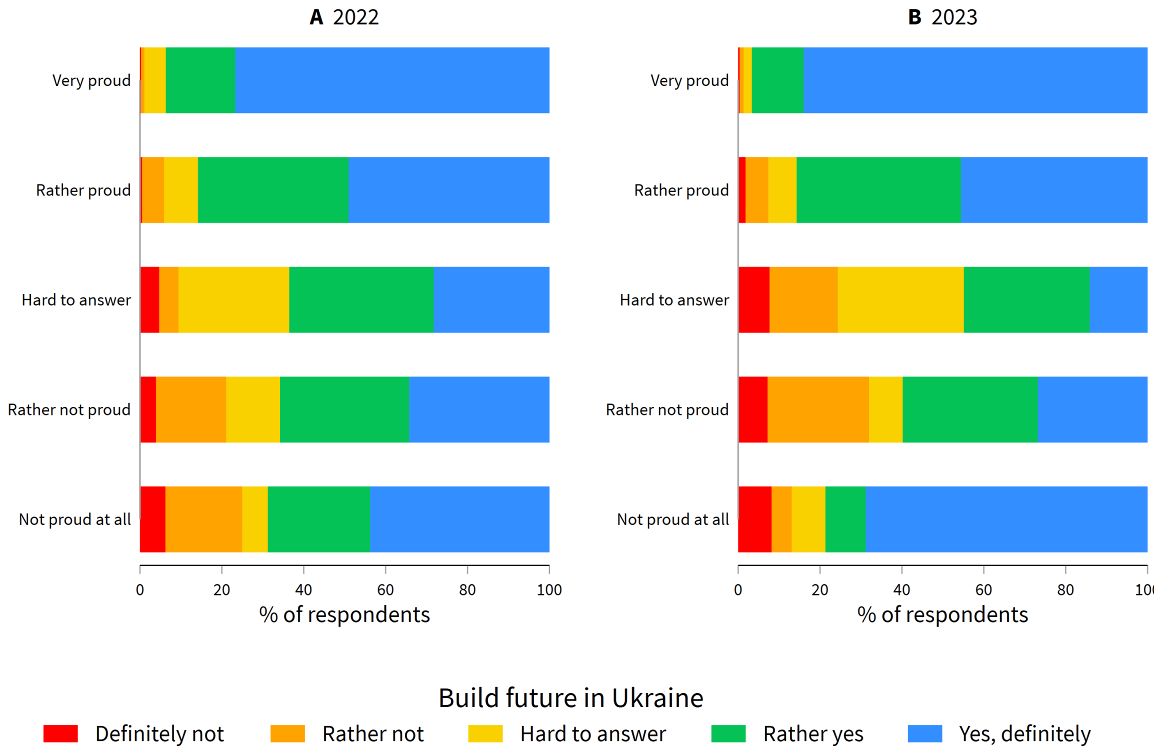

Notes: Plans to build a future in Ukraine in 2022 over levels of national pride. See notes to Figure 3.E.6 about the survey, of which this figure only uses the 2022 and 2023 data. N = 2,024 (2022) and N = 2,002 (2023).

Bibliography

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. 2023. "When should you adjust standard errors for clustering?" *The Quarterly Journal of Economics*, 138(1): 1–35.

Abarcar, Paolo, and Caroline Theoharides. 2024. "Medical worker migration and origin-country human capital: Evidence from U.S. visa policy." *The Review of Economics and Statistics*, 106(1): 20–35.

Abarca, Sergio F., Kristen L. Corbosiero, and Thomas J. Galarneau Jr. 2010. "An evaluation of the Worldwide Lightning Location Network (WWLLN) using the National Lightning Detection Network (NLDN) as ground truth." *Journal of Geophysical Research: Atmospheres*, 115(D18).

Abramenko, Serhii, Vasily Korovkin, and Alexey Makarin. 2024. "Social media as an identity barometer: Evidence from the Russia-Ukraine war." *AEA Papers & Proceedings*, 114: 70–74.

Adamchik, Vera A., Thomas J. Hyclak, Piotr Sedlak, and Larry W. Taylor. 2019. "Wage returns to English proficiency in Poland." *Journal of Labor Research*, 40: 276–295.

Adda, Jérôme, Christian Dustmann, and Joseph-Simon Görlach. 2022. "The dynamics of return migration, human capital accumulation, and wage assimilation." *The Review of Economic Studies*, 89(6): 2841–2871.

Adhikari, Prakash. 2013. "Conflict-induced displacement, understanding the causes of flight." *American Journal of Political Science*, 57(1): 82–89.

Adserà, Alícia, and Ana Ferrer. 2021. "Linguistic proximity and the labour market performance of immigrant men in Canada." *Labour*, 35(1): 1–23.

Adserà, Alícia, and Mariola Pytliková. 2015. "The role of language in shaping international migration." *The Economic Journal*, 125(586): F49–F81.

Adserà, Alícia, and Mariola Pytliková. 2016. "Language and migration." In *The Palgrave Handbook of Economics and Language*. 342–372. London:Palgrave Macmillan UK.

Agersnap, Ole, Amalie Jensen, and Henrik Kleven. 2020. "The welfare magnet hypothesis: Evidence from an immigrant welfare scheme in Denmark." *American Economic Review: Insights*, 2(4): 527–542.

Akbaritabar, Aliakbar, Tom Theile, and Emilio Zagheni. 2024. "Bilateral flows and rates of international migration of scholars for 210 countries for the period 1998-2020." *Scientific Data*, 11(1): 816.

Akbulut-Yuksel, Mevlude. 2014. "Children of war: The long-run effects of large-scale physical destruction and warfare on children." *Journal of Human Resources*, 49: 634–662.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015. "The skill complementarity of broadband internet." *The Quarterly Journal of Economics*, 130(4): 1781–1824.

Aksoy, Cevat Giray, Gaurav Khanna, Victoria Marino, and Semih Tumen. 2024. "Hometown conflict and refugees integration efforts." CEPR Discussion Paper 16862.

Aksoy, Cevat Giray, Panu Poutvaara, and Felicitas Schikora. 2023. "First time around: Local conditions and multi-dimensional integration of refugees." *Journal of Urban Economics*, 137: 103588.

Albert, Christoph, and Joan Monras. 2022. "Immigration and spatial equilibrium: The role of expenditures in the country of origin." *American Economic Review*, 112(11): 3763–3802.

Alesina, Alberto, and Marco Tabellini. 2024. "The political effects of immigration: Culture or economics?" *Journal of Economic Literature*, 62(1): 5–46.

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva. 2023. "Immigration and redistribution." *The Review of Economic Studies*, 90(1): 1–39.

Alesina, Alberto, Johann Harnoss, and Hillel Rapoport. 2016. "Birthplace diversity and economic prosperity." *Journal of Economic Growth*, 21: 101–138.

Al Husein, Nawras, and Natascha Wagner. 2023. "Determinants of intended return migration among refugees: A comparison of Syrian refugees in Germany and Turkey." *International Migration Review*, 57(4): 1771–1805.

Allcott, H., J.C. Castillo, M. Gentzkow, L. Musolff, and T. Salz. 2024. "Sources of Market Power in Web Search: Evidence from a Field Experiment."

Allport, Gordon W. 1954. *The Nature of Prejudice.* Perseus Books.

Alrababa'h, Ala, Daniel Masterson, Marine Casalis, Dominik Hangartner, and Jeremy Weinstein. 2023. "The dynamics of refugee return: Syrian refugees and their migration intentions." *British Journal of Political Science*, 1–24.

Altonji, Joseph G., and David Card. 1991. "The effects of immigration on the labor market outcomes of less-skilled natives." In *Immigration, Trade, and the Labor Market*. 137–170. Routledge.

Amanzadeh, Naser, Amir Kermani, and Timothy McQuade. 2024. "Return migration and human capital flows." NBER Working Paper 32352, National Bureau of Economic Research.

Anelli, Massimo, Gaetano Basso, Giuseppe Ippedico, and Giovanni Peri. 2023. "Emigration and entrepreneurial drain." *American Economic Journal: Applied Economics*, 15(2): 218–252.

Anger, Silke, Jacopo Bassetto, and Malte Sandner. 2022. "Making integration work? Facilitating access to occupational recognition and immigrants labor market performance." IZA Discussion Paper No. 15349.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion.* Princeton University Press.

Arendt, Jacob N., Christian Dustmann, and Hyejin Ku. 2023. "Permanent residency and refugee immigrants' skill investment." *Journal of Labor Economics*, forthcoming.

Arendt, Jacob N., Iben Bolvig, Mette Foged, Linea Hasager, and Giovanni Peri. 2022. "Integrating refugees: Language training or work-first incentives?" NBER Working Paper 26834.

Arnold, Benjamin F., Daniel R. Hogan, John M. Colford, and Alan E. Hubbard. 2011. "Simulation methods to estimate design power: An overview for applied research." *BMC Medical Research Methodology*, 11(1): 1–10.

Autor, David H. 2003. "Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing." *Journal of Labor Economics*, 21(1): 1–42.

Azam, Mehtabul, Aimee Chin, and Nishith Prakash. 2013. "The returns to English-language skills in India." *Economic Development and Cultural Change*, 61(2): 335–367.

Bahar, Dany, and Hillel Rapoport. 2018. "Migration, knowledge diffusion and the comparative advantage of nations." *The Economic Journal*, 128(612): F273–F305.

Bahar, Dany, Andreas Hauptmann, Cem Özgüzel, and Hillel Rapoport. 2022. "Migration and knowledge diffusion: The effect of returning refugees on export performance in the former Yugoslavia." *The Review of Economics and Statistics*, 104(2): 321–337.

Bahar, Dany, Rebecca J. Brough, and Giovanni Peri. 2024. "Forced migration and refugees: Policies for successful economic and social integration." NBER Working Paper 32266.

Bailey, Michael, Drew M Johnston, Martin Koenen, Theresa Kuchler, Dominic Russel, and Johannes Stroebel. 2022. "The social integration of international migrants: Evidence from the networks of Syrians in Germany." NBER Working Paper 29925.

Balcilar, Mehmet, and Jeffrey B. Nugent. 2019. "The migration of fear: An analysis of migration choices of Syrian refugees." *The Quarterly Review of Economics and Finance*, 73: 95–110.

Bansak, Kirk, Jeremy Ferwerda, Jens Hainmueller, Andrea Dillon, Dominik Hangartner, Duncan Lawrence, and Jeremy Weinstein. 2018. "Improving refugee integration through data-driven algorithmic assignment." *Science*, 359(6373): 325–329.

Barsbai, Toman, Andreas Steinmayr, and Christoph Winter. 2024. "Immigrating into a recession: Evidence from family migrants to the U.S." *Journal of Labor Economics*, forthcoming.

Barsbai, Toman, Hillel Rapoport, Andreas Steinmayr, and Christoph Trebesch. 2017. "The effect of labor migration on the diffusion of democracy: Evidence from a former Soviet Republic." *American Economic Journal: Applied Economics*, 9(3): 36–69.

Bassetto, Jacopo, and Teresa Freitas Monteiro. 2024. "Immigrants' returns intentions and job search behavior-when the home country is unsafe." CESifo Working Paper 10908.

Batalova, Jeanne. 2022. "Top statistics on global migration and migrants." https://www.migrationpolicy.org/article/top-statistics-global-migration-migrants (last accessed September 2, 2024).

Battisti, Michele, Giovanni Peri, and Agnese Romiti. 2022. "Dynamic effects of co-ethnic networks on immigrants' economic success." *The Economic Journal*, 132(641): 58–88.

Battisti, Michele, Yvonne Giesing, and Nadzeya Laurentsyeva. 2019. "Can job search assistance improve the labour market integration of refugees? Evidence from a field experiment." *Labour Economics*, 61: 101745.

Bazzi, Samuel. 2017. "Wealth heterogeneity and the income elasticity of migration." *American Economic Journal: Applied Economics*, 9(2): 219–255.

Beaman, Lori A. 2012. "Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the U.S." *The Review of Economic Studies*, 79(1): 128–161.

Beaman, Lori, Harun Onder, and Stefanie Onder. 2022. "When do refugees return home? Evidence from Syrian displacement in Mashreq." *Journal of Development Economics*, 155: 102802.

Becker, Sascha O., and Andreas Ferrara. 2019. "Consequences of forced migration: A survey of recent findings." *Labour Economics*, 59: 1–16.

Becker, Sascha O., Irena Grosfeld, Pauline Grosjean, Nico Voigtländer, and Ekaterina Zhuravskaya. 2020. "Forced migration and human capital: Evidence from post-WWII population transfers." *American Economic Review*, 110(5): 1430–63.

Beine, Michel, Frédéric Docquier, and Çağlar Özden. 2011. "Diasporas." *Journal of Development Economics*, 95(1): 30–41.

Beine, Michel, Frédéric Docquier, and Hillel Rapoport. 2001. "Brain drain and economic growth: Theory and evidence." *Journal of Development Economics*, 64(1): 275–289.

Beine, Michel, Frederic Docquier, and Hillel Rapoport. 2008. "Brain drain and human capital formation in developing countries: Winners and losers." *The Economic Journal*, 118(528): 631–652.

Beine, Michel, Michel Bierlaire, and Frédéric Docquier. 2021. "New York, Abu Dhabi, London or stay at home? Using a cross-nested logit model to identify complex substitution patterns in Migration." *Using a Cross-Nested Logit Model to Identify Complex Substitution Patterns in Migration*.

Beine, Michel, Simone Bertoli, and Jesús Fernández-Huertas Moraga. 2016. "A practitioners' guide to gravity models of international migration." *The World Economy*, 39(4): 496–512.

Bellodi, Luca, Frédéric Docquier, Stefano Iandolo, Massimo Morelli, and Riccardo Turati. 2024. "Digging up trenches: Populism, selective mobility, and the political polarization of Italian municipalities." IZA Discussion Paper 16732.

Belot, Michèle, and Sjef Ederveen. 2012. "Cultural barriers in migration between OECD countries." *Journal of Population Economics*, 25(3): 1077–1105.

Benhabib, Jess, and Boyan Jovanovic. 2012. "Optimal migration: A world perspective." *International Economic Review*, 53(2): 321–348.

Bertoli, Simone, and Jesús Fernández-Huertas Moraga. 2013. "Multilateral resistance to migration." *Journal of Development Economics*, 102: 79–100.

Bertoli, Simone, Frédéric Docquier, Hillel Rapoport, and Ilse Ruyssen. 2022. "Weather shocks and migration intentions in Western Africa: Insights from a multilevel analysis." *Journal of Economic Geography*, 22(2): 289–323.

Bertoli, Simone, Jesús Fernández-Huertas Moraga, and Lucas Guichard. 2020. "Rational inattention and migration decisions." *Journal of International Economics*, 126: 103364.

Besley, Timothy, and Robin Burgess. 2004. "Can labor regulation hinder economic performance? Evidence from India." *The Quarterly Journal of Economics*, 119(1): 91–134.

Bleakley, Hoyt, and Aimee Chin. 2004. "Language skills and earnings: Evidence from childhood immigrants." *The Review of Economics and Statistics*, 86(2): 481–496.

Bleakley, Hoyt, and Aimee Chin. 2010. "Age at arrival, English proficiency, and social assimilation among US immigrants." *American Economic Journal: Applied Economics*, 2(1): 165–192.

Böhme, Marcus H., André Gröger, and Tobias Stöhr. 2020. "Searching for a better life: Predicting international migration with online search keywords." *Journal of Development Economics*, 142: 102347.

Bohra-Mishra, Pratikshya, and Douglas Massey. 2011. "Individual decisions to migrate during civil conflict." *Demography*, 48(2): 401–424.

Bonadio, Barthélémy. 2023. "Migrants, trade and market access." *The Review of Economics and Statistics*, 1–45.

Borjas, George J. 1985. "Assimilation, changes in cohort quality, and the earnings of immigrants." *Journal of labor Economics*, 3(4): 463–489.

Borjas, George J. 1987. "Self-selection and the earnings of immigrants." *The American Economic Review*, 77(4): 531–553.

Borjas, George J. 1999. "Immigration and welfare magnets." *Journal of Labor Economics*, 17(4): 607–637.

Borjas, George J. 2003. "The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market." *The Quarterly Journal of Economics*, 118(4): 1335–1374.

Borjas, George J. 2015. "The slowdown in the economic assimilation of immigrants: Aging and cohort effects revisited again." *Journal of Human Capital*, 9(4): 483–517.

Borjas, George J. 2017. "The wage impact of the Marielitos: A reappraisal." *Industrial and Labor Relations Review*, 70(5): 1077–1110.

Borjas, George J., and Bernt Bratsberg. 1996. "Who leaves? The outmigration of the foreign-born." *The Review of Economics and Statistics*, 78(1): 165–176.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting event-study designs: Robust and efficient estimation." *Review of Economic Studies*, forthcoming.

Brade, Raphael, Oliver Himmler, Robert Jaeckle, and Philipp Weinschenk. 2024. "Helping Students to Succeed–The Long-Term Effects of Soft Commitments and Reminders." CESifo Working Paper 11001.

Bratsberg, Bernt, Oddbjørn Raaum, and Knut Røed. 2017. "Immigrant labor market integration across admission classes." *Nordic Economic Policy Review*, 15–44.

Brell, Courtney, Christian Dustmann, and Ian Preston. 2020. "The labor market integration of refugee migrants in high-income countries." *Journal of Economic Perspectives*, 34(1): 94–121.

Burchardi, Konrad B., Thomas Chaney, and Tarek A. Hassan. 2019. "Migrants, ancestors, and foreign investments." *The Review of Economic Studies*, 86(4): 1448–1486.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with a continuous treatment." *arXiv preprint arXiv:2107.02637*.

Camarena, Kara Ross, and Nils Hägerdal. 2020. "When do displaced persons return? Postwar migration among Christians in Mount Lebanon." *American Journal of Political Science*, 64: 223–239.

Cameron, A Colin, and Douglas L Miller. 2014. "Robust inference for dyadic data." *Unpublished manuscript, University of California-Davis*.

Campante, Filipe, Ruben Durante, and Francesco Sobbrio. 2018. "Politics 2.0: The multifaceted effect of broadband internet on political participation." *Journal of the European Economic Association*, 16(4): 1094–1136.

Card, David. 1990. "The impact of the Mariel boatlift on the Miami labor market." *Industrial and Labor Relations Review*, 43(2): 245–257.

Card, David. 2001. "Immigrant inflows, native outflows, and the local labor market impacts of higher immigration." *Journal of Labor Economics*, 19(1): 22–64.

Caruso, German, Christian Gomez Canon, and Valerie Mueller. 2021. "Spillover effects of the Venezuelan crisis: Migration impacts in Colombia." *Oxford Economic Papers*, 73(2): 771–795.

Chamarbagwala, Rubiana, and Hilcías E. Morán. 2011. "The human capital consequences of civil war: Evidence from Guatemala." *Journal of Development Economics*, 94: 41–61.

Chauvet, Lisa, and Marion Mercier. 2014. "Do return migrants transfer political norms to their origin country? Evidence from Mali." *Journal of Comparative Economics*, 42: 630–651.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with zeros? Some problems and solutions." *The Quarterly Journal of Economics*, forthcoming.

Chiswick, Barry R., and Paul W. Miller. 1994. "The determinants of post-immigration investments in education." *Economics of Education Review*, 13(2): 163–177.

Chiswick, Barry R., and Paul W. Miller. 1995. "The endogeneity between language and earnings: International analyses." *Journal of Labor Economics*, 13(2): 246–288.

Chiswick, Barry R., and Paul W. Miller. 2015. "International migration and the economics of language." In *Handbook of the Economics of International Migration*. Vol. 1, 211–269. Elsevier.

Choudhury, Prithwiraj. 2016. "Return migration and geography of innovation in MNEs: A natural experiment of on-the-job learning of knowledge production by local workers reporting to return migrants." *Journal of Economic Geography*, 16: 585–610.

Clemens, Michael A. 2011. "Economics and emigration: Trillion-dollar bills on the sidewalk?" *Journal of Economic Perspectives*, 25(3): 83–106.

Clemens, Michael A. 2014. "Does development reduce migration?" In *International Handbook on Migration and Economic Development*. Edward Elgar Publishing.

Clemens, Michael A. 2020. "The emigration life cycle: How development shapes emigration from poor countries." IZA Discussion Paper 13614.

Clemens, Michael A., and Jennifer Hunt. 2019. "The labor market effects of refugee waves: Reconciling conflicting results." *Industrial and Labor Relations Review*, 72(4): 818–857.

Colas, Mark, and Dominik Sachs. 2024. "The indirect fiscal benefits of low-skilled immigration." *American Economic Journal: Economic Policy*, 16(2): 515–550.

Collins Bartholomew. 2018. "Mobile Coverage Explorer." https://www.collinsbartholomew.com/mobile-coverage-maps/mobile-coverage-explorer/ (last accessed September 13, 2024).

Conte, Marc N., Pierre Cotterlaz, and Thierry Mayer. 2022. "The CEPII Gravity Database." https://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=8 (last accessed March 2, 2024).

Cooke, Thomas J., and Ian Shuttleworth. 2018. "The effects of information and communication technologies on residential mobility and migration." *Population, Space and Place*, 24(3): e2111.

Cortes, Kalena E. 2004. "Are refugees different from economic immigrants? Some empirical evidence on the heterogeneity of immigrant groups in the United States." *The Review of Economics and Statistics*, 86: 465–480.

Dahlberg, Matz, Johan Egebark, Ulrika Vikman, and Gülay Özcan. 2024. "Labor market integration of refugees: RCT evidence from an early intervention program in Sweden." *Journal of Economic Behavior & Organization*, 217: 614–630.

Dahlberg, Matz, Karin Edmark, and Heléne Lundqvist. 2012. "Ethnic diversity and preferences for redistribution." *Journal of Political Economy*, 120(1): 41–76.

Damm, Anna P. 2009. "Ethnic enclaves and immigrant labor market outcomes: Quasi-experimental evidence." *Journal of Labor Economics*, 27(2): 281–314.

Dao, Thu Hien, Frédéric Docquier, Chris Parsons, and Giovanni Peri. 2018. "Migration and development: Dissecting the anatomy of the mobility transition." *Journal of Development Economics*, 132: 88–101.

Datareportal. 2022. "Digital 2022: Ukraine." https://datareportal.com/reports/digital-2022-ukraine (last accessed Sep 13, 2024).

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2020. "Difference-in-Differences Estimators of Intertemporal Treatment Effects." *arXiv preprint arXiv:2007.04267*.

de Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review*, 110(9): 2964–2996.

de Chaisemartin, Clément, Xavier d'Haultfoeuille, Félix Pasquier, and Gonzalo Vazquez-Bare. 2024. "Difference-in-differences estimators for treatments continuously distributed at every period." arXiv.org Papers 2201.06898.

Dekker, Rianne, and Godfried Engbersen. 2014. "How social media transform migrant networks and facilitate migration." *Global Networks*, 14(4): 401–418.

Demurger, Sylvie, and Hui Xu. 2011. "Return migrants: The rise of new entrepreneurs in rural China." *World Development*, 39: 1847–1861.

Dinesen, Peter Thisted, and Mads Meier Jæger. 2013. "The effect of terror on institutional trust: New evidence from the 3/11 Madrid terrorist attack." *Political Psychology*, 34(6): 917–926.

Di Paolo, Antonio, and Aysit Tansel. 2015. "Returns to foreign language skills in a developing country: The case of Turkey." *The Journal of Development Studies*, 51(4): 407–421.

Di Paolo, Antonio, and Bernat Mallén. 2023. "Does geographical exposure to language learning centres affect language skills and labour market outcomes in a bilingual city?" *Economic Analysis and Policy*, 80: 429–449.

Docquier, Frédéric, and Hillel Rapoport. 2012. "Globalization, brain drain, and development." *Journal of Economic Literature*, 50(3): 681–730.

Docquier, Frédéric, Çağlar Özden, and Giovanni Peri. 2014. "The labour market effects of immigration and emigration in OECD countries." *The Economic Journal*, 124(579): 1106–1145.

Docquier, Frédéric, Elisabetta Lodigiani, Hillel Rapoport, and Maurice Schiff. 2016. "Emigration and democracy." *Journal of Development Economics*, 120: 209–223.

Docquier, Frédéric, Giovanni Peri, and Ilse Ruyssen. 2018. "The cross-country determinants of potential and actual migration." *International Migration Review*, 48(1): 37–99.

Dustmann, Christian. 1999. "Temporary migration, human capital, and language fluency of migrants." *Scandinavian Journal of Economics*, 101(2): 297–314.

Dustmann, Christian. 2003. "Return migration, wage differentials, and the optimal migration duration." *European Economic Review*, 47(2): 353–369.

Dustmann, Christian, and Anna Okatenko. 2014. "Out-migration, wealth constraints, and the quality of local amenities." *Journal of Development Economics*, 110: 52–63.

Dustmann, Christian, and Francesca Fabbri. 2003. "Language proficiency and labour market performance of immigrants in the UK." *The Economic Journal*, 113(489): 695–717.

Dustmann, Christian, Francesco Fasani, Xin Meng, and Luigi Minale. 2023. "Risk attitudes and household migration decisions." *Journal of Human Resources*, 58(1): 112–145.

Dustmann, Christian, Itzhak Fadlon, and Yoram Weiss. 2011. "Return migration, human capital accumulation and the brain drain." *Journal of Development Economics*, 95(1): 58–67.

Dustmann, Christian, Kristine Vasiljeva, and Anna Piil Damm. 2019. "Refugee migration and electoral outcomes." *The Review of Economic Studies*, 86(5): 2035–2091.

Dustmann, Christian, Rasmus Landersø, and Lars Højsgaard Andersen. 2024*a*. "Refugee benefit cuts." *American Economic Journal: Economic Policy*, 16(2): 406–441.

Dustmann, Christian, Rasmus Landersø, and Lars Højsgaard Andersen. 2024*b*. "Unintended consequences of welfare cuts on children and adolescents." *American Economic Journal: Applied Economics*, forthcoming.

Dustmann, Christian, Uta Schönberg, and Jan Stuhler. 2017. "Labor supply shocks, native wages, and the adjustment of local employment." *The Quarterly Journal of Economics*, 132(1): 435–483.

Edin, Per-Anders, Peter Fredriksson, and Olof Åslund. 2003. "Ethnic enclaves and the economic success of immigrants—evidence from a natural experiment." *The Quarterly Journal of Economics*, 118(1): 329–357.

Edo, Anthony. 2019. "The impact of immigration on the labor market." *Journal of Economic Surveys*, 33(3): 922–948.

Edo, Anthony, Yvonne Giesing, Jonathan Öztunc, and Panu Poutvaara. 2019. "Immigration and electoral support for the far-left and the far-right." *European Economic Review*, 115: 99–143.

Ek, Andreas. 2024. "Cultural values and productivity." *Journal of Political Economy*, 132(1): 295–335.

Ersoy, Fulya. 2021. "Returns to effort: experimental evidence from an online language platform." *Experimental Economics*, 24(3): 1047–1073.

Eurostat. 2023. "Decisions granting temporary protection by citizenship, age and sex - annual data ." `https://ec.europa.eu/eurostat/databrowser/view/MIGR_ASYTPFA/default/table?lang=en&category=migr.migr_asy.migr_asytp` (last accessed Apr 28, 2023).

Fackler, Thomas A., Yvonne Giesing, and Nadzeya Laurentsyeva. 2020. "Knowledge remittances: Does emigration foster innovation?" *Research Policy*, 49(9): 103863.

Falck, Oliver, Robert Gold, and Stephan Heblich. 2014. "E-lections: Voting behavior and the Internet." *American Economic Review*, 104(7): 2238–2265.

Fasani, Francesco, Tommaso Frattini, and Luigi Minale. 2021. "Lift the ban? Initial employment restrictions and refugee labour market outcomes." *Journal of the European Economic Association*, 19(5): 2803–2854.

Fasani, Francesco, Tommaso Frattini, and Luigi Minale. 2022. "(The struggle for) Refugee integration into the labour market: Evidence from Europe." *Journal of Economic Geography*, 22(2): 351–393.

Fasani, Francesco, Tommaso Frattini, and Maxime Pirot. 2023. "From refugees to citizens: Labor market returns to naturalization." IZA Discussion Paper 16651.

Feigenberg, Benjamin, Ben Ost, and Javaeria A. Qureshi. 2023. "Omitted Variable Bias in Interacted Models: A Cautionary Tale." *The Review of Economics and Statistics*, 1–47.

Finney, Declan L., Ruth M. Doherty, Oliver Wild, David S. Stevenson, Ian A. MacKenzie, and Alan M. Blyth. 2018. "A projected decrease in lightning under climate change." *Nature Climate Change*, 8(3): 210–213.

Foged, Mette, and Cynthia Van der Werf. 2023. "Access to language training and the local integration of refugees." *Labour Economics*, 102366.

Foged, Mette, and Giovanni Peri. 2016. "Immigrants' effect on native workers: New analysis on longitudinal data." *American Economic Journal: Applied Economics*, 8(2): 1–34.

Foged, Mette, Linea Hasager, and Giovanni Peri. 2024. "Comparing the effects of policies for the labor market integration of refugees." *Journal of Labor Economics*, 42(S1): S335–S377.

Freeman, Cassie, Audrey Kittredge, Hope Wilson, and Bozena Pajak. 2023. "The Duolingo method for app-based teaching and learning." https://duolingo-papers.s3.amazonaws.com/reports/Duolingo_whitepaper_duolingo_method_2023.pdf (last accessed Sep 13, 2024).

Funke, Manuel, Moritz Schularick, and Christoph Trebesch. 2023. "Populist leaders and the economy." *American Economic Review*, 113(12): 3249–3288.

Gehring, Kai. 2021. "Can external threats foster a European Union identity? Evidence from Russia's invasion of Ukraine." *The Economic Journal*, 132: 1489–1516.

Ginsburgh, Victor A., and Juan Prieto-Rodriguez. 2011. "Returns to foreign languages of native workers in the European Union." *Industrial and Labor Relations Review*, 64(3): 599–618.

Ginsburgh, Victor, Jacques Melitz, and Farid Toubal. 2017. "Foreign language learning and trade." *Review of International Economics*, 25(2): 320–361.

Goethe Institute. 2024*a*. "German Courses in Bangladesh: Dates and prices." https://www.goethe.de/ins/bd/en/spr/kur/tup.cfm (last accessed Sep 13, 2024).

Goethe Institute. 2024*b*. "German Courses in Colombia: Dates and prices." https://www.goethe.de/ins/co/de/spr/kur.html (last accessed Sep 13, 2024).

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277.

Görlach, Joseph-Simon. 2023. "Borrowing constraints and the dynamics of return and repeat migration." *Journal of Labor Economics*, 41(1): 205–243.

Gould, David M. 1994. "Immigrant links to the home country: Empirical implications for U.S. bilateral trade flows." *The Review of Economics and Statistics*, 302–316.

Grogger, Jeffrey, and Gordon H. Hanson. 2011. "Income maximization and the selection and sorting of international migrants." *Journal of Development Economics*, 95(1): 42–57.

Grubanov-Boskovic, Sara, Sona Kalantaryan, Silvia Migali, and Marco Scipioni. 2021. "The impact of the internet on migration aspirations and intentions." *Migration Studies*, 9(4): 1807–1822.

GSMA. 2023. "Mobile internet connectivity 2023." https://www.gsma.com/r/wp-content/uploads/2023/10/The-State-of-Mobile-Internet-Connectivity-Report-2023.pdf?utm_source=website&utm_medium=button&utm_campaign=somic23 (last accessed Sep 13, 2024).

Guichard, Lucas, and Joël Machado. 2024. "The externalities of immigration policies on migration flows: the case of an asylum policy." *Journal of Economic Geography*, lbae016.

Guriev, Sergei, Nikita Melnikov, and Ekaterina Zhuravskaya. 2021. "3G internet and confidence in government." *The Quarterly Journal of Economics*, 136: 2533–2613.

Hahm, Sabrina, and Michele Gazzola. 2022. "The value of foreign language skills in the German labor market." *Labour Economics*, 76: 102150.

Halla, Martin, Alexander F. Wagner, and Josef Zweimüller. 2017. "Immigration and voting for the far right." *Journal of the European Economic Association*, 15(6): 1341–1385.

Hangartner, Dominik, Elias Dinas, Moritz Marbach, Konstantinos Matakos, and Dimitrios Xefteris. 2019. "Does exposure to the refugee crisis make natives more hostile?" *American Political Science Review*, 113(2): 442–455.

Harris, John R., and Michael P. Todaro. 1970. "Migration, unemployment and development: A two-sector analysis." *The American Economic Review*, 60(1): 126–142.

Heller, Blake H., and Kirsten Slungaard Mumma. 2023. "Immigrant integration in the United States: The role of adult English language training." *American Economic Journal: Economic Policy*, 15(3): 407–437.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil. 2012. "Measuring economic growth from outer space." *American Economic Review*, 102(2): 994–1028.

Hendricks, Lutz, and Todd Schoellman. 2018. "Human capital and development accounting: New evidence from wage gains at migration." *The Quarterly Journal of Economics*, 133(2): 665–700.

Hjort, Jonas, and Jonas Poulsen. 2019. "The arrival of fast internet and employment in Africa." *American Economic Review*, 109(3): 1032–1079.

Huber, Matthias, and Silke Uebelmesser. 2023. "Presence of language-learning opportunities and migration." *Labour Economics*, 84: 102409.

Huber, Matthias, Ann-Marie Sommerfeld, and Silke Uebelmesser. 2022. "Language learning: human capital investment or consumption?" *Empirica*, 49(4): 897–948.

Hunt, Jennifer. 2017. "The impact of immigration on the educational attainment of natives." *Journal of Human Resources*, 52(4): 1060–1118.

Institute for the Study of War and AEI's Critical Threats Project. 2023. "Russian offensive campaign assessment."

Isphording, Ingo E. 2013. "Returns to foreign language skills of immigrants in Spain." *Labour*, 27(4): 443–461.

Isphording, Ingo E., and Sebastian Otten. 2011. "Linguistic distance and the language fluency of immigrants." Ruhr Economic Papers Working Paper 274.

ITU. 2021. "Measuring digital development: Facts and figures 2021." `https://www.itu.int/en/ITU-D/Statistics/Documents/facts/FactsFigures2021.pdf` (last accessed Sep 13, 2024).

Jaeger, David A., Thomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, and Holger Bonin. 2010. "Direct evidence on risk attitudes and migration." *The Review of Economics and Statistics*, 92(3): 684–689.

Jaschke, Philipp, and Sekou Keita. 2021. "Say it like Goethe: Language learning facilities abroad and the self-selection of immigrants." *Journal of Development Economics*, 149: 102597.

Jiang, Xiangying, Haoyu Chen, Lucy Portnoff, Erin Gustafson, Joseph Rollinson, Luke Plonsky, and Bozena Pajak. 2021*a*. "Seven units of Duolingo courses comparable to 5 university semesters in reading and listening." *Duolingo Research Report DRR-21-03*.

Jiang, Xiangying, Joseph Rollinson, Luke Plonsky, Erin Gustafson, and Bozena Pajak. 2021*b*. "Evaluating the reading and listening outcomes of beginning-level Duolingo courses." *Foreign Language Annals*, 54(4): 974–1002.

Johnson, Paul, and Chris Papageorgiou. 2020. "What remains of cross-country convergence?" *Journal of Economic Literature*, 58(1): 129–175.

Karadja, Mounir, and Erik Prawitz. 2019. "Exit, voice, and political change: Evidence from Swedish mass migration to the United States." *Journal of Political Economy*, 127(4): 1864–1925.

Kerr, William R. 2008. "Ethnic scientific communities and international technology diffusion." *The Review of Economics and Statistics*, 90(3): 518–537.

Krasniqi, Besnik A., and Nick Williams. 2019. "Migration and intention to return: Entrepreneurial intentions of the diaspora in post-conflict economies." *Post-Communist Economies*, 31: 464–483.

Krugman, Paul. 1980. "Scale economies, product differentiation, and the pattern of trade." *The American Economic Review*, 70(5): 950–959.

Kulyk, Volodymyr. 2016. "National identity in Ukraine: Impact of Euromaidan and the war." *Europe-Asia Studies*, 68: 588–608.

León, Gianmarco. 2012. "Civil conflict and human capital accumulation: The long-term effects of political violence in Peru." *Journal of Human Resources*, 47: 991–1022.

Liwiński, Jacek. 2019. "The wage premium from foreign language skills." *Empirica*, 46(4): 691–711.

Lochmann, Alexia, Hillel Rapoport, and Biagio Speciale. 2019. "The effect of language training on immigrants' economic integration: Empirical evidence from France." *European Economic Review*, 113: 265–296.

Manacorda, Marco, Alan Manning, and Jonathan Wadsworth. 2012. "The impact of immigration on the structure of wages: Theory and evidence from Britain." *Journal of the European Economic Association*, 10(1): 120–151.

Manacorda, Marco, and Andrea Tesei. 2020. "Liberation technology: Mobile phones and political mobilization in Africa." *Econometrica*, 88(2): 533–567.

Manchin, Miriam, and Sultan Orazbayev. 2018. "Social networks and the intention to migrate." *World Development*, 109: 360–374.

Massey, Douglas S., and Emilio A. Parrado. 1998. "International migration and business formation in Mexico." *Social Science Quarterly*, 79: 1–20.

Massey, Douglas S., Joaquin Arango, Graeme Hugo, Ali Kouaouci, Adela Pellegrino, and J. Edward Taylor. 1993. "Theories of international migration: A review and appraisal." *Population and Development Review*, 19(3): 431–466.

Ma, Wenyue, and Paula Winke. 2019. "Self-assessment: How reliable is it in assessing oral proficiency over time?" *Foreign Language Annals*, 52(1): 66–86.

Mayda, Anna-Maria. 2010. “International migration: A panel data analysis of the determinants of bilateral flows.” *Journal of Population Economics*, 23: 1249–1274.

Mayda, Anna-Maria, Christopher Parsons, Han Pham, and Pierre-Louis Vézina. 2022. “Refugees and foreign direct investment: Quasi-experimental evidence from US resettlements.” *Journal of Development Economics*, 156: 102818.

Mayda, Anna-Maria, Giovanni Peri, and Walter Steingress. 2022. “The political impact of immigration: Evidence from the United States.” *American Economic Journal: Applied Economics*, 14(1): 358–389.

Mayer, Thierry, and Soledad Zignago. 2011. “Notes on CEPII’s distances measures: The GeoDist database.” CEPII Working Paper 2011-25.

McFadden, Daniel. 1978. “Modelling the choice of residential location.” In *Spatial Interaction Theory and Residential Location by F. Snickars and J. Weibull*. 75–96. North Holland, Amsterdam.

McKenzie, David, and Hillel Rapoport. 2010. “Self-selection patterns in Mexico-US migration: The role of migration networks.” *The Review of Economics and Statistics*, 92(4): 811–821.

McKenzie, David, John Gibson, and Steven Stillman. 2013. “A land of milk and honey with streets paved with gold: Do emigrants have over-optimistic expectations about incomes abroad?” *Journal of Development Economics*, 102: 116–127.

Melitz, Jacques, and Farid Toubal. 2014. “Native language, spoken language, translation and trade.” *Journal of International Economics*, 93(2): 351–363.

Migali, Silvia, and Marco Scipioni. 2019. “Who’s About to Leave? A Global Survey of Aspirations and Intentions to Migrate.” *International Migration*, 57(5): 181–200.

Mincer, Jacob. 1978. “Family migration decisions.” *Journal of Political Economy*, 86(5): 749–773.

Moriconi, Simone, Giovanni Peri, and Riccardo Turati. 2019. “Immigration and voting for redistribution: Evidence from European elections.” *Labour Economics*, 61: 101765.

Mueller, John E. 1970. “Presidential popularity from Truman to Johnson.” *American Political Science Review*, 64: 18–34.

Munk, Martin D., Till Nikolka, and Panu Poutvaara. 2022. “International family migration and the dual-earner model.” *Journal of Economic Geography*, 22(2): 263–287.

Nagengast, Arne, and Yoto V Yotov. 2023. “Staggered difference-in-differences in gravity settings: Revisiting the effects of trade agreements.” CESifo Working Paper No. 10782.

Nocito, Samuel. 2021. “The effect of a university degree in english on international labor mobility.” *Labour Economics*, 68: 101943.

Ortega, Francesc, and Giovanni Peri. 2013. “The effect of income and immigration policies on international migration.” *Migration Studies*, 1(1): 47–74.

Oster, Emily. 2019. “Unobservable selection and coefficient stability: Theory and evidence.” *Journal of Business & Economic Statistics*, 37(2): 187–204.

Ottaviano, Gianmarco I. P., and Giovanni Peri. 2012. “Rethinking the effect of immigration on wages.” *Journal of the European Economic Association*, 10(1): 152–197.

Ottaviano, Gianmarco I. P., Giovanni Peri, and Greg C. Wright. 2013. “Immigration, offshoring, and American jobs.” *American Economic Review*, 103(5): 1925–1959.

Parey, Matthias, and Fabian Waldinger. 2011. “Studying abroad and the effect on international labour market mobility: Evidence from the introduction of ERASMUS." *The Economic Journal*, 121(551): 194–222.

Parsons, Christopher, and Pierre-Louis Vézina. 2018. "Migrant networks and trade: The Vietnamese boat people as a natural experiment." *The Economic Journal*, 128(612): F210–F234.

Peri, Giovanni. 2012. "The effect of immigration on productivity: Evidence from US states." *The Review of Economics and Statistics*, 94(1): 348–358.

Peri, Giovanni, and Chad Sparber. 2009. "Task specialization, immigration, and wages." *American Economic Journal: Applied Economics*, 1(3): 135–169.

Peri, Giovanni, and Vasil Yasenov. 2019. "The labor market effects of a refugee wave: Synthetic control method meets the Mariel boatlift." *Journal of Human Resources*, 54(2): 267–309.

Pesando, Luca Maria, Valentina Rotondi, Manuela Stranges, Ridhi Kashyap, and Francesco C. Billari. 2021. "The internetization of international migration." *Population and Development Review*, 47(1): 79–111.

Pischke, Steve. 2007. "Lecture notes on measurement error." *London School of Economics, London*.

Porcher, Charly. 2020. "Migration with Costly Information." Princeton University Unpublished working paper.

Prato, Marta. 2022. "The global race for talent: Brain drain, knowledge transfer, and growth." *Knowledge Transfer, and Growth (November 27, 2022)*.

Prentice, Sydney A., and David Mackerras. 1977. "The ratio of cloud to cloud-ground lightning flashes in thunderstorms." *Journal of Applied Meteorology (1962-1982)*, 16(5): 545–550.

Priem, Jason, Heather Piwowar, and Richard Orr. 2022. "OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts." *arXiv preprint arXiv:2205.01833*.

Rachels, Jason R, and Amanda J Rockinson-Szapkiw. 2018. "The effects of a mobile gamification app on elementary students' Spanish achievement and self-efficacy." *Computer Assisted Language Learning*, 31(1-2): 72–89.

Raleigh, Clionadh, Andrew Linke, Håvard Hegre, and Joakim Karlsen. 2010. "Introducing ACLED: An armed conflict location and event dataset." *Journal of Peace Research*, 47(5): 651–660.

Raleigh, Clionadh, and Roudabeh Kishi. 2019. "Comparing conflict data: Similarities and differences across conflict datasets." *ACLED Documentation*, 651–660. www.acleddata.com/wp-content/uploads/2019/09/ACLED-Comparison_8 (last accessed Sep 13, 2024).

Rapoport, Hillel, Sulin Sardoschau, and Arthur Silve. 2021. "Migration and cultural change." IZA Discussion Paper 14772.

Ravenstein, Ernst Georg. 1876. *The birthplaces of the people and the laws of migration.* Trübner.

Ravenstein, Ernst Georg. 1885. "The laws of migration." *Journal of the Statistical Society of London*, 48(2): 167–235.

Ravenstein, Ernst Georg. 1889. "The laws of migration." *Journal of the Royal Statistical Society*, 52(2): 241–305.

Rodriguez, Catherine, and Fabio Sanchez. 2012. "Armed conflict exposure, human capital investments, and child labor: Evidence from Colombia." *Defence and Peace Economics*, 23: 161–184.

Roy, Andrew D. 1951. "Some thoughts on the distribution of earnings." *Oxford Economic Papers*, 3(2): 135–146.

Ruyssen, Ilse, and Sara Salomone. 2018. “Female migration: A way out of discrimination?” *Journal of Development Economics*, 130: 224–241.

Saiz, Albert, and Elena Zoido. 2005. “Listening to what the world says: Bilingualism and earnings in the United States.” *The Review of Economics and Statistics*, 87(3): 523–538.

Sarvimäki, Matti. 2017. “Labor market integration of refugees in Finland.” *Nordic Economic Policy Review*, 7(1): 91–114.

Serdar, Kaya, and Phil Orchard. 2020. “Prospects of return: The case of Syrian refugees in Germany.” *Journal of Immigrant & Refugee Studies*, 18(1): 95–112.

Shemyakina, Olga. 2011. “The effect of armed conflict on accumulation of schooling: Results from Tajikistan.” *Journal of Development Economics*, 95: 186–200.

Shortt, Mitchell, Shantanu Tilak, Irina Kuznetcova, Bethany Martens, and Babatunde Akinkuolie. 2023. “Gamification in mobile-assisted language learning: A systematic review of Duolingo literature from public release of 2012 to early 2020.” *Computer Assisted Language Learning*, 36(3): 517–554.

Silva, J. M. C. Santos, and Silvana Tenreyro. 2006. “The log of gravity.” *The Review of Economics and Statistics*, 88(4): 641–658.

Sjaastad, Larry A. 1962. “The costs and returns of human migration.” *Journal of Political Economy*, 70(5, Part 2): 80–93.

Spilimbergo, Antonio. 2009. “Democracy and foreign education.” *American Economic Review*, 99(1): 528–543.

Statcounter. 2024. “Search Engine Market Share Worldwide.” https://gs.statcounter.com/search-engine-market-share (last accessed Sep 13, 2024).

Steinmayr, Andreas. 2021. “Contact versus exposure: Refugee presence and voting for the far right.” *The Review of Economics and Statistics*, 103(2): 310–327.

Stöhr, Tobias. 2015. “The returns to occupational foreign language use: Evidence from Germany.” *Labour Economics*, 32: 86–98.

Storesletten, Kjetil. 2000. “Sustaining fiscal policy through immigration.” *Journal of Political Economy*, 108(2): 300–323.

Strang, Alison, and Alastair Ager. 2010. “Refugee integration: Emerging trends and remaining agendas.” *Journal of Refugee Studies*, 23(4): 589–607.

Strezhnev, Anton. 2023. “Decomposing Triple-Differences Regression under Staggered Adoption.” *arXiv preprint arXiv:2307.02735*.

Sundberg, Ralph, and Erik Melander. 2013. “Introducing the UCDP georeferenced event dataset.” *Journal of Peace Research*, 50(4): 523–532.

Sun, Liyang, and Sarah Abraham. 2021. “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects.” *Journal of Econometrics*, 225(2): 175–199.

Tabellini, Marco. 2020. “Gifts of the immigrants, woes of the natives: Lessons from the age of mass migration.” *The Review of Economic Studies*, 87(1): 454–486.

The Gallup Organization. 2022. “Gallup World Poll .” https://www.gallup.com/178667/gallup-world-poll-work.aspx (last accessed Nov 18, 2022).

The World Bank. 2016. "World Development Report 2016: Digital dividends." The World Bank.

Tjaden, Jasper, Daniel Auer, and Frank Laczko. 2019. "Linking migration intentions with flows: Evidence and potential use." *International Migration*, 57(1): 36–57.

Tyazhelnikov, Vladimir, and Xinbei Zhou. 2021. "PPML, Gravity, and Heterogeneous Trade Elasticities." Unpublished Working Paper.

Uebelmesser, Silke, Ann-Marie Sommerfeld, and Severin Weingarten. 2022. "A macro-level analysis of language learning and migration." *German Economic Review*, 23(2): 181–232.

UNHCR. 2023. "UNHCR Ukraine refugee situation ." https://data2.unhcr.org/en/situations/ukraine (last accessed Oct 17, 2023).

UNHCR. 2023a. "Global trends: Forced displacement in 2023." https://www.unhcr.org/refugee-statistics/ (last accessed Oct 17, 2023).

UNHCR. 2023b. "Eighth regional survey on Syrian refugees' perceptions and intentions on return to Syria." https://data.unhcr.org/en/documents/details/100851 (last accessed Oct 17, 2023).

UNHCR. 2023c. "Syria Regional Refugee Response: Durable solutions." https://data.unhcr.org/en/situations/syria_durable_solutions.

United Nations. 2023. "United Nations Ukraine - subnational administrative boundaries ." https://data.humdata.org/dataset/cod-ab-ukr? (last accessed Apr 28, 2023).

United Nations Population Division. 2023. "International migrants: numbers and trends." https://worldmigrationreport.iom.int/what-we-do/world-migration-report-2024-chapter-2/international-migrants-numbers-and-trends (last accessed Sep 13, 2024).

US Census Bureau. 2015. "How Well Do You Speak English? Assessing the Validity of the American Community Survey English-Ability Question." https://www.census.gov/newsroom/blogs/research-matters/2015/10/how-well-do-you-speak-english-assessing-the-validity-of-the-american-community-survey-englis html (last accessed Oct 17, 2023).

Verwimp, Philip, and Jan Van Bavel. 2014. "Schooling, violent conflict, and gender in Burundi." *The World Bank Economic Review*, 28: 384–411.

Vesselinov, Roumen, and John Grego. 2012. "Duolingo effectiveness study." *City University of New York, USA*, 28(1-25).

Wang, Haining, Russell Smyth, and Zhiming Cheng. 2017. "The economic returns to proficiency in English in China." *China Economic Review*, 43: 91–104.

Weidner, Martin, and Thomas Zylkin. 2021. "Bias and consistency in three-way gravity models." *Journal of International Economics*, 132: 103513.

White, Roger, and David Buehler. 2018. "A closer look at the determinants of international migration: Decomposing cultural distance." *Applied Economics*, 50(33): 3575–3595.

Wong, Lorraine. 2023. "The effect of linguistic proximity on the labour market outcomes of the asylum population." *Journal of Population Economics*, 36: 609–652.

Wooldridge, Jeffrey M. 2015. "Control function methods in applied econometrics." *Journal of Human Resources*, 50(2): 420–445.

Wooldridge, Jeffrey M. 2021. "Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators." *Available at SSRN 3906345*.

Wooldridge, Jeffrey M. 2023. "Simple approaches to nonlinear difference-in-differences with panel data." *The Econometrics Journal*, 26(3): C31–C66.

Yang, Dean. 2008. "International migration, remittances and household investment: Evidence from Philippine migrants' exchange rate shocks." *The Economic Journal*, 118(528): 591–630.

Yang, Dean, and HwaJung Choi. 2007. "Are remittances insurance? Evidence from rainfall shocks in the Philippines." *The World Bank Economic Review*, 21(2): 219–248.

Yarkin, Alexander. 2024. "Does the 'melting pot' still melt? Internet and immigrants' integration." Unpublished Working Paper.

Yotov, Yoto V, Arne Nagengast, and Fernando Rios-Avila. 2024. "The European Single Market and Intra-Eu Trade: An Assessment with Heterogeneity-Robust Difference-In-Differences Methods." Drexel University School of Economics Working Paper Series 2024-5.

Young, Alwyn. 2019. "Channeling fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results." *The Quarterly Journal of Economics*, 134(2): 557–598.

Zaiour, Reem. 2023. "Violence in Mexico, return intentions, and the integration of Mexican migrants in the US." Unpublished Working Paper.

Zelinsky, Wilbur. 1971. "The hypothesis of the mobility transition." *Geographical Review*, 219–249.

Zuo, George W. 2021. "Wired and hired: Employment effects of subsidized broadband internet for low-income Americans." *American Economic Journal: Economic Policy*, 13(3): 447–82.