Frictions in International Trade

The Role of Culture, Regulations and Information

Matthias Beestermöller

Ludwig-Maximilians-Universität

Munich 2016

Frictions in International Trade

The role of culture, regulations and information

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

2016

vorgelegt von

Matthias Gerhard Beestermöller

Referent:	Prof. Dr. Carsten Eckel
Korreferent:	Prof. Gabriel Felbermayr, Ph.D.
Promotionsabschlussberatung:	10. Mai 2017

Datum der mündlichen Prüfung: 25. April 2017

Namen der Berichterstatter: Carsten Eckel, Gabriel Felbermayr, Lionel Fontagné

This dissertation
is dedicated to
my parents
Andrea and Gerhard,
who encouraged
me to pursue
my dreams
and supported me
in all things great
and small.

Acknowledgements

First and foremost, I would like to thank my supervisor Carsten Eckel. I am grateful for his critical questions, constructive comments and continuous support that have guided me while writing this dissertation. I would further like to thank my co-supervisor, Gabriel Felbermayr, for his guidance and continuous encouragement during the past three years. Thanks also goes to Lionel Fontagné for agreeing to serve on my committee, and inviting me to spend five months at the Université Paris 1 Panthéon-Sorbonne. The collaboration with Lionel and Anne-Célia Disdier has been, and continues to be, a source of inspiration and motivation. In addition, a special word of thanks goes to my co-author Ferdinand Rauch. Our collaboration on the first chapter of this dissertation, as well as his continuous academic and personal support, have been the driving factor behind my decision to pursue a doctoral degree in economics.

I benefited greatly from the exchange with numerous colleagues at the Munich Graduate School of Economics, the University of Paris and the members of the Chair of International Trade and Trade Policy in particular. They provided insightful comments and discussions, and fostered a stimulating research environment.

Special thanks go to the MeinFernbus Flixbus team, especially Tina Bosler and Anna Humpert, for trusting me to work with their booking data.

I also gratefully acknowledge the generous financial support from the Egon-Sohmen-Foundation, the Deutsche Forschungsgemeinschaft through GRK 1928 and the European Commission's 7th Framework Programme under Grant 613504.

Last but not least, I am deeply indebted to my family and friends, who endured my enthusiastic conversations about the Habsburg Empire, EU food safety regulations and German inter-city buses. They encouraged me throughout, and supported me in all my pursuits.

Contents

Preface

"The lowering of tariffs has, in effect, been like draining a swamp. The lower water level has revealed all the snags and stumps of non-tariff barriers that still have to be cleared away." (Baldwin, 1970)

Trade theory stresses the gains from trade: through comparative advantage and specialization (Dixit and Norman, 1980); economies of scale (Krugman 1979, 1980); inter-firm reallocations and selection into exporting (Melitz, 2003); intra-firm reallocations (Eckel and Neary, 2010; Bernard et al., 2011); input- and task-sourcing (Hummels et al. 2001; Grossman and Rossi-Hansberg, 2008); and innovation (Verhoogen, 2008; Lileeva and Trefler, 2010). Indeed, although globalization creates winners and losers (Felbermayr et al., 2011; Autor et al., 2013), the overall gains from trade are now empirically well documented (Arkolakis et al. 2012; Feyrer, 2009a,b).

There remains a puzzle, however: most trade models predict significantly more trade than can be observed empirically (Head and Mayer, 2013; Anderson, 2000; Trefler, 1995). The economics literature has gained many insights into the quantitative effect of tariffs and quotas, as highlighted in reviews by Yeaple (2013) as well as Hornok and Koren (2016). However, tariffs and quotas are far too low to account for the difference between actual and predicted trade volumes (Grossman, 1998). This 'missing' trade is also visible at the extensive margin: the matrix of disaggregated bilateral trade flows displays a large number of zeros (Helpman et al., 2008; Armenter and Koren, 2014). If there are untapped gains from further market integration, the missing trade requires an explanation. A prominent suggestion is that non-tariff border frictions remain sizeable – the 'dark matter' of trade costs (Anderson and van Wincoop, 2004; Head and Mayer, 2015). The nature and importance of these trade frictions are poorly understood. What is the true extent of non-tariff barriers? How important are intangible factors such as culture, regulatory costs, and information? Answers to these questions can help to reconcile the gap between predicted and actual trade flows and ultimately inform the debate on globalization.

The trade literature on border frictions is in its infancy; many questions remain open for future research. However, there have been some significant advances in the recent literature as models have moved away from treating trade costs exclusively as *ad-valorem* costs. New approaches include fixed entry costs (Das et al., 2007), time costs (Hummels and Schaur, 2013), per-unit costs (Irrazabal et al., 2013), and per-shipment costs (Hornok and Koren, 2015). Yet despite these advances, a better and more quantitative understanding of the causes of cross-border frictions is still needed.

This dissertation comprises three empirical chapters, which work towards a better understanding of non-tariff and non-quota frictions to trade and, more generally, to economic exchange. Under this common theme, I study a diverse range of topics. In particular, the dissertation contributes insights into the importance of frictions, such as the effects of cultural familiarity (Chapter 1), import regulations (Chapter 2), and the role of information (Chapter 3). Chapter 1 provides an examples of how frictions may re-direct trade. It highlights the importance of cultural links between former members of the Habsburg monarchy. Chapter 2 provides a direct measure of cross-border barriers. Chapter 3, which focuses on the German inter-city bus market, is not a trade paper. However, the studied effect, namely frictions due to habits or lack of information, offers a broader applicability that extends to international trade. From a methodological point of view, all chapters share a strong emphasis on empirical analysis using newly collected datasets. Although thematically related, the three chapters in this dissertation are self-contained and can be read independently.

Chapter 1, which is based on joint work with Ferdinand Rauch, studies trade in Europe after the fall of the Iron Curtain. We show that the countries of the former Austro-Hungarian monarchy traded significantly more with one another after 1989 than predicted by a standard gravity model. The surplus trade is approximately four times the effect of a monetary union in 1990. This surplus then declines linearly and monotonically and becomes statistically insignificant after two decades. Both the initial surplus trade between the former members of the Austro-Hungarian monarchy after 1989, and its subsequent decline need to be accounted for.

We argue that these results can best be explained by dissolving ‘trading capital’, a term coined by Head et al. (2010). They find that after independence former colonies continue to trade for a long period with their colonizers, but at a declining rate. Trading capital is built up during colonization, and deteriorates after independence. In Chapter 1, we think of trading capital in three broad categories: physical capital, such as roads or railways; capital relating to direct human interaction; and a third category capturing all other factors facilitating trade such as notions of cultural familiarity.

The collapse of the Austro-Hungarian monarchy offers a natural experiment setting in which we can observe some components of trading capital. Prior to its collapse in 1918, the monarchy was a well-integrated and interconnected market with significant trading capital. The Iron Curtain divided the East and West of the old monarchy between 1945 and 1989. As a result, all formal and business relationships between East and West were severed, almost all trade ceased, and maintaining personal contacts became very costly. Transport infrastructure linkages were left to deteriorate. Institutions and norms diverged into two distinct blocks.

We argue that the surplus trade observed between East and West after 1989 overwhelmingly results from the third category of trading capital; historical legacies and cultural linkages persisted. Trading capital, established under Habsburg rule, survived over four decades of separation and provided an initial boost to trade. This proved short-lived: the surplus trade disappeared rapidly as countries rearranged themselves according to changing geopolitical circumstances.

Chapter 1 contributes to the literature by showing that the degree to which cultural forces influence trade appears to be significant. While trade that is once interrupted takes a long time to recover (Felbermayr and Gröschl, 2013; Nitsch and Wolf, 2013), we demonstrate that linkages between countries are highly resilient once built up. This chapter thus adds to the growing literature which emphasizes the long persistent effects of borders, institutions and culture (Guiso et al., 2009; Becker et al., 2014). Further, we contribute to this literature by providing an example and new measure of both the resilience of such historical and cultural effects on trade, as well as on its decline.

Chapter 2, which is a joint work with Anne-Célia Disdier and Lionel Fontagné, studies the microeconomic impact of rejection risk at European borders on safety grounds. We examine how the risk affects Chinese agri-food exporters. Despite low tariffs, access to the EU remains difficult because individual exporters are required to meet stringent safety regulations.

Using a rarely exploited dataset of information from the European Rapid Alert System for Food and Feed (RASFF) combined with Chinese firm-level export data, we analyse the impact of border rejections on firms' export decisions. We find that Chinese exporters of agri-food products are more likely to exit the European market if the product they export has been rejected in previous years. At the same time, rejections favour the entry of new firms. Thus, border rejections increase turnover at the extensive margin of trade. Furthermore, the impact is heterogeneous across firms. Small firms are affected more strongly than big firms by this turnover. At the intensive margin, border rejections boost the exports of surviving firms. This suggests some re-allocation effect towards big and productive exporters.

Chapter 2 contributes to the literature in three ways. First, we provide a more nuanced understanding of the uncertainty component of non-tariff measures (NTMs), which has, somewhat surprisingly, been largely overlooked in the literature on NTMs and border inspections. In this regard, we particularly highlight the importance of information externalities and reputation effects. Second, whilst details on the occurrence of regulations give evidence on *de jure* NTMs, knowledge about rejections sheds light on their *de facto* trade impact. Border rejections represent an example of a specific trade-impeding NTM where regulations are enforced. Third, to the best of our knowledge, this chapter is the first to study the effect of sanitary measures on firm-level exports from a large and significant developing economy. We pay explicit attention to the role of firm heterogeneity, and show that big firms are more resilient to the risk of border rejections.

Chapter 3 studies the effect of the 2014-2015 rail strikes on German inter-city buses. I combine three novel and extremely rich datasets: detailed booking data provided by Germany's largest bus provider MeinFernbus (MFB), emergency timetables published by Deutsche Bahn (DB) during the strikes, and a dataset of all rail itineraries. This data is used to study how the rail strikes affected bus ticket sales and to test for persistence as rail operations returned to normal.

Firstly, I ask which bus routes were most affected *during* the rail strike. While the exposure of *rail routes* to the strike can be deduced from the emergency timetables, the exposure of *bus routes* is not ex-ante clear to the researcher. On the one hand, travellers might not have had sufficient information about their route's exposure to the rail strike to decide if they could remain with DB services. On the other hand, travellers may only switch to inter-city buses if the bus service is a close enough substitute to rail. I find that the primary channel that drives ticket sales during the strike is whether the absolute bus travel time was sufficiently short. The variation in rail service cancellations across routes does not explain increased bus bookings: travellers switched to buses even on routes with little or no service cancellations. It follows that either travellers were not well informed about their exposure to the rail strike, or they had no trust in DB's ability to implement the emergency timetables.

Secondly, I study whether the effect of the strike was persistent. Did short routes have higher ticket sales *after* rail operations returned to normal? In a difference-in-differences framework, I compare the change in the number of customers between high and low strike-exposed routes to identify any demand persistence. Although the common trend assumption does not seem to be completely tenable in the given context, my results point to a persistent effect on the ticket sales for inter-city buses on the affected routes. I follow the methodology of Nunn and Qian (2011), who employ a similar strategy in a different setting. They estimate period-specific treatment effects for the *pre-period* in order to compare these to the post-treatment coefficients. Following their methodology, my results also remain largely unaltered to a number of alternative specifications and robustness checks.

Chapter 3 contributes to the literature by highlighting an unintended and potentially positive effect of a rail strike. If the strike revealed information about an alternative transport mode, it may have been welfare improving (Larcom et al., 2016). Chapters 3 is not a trade paper. However, the focus on the German inter-city bus market offers a broader applicability that extends to international trade. It provides an example of how information asymmetries may re-direct trade. Some customers, who were forced to experiment with buses, discovered that their previous choice of rail was not optimal. This chapter supplements the classic literature relating to the way in which individuals decide between alternatives (Weitzman, 1979; Morgan and Manning, 1985). My results cannot be reconciled with the classical economic assumption of perfectly informed and rational consumers. I contribute to this literature by providing an example of the Porter hypothesis: exogenous shocks may help individuals find better choices by triggering experimentation (Porter, 1991). Chapter 3 complements the findings by Larcom et al. (2016), who study the effect of a London Underground strike, in two ways. Firstly, I study inter-modal switching across transport modes for inter-city transport – a less frequent travel decision than daily commuting. Secondly, the longer post-strike period allows me to better understand the short- and medium-term impacts of the strikes.

To summarize, this dissertation adds three empirical chapters to the current debate on the role of non-tariff barriers to economic exchange, filling several gaps in the academic literature. The findings contribute to a better understanding of cross-border frictions, making it easier for researchers to give informed answers to policymakers. Using the example of trade among countries of the former Austro-Hungarian monarchy after 1989, Chapter 1 shows that the degree to which cultural forces influence trade appears to be large. Chapter 2 studies the effect of border rejection risk on Chinese firms. Following a spell of rejections, the number of firms tends to decrease but the size of the surviving firms increases. Chapter 3 estimates the persistent effect of the 2014-2015 rail strikes on the demand for inter-city buses in Germany. The strike induced some customers, who would have routinely stayed with rail, to permanently switch to buses. In these ways, this thesis seeks to contribute a valuable piece of research towards a better understanding of all those "snags and stumps of non-tariff barriers" (Baldwin, 1970, *op. cit.*) and ultimately to close the gap between predicted and actual trade flows.

Chapter 1

A Dissection of Trading Capital: Trade in the Aftermath of the Fall of the Iron Curtain

This chapter is based on joint work with Ferdinand Rauch. We are grateful for comments and suggestions of numerous colleagues and seminar participants including but not limited to James Anderson, Daniel Baumgarten, Tibor Besedes, Johannes Van Biesebroeck, Carsten Eckel, Peter Egger, Thibault Fally, Gabriel Felbermayr, Lisandra Flach, Lionel Fontagné, James Harrigan, Keith Head, Harald Heppner, Michael Irlacher, Beata Javorcik, Amid Khandelwaal, Helmut Konrad, Anna Koukal, Thierry Mayer, Guy Michaels, Peter Neary, Volker Nitsch, Chris Parsons, Steven Poelhekke, Monika Schnitzer, Jens Südekum, Pierre-Louis Vezina, Daniel Wissmann as well as seminar and conference participants from the LETC conference in Slovenia, EGIT Düsseldorf, Paris, Göttingen, ETSG Munich, Vienna and Zurich. We particularly thank two very helpful anonymous referees.

1.1 Introduction

In 1989 the Iron Curtain fell quickly and unexpectedly, ending the separation between Western Europe and the Soviet Union. After 44 years of an almost completely sealed border, trade was suddenly free to reconnect. Despite the political and economic turmoil within the Eastern regimes, trade between West and East almost doubled within five years after 1990. By the year 2000 it had almost tripled. We study this trade in the aftermath of the collapse of the Soviet Union. We pay special attention to Austria, a country that has engaged in trading opportunities beyond what would be expected given its size and geographic location, and might have been the main western beneficiary of Europe's economic expansion eastwards.

In a standard gravity equation setting we document that Austria indeed trades more with countries east of the Iron Curtain after 1990 than gravity would predict. However, we find that this effect is only found for the members of the former Habsburg Empire[1]. It declines linearly and monotonically, and in our preferred specification becomes statistically insignificant after a decade while the predicted magnitude becomes zero after two decades. The magnitude of the Habsburg surplus trade in 1990 is very large, about four times the effect of a monetary union. We find no similar surplus trade for other western countries with the East.

We argue that these results can best be explained by assuming a deterioration of specific components of 'trading capital' built up during the Habsburg years. The 44 years of Iron Curtain division severed all formal and business relationships, almost all trade between East and West, and made personal contacts very costly.

[1] Throughout this chapter we use the terms 'Habsburg monarchy', 'Habsburg Empire' and 'Austro-Hungarian monarchy' interchangeably, knowing that Austro-Hungary is only valid since 1867. We usually refer to the Empire in its extension shortly before World War I, as displayed in Figure 1.1. Former Habsburg members include Austria, Bosnia and Herzegovina, Croatia, Czech Republic, Hungary, Italy, Poland, Romania, Serbia, Slovakia, Slovenia and the Ukraine to differing degrees as detailed in Table 1.1 and Figure 1.1.

However, historical legacies and cultural linkages persist and a Habsburg surplus trade survives. Its decline reflects the continued dissolution of trading capital and the build-up of trading capital with other countries in Western Europe.

The term 'trading capital' is introduced by Head, Mayer and Ries (2010, from here on we refer to this paper as HMR) who show that after independence former colonies continue to trade for a long period with their colonizers, at a declining rate. They suggest that this observation might point to the presence of trading capital that is built up during colonization, and deteriorates after independence. Trading capital consists of various components that we can divide into three broad categories that facilitate trade: (i) physical capital, such as roads, railway lines or pipelines that connect countries and directly facilitate trade through reduced bilateral trade costs; (ii) capital relating to personal communication, direct human interaction and contacts or trust built up in repeated games, such as provided in structures of multi-national firms, joint ventures or by frequent personal contacts and trust won through repeated interaction; and (iii) all other variables that facilitate trade that are not based on personal interaction and formal or physical structures. These include all notions of cultural familiarity, such as those facilitated by cultural norms, language, history, consumers' familiarity with products, trust based on similarity and familiarity of people with each other. In the case of the Habsburg Empire this may relate to people in whose minds the Habsburg monarchy was the last functioning state before the hardship of the wars and communism. This may have created a brief nostalgic impulse to return to the old state of affairs when the possibility came. Indeed, below we verify a positive Habsburg bias in the cultural data by Felbermayr and Toubal (2010). Category (iii) may also include past decisions on institutional design and standards as basic as which side of the road to drive on or what type of light bulbs to adopt. However these latter effects are less relevant in the present example as such standards were fully harmonized across continental Europe by 1990.

We argue that the declining surplus trade of Habsburg countries after 1989 is comparable to the dissolving trading capital described by HMR, but given the history of Central Europe only relates to that part of trading capital that was not isolated by the Iron Curtain, the elements described in point (iii). At the beginning of the century the Habsburg monarchy was a politically and economically well integrated country. In the second half of the century it was split into two parts that were strictly separated for 44 years by the Iron Curtain. During the separation all formal institutions of the Empire ceased to exist as there were several waves of drastic institutional changes especially east of the Iron Curtain. Personal relationships were hard to maintain, and multinational firms connecting East and West as well as other formal institutions were broken apart. Physical transport capital such as railway lines, pipelines and roads – already badly damaged in WWII – were deliberately destroyed, or left to deteriorate. At the same time institutions and norms converged both within the East and within the West of the Iron Curtain into two distinct blocks. The historical circumstances thus offer a natural experiment setting in which we can observe some components of trading capital only between members of the former Habsburg Empire. In particular, any surplus trade observed after 1989 will overwhelmingly include those parts of trading capital that relate to point (iii) above. Comparing these effects to HMR we find that these forces explain a quantitatively large part of trading capital, and that they deteriorate at a rate smaller than suggested for all trading capital by HMR.

We add direct evidence for this hypothesis in five ways. First, we show that this surplus trade appears for the Habsburg countries, but not for a number of placebo combinations between western and eastern countries in Europe. We also verify that our main finding, the declining surplus trade for Habsburg countries is highly robust to alternative empirical strategies. When looking at product level, we see the effect mainly for homogeneous rather than heterogeneous goods. We would expect this if countries follow a heuristic not based on economic rationale alone, since homogeneous goods make substitution less costly. Fourth, we see that the effect is stronger for those goods that were traded in the Habsburg Monarchy. We rule out a number of possible alternative explanations. Finally, we cite research that points to some more general Habsburg nostalgia in the 1990s.

Our chapter adds to the literature showing that the degree to which such cultural forces influence trade seems to be large (for example, Algan et al., 2010; Disdier and Mayer, 2007; and Michaels and Zhi, 2010), linkages between countries are highly persistent once built up (Djankov and Freund, 2002 and Thom and Walsh, 2002) and trade once interrupted takes a long time to recover (Felbermayr and Gröschl, 2013; Nitsch and Wolf, 2013). There have been suggestions that culture matters more for trade than either institutions or borders (Becker et al., 2014). Our chapter also adds to a growing literature which emphasizes the long persistent effects of borders, institutions and culture. For example, Guiso et al. (2009) establish the importance of trust and cultural similarity on economic exchange. Meanwhile, Egger and Lassmann (2015) and Melitz and Toubal (2014) document the importance of common languages. However, it is difficult to distinguish between cultural similarity and ease of communication. Cultural proximity is inherently difficult to measure. A number of recent studies have thus used proxy measures for cultural proximity such as voting behaviour in the Eurovision Song Contest (Felbermayr and Toubal, 2010) or the United Nations General Assembly (Dixon and Moon, 1993; Umana Dajud, 2012). Lameli et al. (2015) show that the similarity of German dialects is an important predictor of trade within Germany. We add to this literature by providing an example and new measure of both the resilience of such historic and cultural effects on trade, as well as its decline.

Our chapter's methodology is related to Redding and Sturm (2008), who study the development of towns in West Germany and use the fall of the Iron Curtain as a natural experiment. Nitsch and Wolf (2013) document that it takes between 33 to 40 years to eliminate the impact of the Iron Curtain on trade within Germany. Our chapter mirrors Nitsch and Wolf (2013): While they show that borders remain visible in trade statistics long after they have been abolished, we demonstrate that borders take a long time to diminish trade when newly constructed. Djankov and Freund (2002) document that Russian regions continued to trade with each other 60 percent more in the period from 1994 to 1996, which is broadly consistent with our findings. Other studies that use a similar setting to our chapter are Schulze and Wolf (2009) who examine trade within the Habsburg monarchy in the late 19^{th} century and find that borders that later emerge become visible in price data long before the collapse of the Empire. Thom and Walsh (2002) study the trade effect of Anglo-Irish monetary dissolution and find little effect on trade. Becker et al. (2014) also present evidence on the importance of the Habsburg Empire on cultural norms. When comparing individuals living east and west of the long-gone Habsburg border, they find that people living on territory of the former Habsburg Monarchy have higher trust in courts and police. They argue that the former Empire had an enduring effect on people's values through its decentralized, honest and widely accepted state bureaucracy.

Trade is only one of many possible measures that could be influenced by historical legacies and cultural persistence. Migration and FDI might be others. Like HMR we choose to discuss this effect in terms of trade given that trade is recorded in a more consistent way and at a higher frequency than the aforementioned other measures. It is also less influenced by political decisions. For example, migration in Europe remained highly politically regulated until the EU enlargement, and migration numbers are thus politically constrained.

This chapter proceeds as follows: after a brief historical overview concerning the decline of the Habsburg Empire, the Iron Curtain and the reunion of the continent as far as these events concern our study in Section 1.2, we discuss our empirical strategy in Section 1.3. We then present our estimates of the surplus trade and its decline among former Habsburg countries in Section 1.4 and Section 1.5, which focus on product level results. Section 1.6 discusses the implications of the surplus trade and Section 1.7 concludes. Appendix A provides more details on the construction of the dataset, and shows a few additional results and robustness tests.

1.2 Historical overview

We focus on the borders of the Habsburg Empire just before the outbreak of World War I as displayed in Figure 1.1. While the Habsburg family had ruled the Empire for many centuries with changing borders, unification attempts and the introduction of a centralized administration came fairly late in the course of the 18^{th} century.[2] For our purposes it is important that the monarchy maintained a large, stable and well integrated market with large internal trade flows throughout its last decades:

In 1913 the Austro-Hungarian Empire had a large degree of ethnic and linguistic diversity, not only across the empire as a whole, but also within major sub-state regions and cities. All parts of the monarchy were linked by a common official language, common legal institutions and administration, as well as an expanding rail network. A strong emphasis on free trade strengthened the economic integration and trade flows within the country throughout the 19^{th} century (Good, 1984). The monarchy possessed a fully integrated monetary union with full control maintained by the Austro-Hungarian Bank in Vienna. Fiscal policy of the Empire was run as a

[2] In the 13^{th} century Rudolf von Habsburg acquired the thrones of Austria and Styria, which his family held until the first half of the 20^{th} century. The Habsburg monarchy expanded over the centuries mainly through skilful marriage policy, but also frequently lost territory in battle. The territory ruled by this family always incorporated different languages, customs and religions, which especially in the early years were allowed to flourish locally. There was little superstructure until the reforms under Maria Theresia and Josef II. helped by chancellors Kaunitz and Metternich in the course of the 18^{th}.

FIGURE 1.1
Austro-Hungarian Empire in 1910 and modern country boundaries

Source: Habsburg map is from Jeffreys (2007), and the modern country boundaries come from Eurostat (2013).

joint operation with separate budgets in Austria and Hungary contributing to the same common imperial expenditures and debt services (Dornbusch, 1992).

The monarchy consisted of 53 million people, numbering 13 percent of the total European population and producing 10 percent of Europe's GDP. As these figures imply, the economic condition of the Austro-Hungarian monarchy in its final decades prior to 1913 was poor in comparison to other European countries.[3] Before the collapse of the Empire some internal trade barriers became visible in price data at the end of the 19^{th} century, and nationalism was on the rise long before the collapse contributing to it (Schulze and Wolf, 2009 and 2012). Yet these studies highlight that the Empire possessed a heavily integrated internal market at the beginning of the 20^{th} century regardless of these tendencies. The monarchy further

[3]For example, Schulze (2010) documents poor performance in terms of GDP per capita growth for the monarchy between 1870 and 1913, and even uses the term 'great depression' to describe the situation in the western half of the Empire in 1873.

consisted of a well-functioning administration that unified the workings of many institutions across the countries it governed. The importance of the attachment of people to the imperial administration and its government, and the political, economic and cultural integration of its parts is highlighted by Clark (2013)[4] and Boyer (1989)[5] among other historians.

The end of World War I brought about a number of declarations of independence, which were sealed by the treaties of Saint Germain (1919) and Trianon (1920). New borders were drawn and new countries appeared, following considerations of ethnicity, language and trade networks. All the newly founded democracies on the territory of the former monarchy now included large numbers of ethnic and linguistic minorities. The newly founded Republic of Austria was left with 23 percent of the population of the former monarchy. Trade between countries of the former monarchy remained high in the 1920s. De Menil and Maurel (1994) present some evidence for strong trade in the years 1924-26 among successor states of the former monarchy, roughly of the magnitude of trade within the British Empire at that time. They explain the persistence of trade pointing to common history, shared linguistic and cultural ties, and mention the importance of business and personal relations as well as networks – all parts of trading capital. Institutional drift, however, started. New and different currencies were introduced. For example, Hungary replaced the Austro-Hungarian korona by its own korona after independence only to replace it again by the pengo in 1925 and forint in 1946 following hyperinflation. The Austrian-Hungarian national railways was also split

[4] "[The administration] was an apparatus of repression, but a vibrant entity commanding strong attachments, a broker among manifold social, economic and cultural interests. [...] most inhabitants of the empire associated the Habsburg state with the benefits of orderly government: public education, welfare, sanitation, the rule of law and the maintenance of a sophisticated infrastructure."

[5] " [...] competing popular and ethnic groups all had access to these public institutions [...] and these social groups quietly obtained some of their most sought after cultural attainments by means of these mechanisms, one might argue that the political and institutional history of the Empire presents [...] a state system that was not only more than the sum of its social parts, but was also psychologically consubstantial with those parts."

into multiple corporations, though traffic across the former monarchy continued at a significant pace.

World War II disrupted trade substantially, and it did not recover in the aftermath. Beginning in 1947, communist regimes in Central and Eastern Europe emerged under Soviet rule. The Sovietization of these economies caused a breakdown of their trade relations with the West, and foreign trade was organized as a strict state monopoly. Much of this remaining trade was arranged from Moscow, and negotiated at the highest political level, often as part of political bargains. An example for this was the export of goods worth 6.6 billion Austrian schillings in the aftermath of its independence in 1955 to the Soviet Union (Resch, 2010). Pogany (2010) writes on the relationship between Austria and Hungary: "Economic ties [...] became insignificant in the years following World War II. Centuries-old relations were reduced to a minimal level [...]." While Moscow took control of trade in the Eastern countries, on the western side trade was also heavily politically influenced. The main driver of this was the Co-ordinating Committee for Multilateral Export Controls (COCOM), established in 1949, an institution to organize embargoes against Soviet countries. Austria did not formally become a COCOM member, but its Eastern trade was influenced heavily by it under the obligations coming with Marshall aid (Resch, 2010). Economic cooperation was politically motivated and largely symbolic.

Large parts of infrastructure, especially the railways, were destroyed by the war – they would only partially be rebuilt taking into account the new borders that had emerged. An anecdote might highlight the poor recovery of infrastructure: The two capitals closest to each other in Europe are Vienna and Bratislava, at a distance of less than 60 kilometres. During the time of the monarchy there was a tramway that connected both cities, the 'Pressburger Bahn'. There has been no similar connection attempt since 1990, and thus the time to travel from one city to the other is now larger than it was in 1900.[6]

The Iron Curtain was an ideological boundary, but also primarily a geographical border. The most substantial cut to trade relations was brought about by the erection of the physical Iron Curtain, whose construction begun in 1949. The new border ran right through the former Habsburg countries, splitting Austria and the formerly Austrian parts of Italy from the rest. After the Hungarian Uprising of 1956 the already very limited possibility of transit ceased and all activity crossing this border was further suppressed. The border was sealed by barbed wire, land mines, high voltage fences, self shot systems and other means. Only people with appropriate restrictions were allowed close to the border. As such the Iron Curtain thus presented a completely sealed border that cut off all former local economic activity between the two sides (Redding and Sturm, 2008).

Furthermore, the economies of Hungary and Czechoslovakia switched to central planning. Multinational companies were split, personal interaction and communication over the border became increasingly difficult and rare. To put the decline of trade in numbers, Austrian imports from Hungary fell from 10% in 1929 to 2% in 1959 and 1% in 1988, and imports from Czechoslovakia fell from 18% to 4% and 1% in the same period (Butschek, 1996; Lazarevic, 2010); numbers indicate shares of total Austrian imports). At the same time, Hungarian imports from Austria went from 77% in 1911-13 to 60% in 1920, to 5% in 1946 and then to below 4% in 1974 (Pogany, 2010). This collapse in trade includes estimates of black market activity. As we show in Appendix A.2, the trade relationship of Austria with Poland, Hungary and Czechoslovakia was essentially flat compared to the relationship with Germany in the years before 1990.

[6] In the discussion of the results below we provide further examples of abandoned infrastructure between East and West.

The relationships of the West with Yugoslavia were different to those with Hungary and Czechoslovakia as Yugoslavia – despite being socialist and autocratic – maintained looser ties with Moscow (Lazerevic, 2010). This allowed the United States to contribute to aid programmes from 1952. Eventually this even led to the accession of Yugoslavia to GATT in 1966. Yugoslavia maintained sizeable trade relationships with the West, which in some years even exceeded its trade levels with the Comecon countries. Given its coastal location, its main trade partners in the West between 1955 and 1986 were the EEA countries (Belgium, Luxembourg, France, Italy, the Netherlands, West Germany, Great Britain, Denmark and Ireland). For example, in 1986 Yugoslav exports to the EEA countries were over 7 times as large as exports to EFTA (Austria, Norway, Portugal, Sweden and Switzerland) (Lazerevic, 2010), which suggests that trade between Yugoslavia and Austria was not particularly developed during the Cold War.

We mention only two properties of the fall of the Iron Curtain which are important here, namely that it happened fast and that it was received by almost everyone on either side of the border with surprise (Redding and Sturm, 2008).

These large changes of the map of Central Europe in the course of the 20^{th} century are displayed in Figure 1.1. The map shows modern country boundaries and a map of the Habsburg Empire as of 1910. Table 1.1 displays the percentage of modern territory that was part of the Austro-Hungarian Empire for modern countries. Most of the countries that were part of the Empire are in the east, by which we indicate countries that were on the eastern side of the Iron Curtain, to which we count the countries of former Yugoslavia. These countries are Bosnia and Herzegovina, Croatia, the Czech Republic, Hungary, Slovakia, Slovenia as well as parts of Poland, Romania, Serbia and the Ukraine. On the western side of the Iron Curtain we only find Austria and South Tyrol, which is now part of Italy.

Table 1.1
Habsburg Members

Country	Share of land that was Habsburg	East	Year of EU accession	Year of Euro adoption
Austria	1	0	1995	1999
Bosnia and Herzegovina	1	1	–	–
Croatia	1	1	2013	–
Czech Republic	1	1	2004	–
Hungary	1	1	2004	
Italy	0.05	0	1952	1999
Poland	0.12	1	2004	–
Romania	0.44	1	2007	–
Serbia	0.25	1	–	–
Slovakia	1	1	2004	2009
Slovenia	1	1	2004	2007
Ukraine	0.12	1	–	–

Notes: Share of land that was Habsburg denotes the share of the area of the modern country that was part of the Habsburg monarchy in the year 1910. The Habsburg dummy consists of countries with values of 1 in Column 1. Missing values in the last two columns indicate no membership in 2013.

There is plenty anecdotal evidence on Habsburg nostalgia after 1990 in former members of the monarchy. Wank (1997) describes a consensus view of historians of the 90s that was nostalgic of Habsburg and run the risk of "distorting historical reality [...] by emphasizing the monarchy's positive qualities [...]". Furthermore, historians of the time also implied that "some substitute for Austria-Hungary in Central Europe must be created" and "there is a legacy of positive lessons that the Habsburg Empire has bequeathed to Europe." Hartmuth (2011) writes that the monarchy was not remembered as ‘prison of nations’ any longer but a multicultural empire, and they point to Hungarian Salami meat being sold with the counterfeit of Franz Joseph. Becker et al. (2014) discuss the long cultural legacy of Habsburg on outcomes such as trust. For some speculation on the reasons behind Habsburg nostalgia see Schlipphacke (2014).

1.3 Empirical strategy and data

To investigate persistence after decades of Cold War of Austrian trade with countries east of the Curtain (Austria-East[7]) and members of the former Habsburg monarchy, we largely follow the methodology applied by HMR. They develop a method to address a closely related question, and the similarity allows us to compare our estimates to theirs. We estimate gravity equations to which we add $(Austria \times East) \times year$ and $Habsburg \times year$ dummies, which are our principal variables of interest. We run the estimations once jointly with Austria-East and Habsburg dummies and once separately only including one set of dummies interacted with year. We use the boundaries of the Habsburg Empire in its last days. The gravity framework captures the counterfactual multinational trade had there been no Habsburg relationship. The $(Austria \times East) \times year$ and $Habsburg \times year$ indicators capture any trade in excess of what the gravity model alone would predict.

The well-known empirical and theoretical formulations of the gravity equation can be represented in the following form:

$$X_{int} = C^{ex}_{it} C^{im}_{nt} \phi_{int} \tag{1.1}$$

where X_{int} denotes importer n's total expenditure on imports from origin i in year t, C^{ex}_{it} and C^{im}_{nt} are origin and destination attributes in a specific year, and ϕ_{int} measures bilateral effects on trade.[8] Since there is no set of parameters for which Equation 1.1 will hold exactly, the conventional approach is to add a stochastic term and estimate after log-linearizing. We follow the commonly practised gravity approach. Head and Mayer (2013) and Egger (2000) provide overviews of this tech-

[7]A variable indicating a trade flow between Austria and a country east of the Iron Curtain
[8]We follow HMRs notation here.

nique including a number of theoretical foundations which yield gravity equations. In particular, we estimate the following equation:

$$\begin{aligned}\ln(X_{int}) &= \mu_{it} + \mu_{nt} + \gamma D_{int} \\ &+ \delta^{(Aus\times East)}(Aus \times East)_{int} \\ &+ \delta H_{int} + \delta^{east} H^{east}_{int} + \epsilon_{int},\end{aligned} \tag{1.2}$$

where μ_{it} and μ_{nt} denote origin- and destination-year fixed effects respectively and δ coefficients to be estimated. The inclusion of sets of fixed effects interacted with year makes separate time fixed effects as in Equation 1.1 multicollinear and thus redundant. Matrix D_{int} denotes pairwise covariates that may be time varying or not. In an effort to distil the main effect of interest as precisely as possible, we include as detailed fixed effects as possible. In particular, we include the variables shared border, common official and spoken language and common legal institutions as time varying dummy variables to flexibly account for the many possible changes in the cultural and political climate in Europe during this period. These sets of control variables make it redundant to control for the standard right hand side variables measuring the size of countries, such as population and income, and allow only to include bilateral variables that vary over time. We include bilateral indicators for the distance between both countries, indicators for a shared border, an officially joint language, a joint spoken language, common legal institutions, common religion, common currency, the presence of a regional trade agreement as well as indicators if both are members of the EU, the Eurozone, or to the east of the Iron Curtain. All these standard bilateral control variables are taken from the standard source for this type of estimation, and precise definitions are given there (Mayer and Zignago, 2011).[9]

[9] To summarize a few key properties of these control variables: distance is measured as the crow flies. Common legal institutions are countries that share Civil Law, Common Law or Muslim Law. The shared religion variable relies on a breakdown for Buddhist, Christian Roman Catholic, Christian Orthodox, Christian Protestant, Hindu, Muslim. Having at least 9 percent of

The main variables of interest are the bilateral coefficients on the interaction term $(Aus \times East)_{int}$, dummies indicating if the observed flow is between Austria and a country east of the former Iron Curtain, and H_{int} which indicates if both countries were once part of the Austro-Hungarian monarchy in year t. Since we are only interested in Habsburg trade that crosses the Iron Curtain, we also include a H_{int}^{east} variable, which captures all trade east of the Curtain (there is only Austria west of the Curtain in our baseline specification). Intuitively we estimate how the fraction of Austria-East and Habsburg surplus trade evolves over time. We use a comprehensive set of indicators to capture the different types of Habsburg trade. For our main variable we restrict our measure of Habsburg economies to only those which were fully part of the Habsburg monarchy: Austria, Hungary and former Czechoslovakia. We argue that this is the safest approach as including other economies which were only partly part of the Empire, such as Italy, may pick up effects not specific to the Habsburg relationship. In Appendix A.2 we show robustness to different choices of this Habsburg definition.

If we were to control for attributes of the exporter and importer using GDP per capita and populations, our specification would suffer from bias caused by omission of 'multilateral resistance' terms (Anderson and van Wincoop, 2003). Multilateral resistance terms are functions of the entire set of ϕ_{int} from Equation 1.1. We thus adopt the preferred method of the literature which is to introduce exporter-year and importer-year fixed effects.[10] This full fixed effects approach absorbs the exporting and importing specific effects.[11] Exporter- and importer-year fixed effects do not work for unbalanced two-way panels as pointed out by Baltagi (1995). If actual bilateral data are not balanced, as is the case in HMR, one should use the least square dummy variable (LSDV) approach. However, this concern is not relevant

the population with a shared language has become a standard threshold to measure a significant part of population in similar settings since Mayer and Zignago (2011).

[10]See Feenstra (2004) who addresses different techniques to take care of multilateral resistance within the gravity framework.

[11]See Egger (2000).

to our aggregated European data set which is balanced.[12] We therefore adopt the full fixed effects approach, even though this approach has the disadvantage that we cannot observe the coefficients of some the right-hand side variables typically used in gravity models .

We also address the issue of missing and zero trade observations. Zero and missing observations may be due to mistakes or reporting thresholds, but bilateral trade can actually be zero. We treat all missing trade observations as zero trade. Our linear-in logs specification of Equation 1.2 removes all observations of zero trade, thus introducing a potential selection bias. In the literature it has been common to either drop the pairs with zero trade or estimate the model using $X_{int} = 1$ for observations with $X_{int} = 0$ as the dependent variable.[13] In our baseline specification we choose to drop the zero pairs, but also run a robustness check replacing zeros as ones. We also adopt the Poisson Pseudo-Maximum-Likelihood (PPML) estimation technique. A natural step would be to use Tobit which incorporates the zeros, but it assumes log normality and homoskedasticity on the error term, so we prefer PPML. PPML incorporates zeros, and parameters can be estimated consistently with structural gravity as long as the data are consistent; i.e. provided the expectation of ϵ conditional on the covariates equals one.[14] The estimation method is consistent in the presence of heteroskedasticity.[15] Thus, it provides a natural way to deal with zero values of the dependent variable. We believe this preferable to other estimators without further information on the heteroskedasticity. However, it may be severely biased when large numbers of zeros are handled in this way (Martin and Pham, 2015). There are only 53 missing trade observations out of 13,200 observations in our data since we focus on estimating trade among

[12] Appendix A.1 lists our data sources and discusses our approach to minimize data inaccuracies.

[13] See, for example, Felbermayr and Kohler (2006).

[14] See Silva and Tenreyro (2006).

[15] Consistency of estimating Equation 1.2 depends critically on the assumption that ϵ_{int} is statistically independent of the explanatory variables.

European economies. The majority of missing trade values involve Albania as a trading partner for which trade may indeed be zero or so small that it falls below a minimum reporting threshold.[16]

The estimation equation for the Poisson Pseudo-Maximum-Likelihood (PPML) estimator expresses Equation 1.2 as

$$\begin{aligned} X_{int} = \exp(& \mu_{it} + \mu_{nt} + \gamma D_{int} \\ & + \delta^{(Aus \times East)} (Aus \times East)_{int} \\ & + \delta H_{int} + \delta^{east} H_{int}^{east}) u_{int}, \end{aligned} \tag{1.3}$$

where $u_{int} = \exp(\epsilon_{int})$.

Even though we include all the usual controls, our vector of bilateral variables may remain incomplete, so unobserved linkages end up in the error term. To capture possible omitted variables in ϵ_{int}, we estimate two additional econometric techniques: a lag dependent variable specification and a specification with origin-destination (bilateral or dyad) fixed effects. The lagged dependent variable would absorb unobserved influences on trade that evolve gradually over time. Including a lagged dependent variable biases coefficient estimates in short panel models.[17] Monte Carlo experiments suggest that the bias can be non-negligible with panel lengths of T=10 or even T=15 (Dell et al., 2014). However, the time series dimension of our panel (T=22) is likely long enough such that biases can be safely considered second-order. Furthermore, the lagged dependent variable technique will not deliver consistent estimates if there is a fixed component in the error term that is correlated with the control variables. We thus also run a specification with bilateral fixed effects. We can still obtain estimates of our coefficients of interest as our variation of interest is also varying over time (the Habsburg and Austria-East dummies are interacted by year). The bilateral fixed effects specification identifies

[16] See Appendix A.1 for more details on the data set.

[17] Nickell (1981) shows that the bias declines at rate $\frac{1}{T}$.

the effect of Habsburg membership based on temporal (within-bilateral) variation. In the bilateral fixed effects specification all time invariant bilateral variables drop out.

To summarize, we estimate the Habsburg and Austria-East coefficients of interest using four different estimation techniques closely following HMR: simple OLS, Poisson Pseudo Maximum Likelihood (PPML), lag dependent variable specification and bilateral fixed effects (Dyad FE), each with a strong set of fixed effects. Our typical estimation has in excess of 13,000 observations and is robust to heteroskedasticity. We run these four estimations on the joint set of Habsburg and Austria-East dummies and separately with one set of dummies interacted with year. In the product level regressions we run the same specifications, but restrict the set of products for which we run the regression in various ways. For example, we analyse homogeneous and heterogeneous products separately to compare estimates.

The sources and details related to the construction of our dataset are documented in Appendix A.1. All data we use and our treatment of them is standard throughout the related literature. Here we summarize a few decisions that we take. The dataset we use contains all European countries in the years from 1990 until 2011, the first year for which Comtrade data is available for all the countries of Europe after the fall of the Iron Curtain and the last year for which we found a complete set of data when we embarked on this project. We clean Comtrade data using the methodology of Feenstra et al. (2005). Trade data for the years before 1990 are available from sources other than Comtrade, which we do not use given concerns about the comparability of data. We use data for Europe only as we think that it provides a cleaner sample of countries to run the proposed tests than the entire world would, given greater similarity of shipping and other technology in Europe. The first OLS assumption that the correct model is specified is easier to justify in a sample of more similar countries. We aggregate a few countries to maintain a

FIGURE 1.2
Descriptive GDP and trade ratios
(ratios on year)

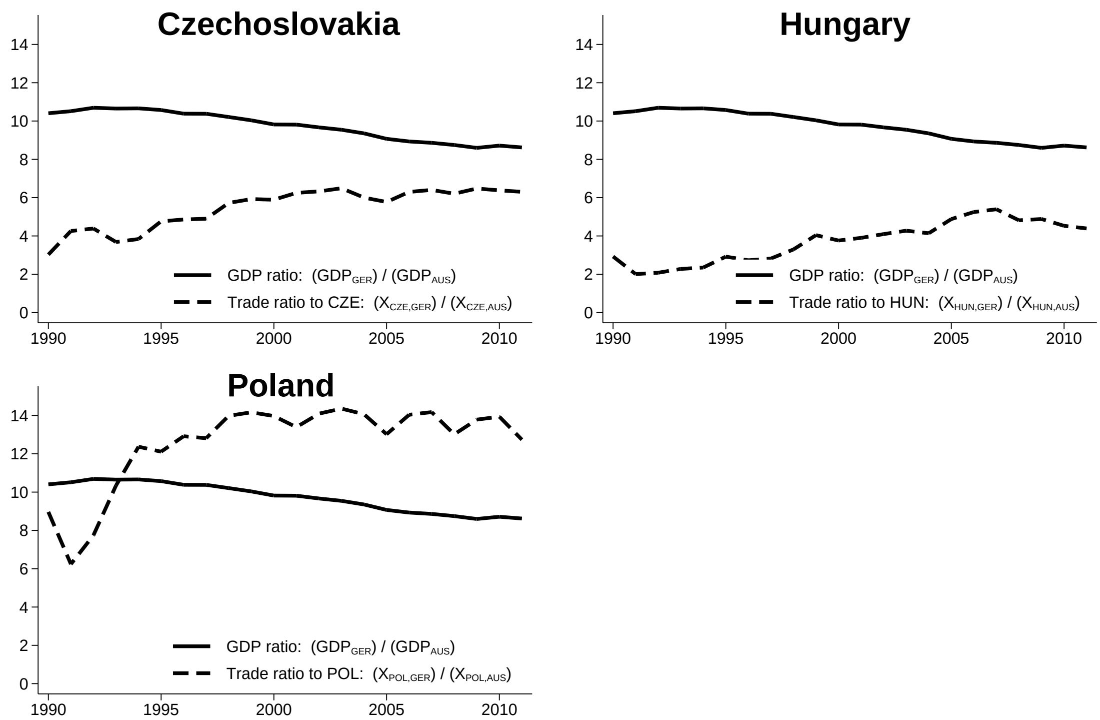

(*Datasource:* trade data from UN Comtrade, 2013; GDP data from World Bank WDI, 2013)

balanced panel, see details of this in Table A.1 in Appendix A.1. For the product regressions we use the well known BACI dataset from CEPII, details described in Appendix A.1. CEPII provided a BACI version that starts in 1992 for our countries, thus our product level analyses begin only in 1993 throughout.

Before turning to the regression results, we present some descriptive statistics which document the Habsburg trading surplus relative to Germany.[18] Figure 1.2 considers trade of Germany and Austria with Czechoslovakia, Poland and Hungary. Czechoslovakia borders on both Germany (both East and West) and Austria, thus differences in distance seem negligible. Moreover, changes in multilateral resistance

[18]We later use Germany as a placebo as it shares the language with Austria, and also directly borders many eastern countries. A risk of using that placebo might be that Germany could have also integrated faster with the East for its own particular history. However, as Nitsch and Wolf (2013) observe, there was "remarkable persistence in intra-German trade patterns along the former East-West border".

should also be fairly similar.[19] We plot the ratio of German to Austrian GDP $\left(\frac{GDP_{Gt}}{GDP_{At}}\right)$ and the ratio of German trade with Czechoslovakia to Austrian trade with Czechoslovakia $\left(\frac{X_{Ger,Cze,t}}{X_{Aus,Cze,t}}\right)$. If Habsburg did not matter, we would expect the ratio of trade to mirror the ratio of GDP (using GDP as measure for market and production size). However, we observe a large gap. In 1990 the German economy is roughly ten times as large as the Austrian economy. At the end of our sample period this ratio falls to about 8.5. However, trade with Czechoslovakia is only three times as large for Germany and this ratio rises to just over 6 over the sample period. We also conduct the same exercise for Hungary and Poland. On the one hand, Hungary – yet another core Habsburg member – displays an even starker gap. The trade ratio rises from approximately 2 to 4.5. These graphs highlight that Austria's trade with these two eastern countries was highly over-proportional given its size relative to Germany, but that this surplus steadily lowered over time. Even Poland, which we do not regard as a Habsburg member, since only 10 percent of its mass belonged to the monarchy, and which does not share a border with Austria, exported less than ten times its Austrian exports to Germany in 1990. All the countries show the central empirical finding in this figure, a strong Austrian surplus trade that weakens over time. We now turn to a more rigorous exploration of these suggested observations.

[19] A surge in French or Spanish GDP would have similar effects on Germany and Austria.

1.4 Results

We run three sets of regressions. First, we restrict the sample to Habsburg countries. Second, we include Austria-East dummies to investigate surplus trade with all of the East. Third, we control for Austria-East and Habsburg jointly and find that the effect for Austria-East becomes insignificant once we control for Habsburg. The first of these specifications is most important for our conclusion. We present it in detail and focus on the main elements of the other two.[20] It is worth emphasizing that we use origin interacted with year fixed effects and destination times year fixed effects separately in all of these regressions. The Habsburg surplus trade coefficients are bilateral and vary annually by construction. Thus, they are not multicollinear with the inclusion of this strong set of control variables and fixed effects.

In Table 1.2 we plot the Habsburg - year coefficients, which we interpret to be the surplus trade of Habsburg countries relative to what we would expect if trade followed our gravity model. These coefficients are also depicted in Figure 1.3. All four estimation methods display a steady decrease of the Habsburg surplus trade over time. We confirm that the first and last estimated coefficients are statistically significantly different to each other.[21] The downward slope of the trend given in Figure 1.3 is strongly significant in all of the specifications, and the slope is remarkably similar. It shows a strongly statistically significant, monotonic decline with a slope of around -0.044. Thus the main results, namely that the cultural component of trading capital declines over time, is insensitive to our estimation method. Note that the Habsburg trade bonus is large in the first year after the collapse of the Iron Curtain. For example, in the specification of column 1 the

[20] Tables reporting coefficients of control variables and the exact Habsburg and Austria-East coefficients are omitted for length but available upon request.

[21] F-test Probability $>$ F values are OLS: .008; PPML: .001; Lag DV: .768; and Dyad FE: .000.

TABLE 1.2
Estimation with Habsburg - year fixed effects only
Habsburg coefficients

	(1) OLS	(2) PPML	(3) Lag DV	(4) Dyad FE
Dependent variable:	$\ln(x_{int})$	x_{int}	$\ln(x_{int})$	$\ln(x_{int})$
1990	0.687***	0.919***		0.854***
	(0.257)	(0.199)		(0.253)
1991	0.613***	1.065***	0.00457	0.771***
	(0.227)	(0.151)	(0.113)	(0.220)
1992	0.477**	0.885***	0.0131	0.609***
	(0.232)	(0.154)	(0.108)	(0.206)
1993	0.514**	0.732***	0.150	0.612***
	(0.210)	(0.143)	(0.116)	(0.160)
1994	0.351	0.784***	-0.149*	0.459***
	(0.219)	(0.136)	(0.0812)	(0.158)
1995	0.367*	0.783***	0.00948	0.501***
	(0.216)	(0.164)	(0.0804)	(0.149)
1996	0.498***	0.750***	0.171*	0.639***
	(0.192)	(0.105)	(0.0997)	(0.153)
1997	0.506**	0.795***	0.0584	0.650***
	(0.203)	(0.114)	(0.0921)	(0.153)
1998	0.363*	0.634***	-0.0761	0.509***
	(0.215)	(0.122)	(0.0740)	(0.132)
1999	0.212	0.521***	-0.0477	0.412***
	(0.212)	(0.135)	(0.0831)	(0.136)
2000	0.205	0.531***	0.00470	0.392***
	(0.199)	(0.110)	(0.0690)	(0.136)
2001	0.134	0.485***	-0.0399	0.316**
	(0.204)	(0.112)	(0.0712)	(0.142)
2002	0.0599	0.388***	-0.0714	0.242
	(0.194)	(0.113)	(0.0805)	(0.149)
2003	-0.0428	0.334***	-0.110	0.137
	(0.199)	(0.114)	(0.0675)	(0.155)
2004	0.112	0.405***	0.123	0.294**
	(0.209)	(0.132)	(0.0969)	(0.147)
2005	-0.0520	0.265*	-0.151**	0.131
	(0.211)	(0.157)	(0.0712)	(0.160)
2006	-0.111	0.176	-0.102*	0.0691
	(0.208)	(0.123)	(0.0617)	(0.146)
2007	-0.209	0.203	-0.154**	-0.0448
	(0.210)	(0.131)	(0.0786)	(0.149)
2008	-0.159	0.271**	-0.000727	0.00778
	(0.202)	(0.115)	(0.0614)	(0.145)
2009	-0.215	0.177	-0.109	-0.0509
	(0.230)	(0.128)	(0.0895)	(0.161)
2010	-0.179	0.201*	-0.0225	-0.0150
	(0.216)	(0.122)	(0.0702)	(0.163)
2011	-0.167	0.206*	-0.0325	
	(0.196)	(0.115)	(0.0554)	

Notes: This table and Table 1.3 display different coefficients from the same regressions. Columns 1, 2 and 4 provide estimates of Equation 1.2, Column 2 from Equation 1.3. Coefficients are depicted in Figure 1.3. Stars denote statistical significance on the level of one (***), five (**) and ten (*) percent. Robust standard errors used.

Figure 1.3
Estimation with Habsburg - year fixed effects only
Habsburg coefficient plots

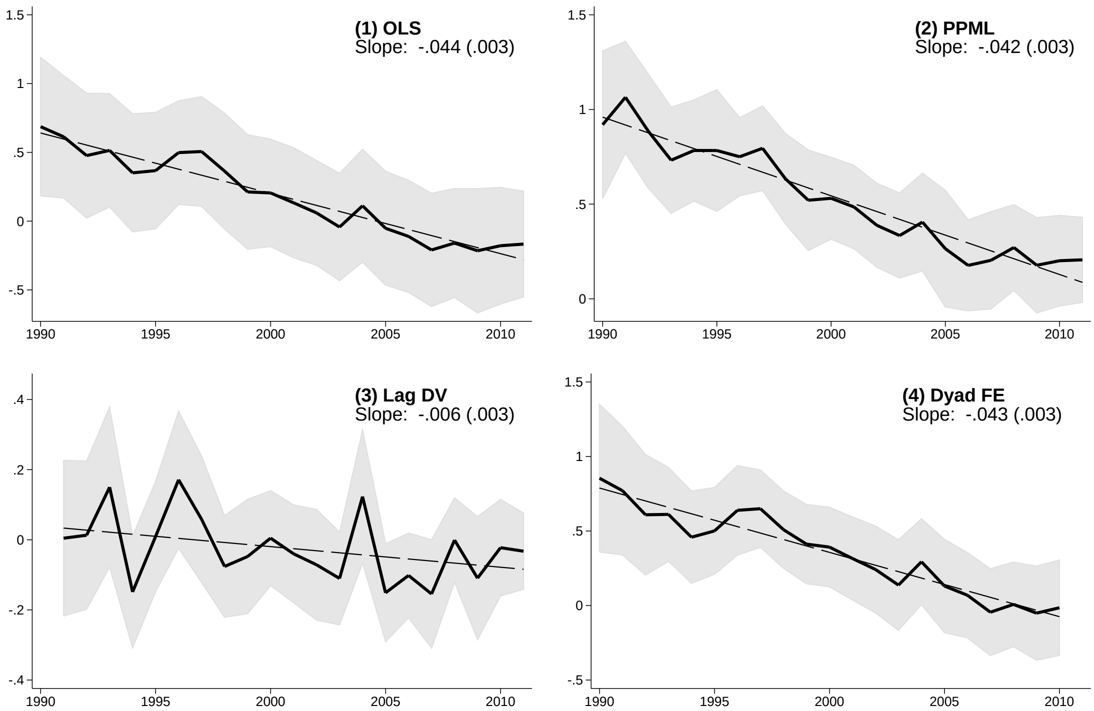

Notes: Coefficients of the Habsburg by year interaction term H_{int} in Equation 1.2 and Equation 1.3 with 95 percent confidence intervals. Line of best fit with slope and s.e. are also recorded. Restricted sample: includes only countries that were fully part of the Habsburg monarchy: Austria, Hungary and former Czechoslovakia. Coefficients of control variables are reported in Table 1.2.

additional trade in the year 1990 is 0.69, which is about three times as large as the trade bonus from two countries having a regional trade agreement (0.24), twice as large as both countries having the same religion (0.34) and 1.6 times as large as both countries being located in Eastern Europe. This magnitude also corresponds to additional trade by a factor of $e^{0.69}$, which is close to two. The surplus trade declines steadily and becomes statistically insignificant about ten years after the fall of the Iron Curtain. Note that the coefficients with Habsburg alone show stronger effects, smaller margins of error, and are more precisely estimated than the Austria-East coefficients.

Table 1.3
Estimation with Habsburg - year fixed effects only
Coefficients of control variables

	(1) OLS	(2) PPML	(3) Lag DV	(4) Bilateral FE
Dependent variable:	$\ln(x_{int})$	x_{int}	$\ln(x_{int})$	$\ln(x_{int})$
Variable of interest:				
Habsburg - year fixed effects	*– Coeff. are reported in Table 1.2 and Figure 1.3 –*			
Time fixed dyadic effects:				
Log distance	-1.181***	-0.641***	-0.213***	
	(0.0239)	(0.0113)	(0.0215)	
Common religion	0.344***	0.108***	0.0614***	
	(0.0336)	(0.108)	(0.0162)	
Both East	0.419***	0.116***	-0.0358	
	(0.0491)	(0.0455)	(0.0304)	
Shared border - year	Yes	Yes	Yes	Yes
Official common language - year	Yes	Yes	Yes	Yes
Common language spoken - year	Yes	Yes	Yes	Yes
Common legal institutions - year	Yes	Yes	Yes	Yes
Time varying dyadic effects:				
Common currency	-0.197***	0.00541	-0.00482	-0.0192
	(0.0358)	(0.0339)	(0.0188)	(0.0307)
Regional trade agreement	0.237***	0.288***	0.0576	0.344***
	(0.0560)	(0.0531)	(0.0411)	(0.0570)
Both EU	-0.0119	-0.108***	0.0175	-0.00553
	(0.0396)	(0.0319)	(0.0198)	(0.0222)
Both Euro	-0.0862***	0.271***	-0.0451***	-0.0302
	(0.0280)	(0.0311)	(0.0157)	(0.0363)
Lagged exports			0.831***	
			(0.0126)	
Origin country - year fixed effects	Yes	Yes	Yes	Yes
Destination country - year fixed effects	Yes	Yes	Yes	Yes
Bilateral fixed effects	No	No	No	Yes
Habsburg - east - year fixed effects	Yes	Yes	Yes	Yes
Observations	13,147	13,200	12,518	13,147
R-squared	0.937	0.966	0.982	0.976

Notes: This Table and Table 1.2 display different coefficients from the same regressions. Columns 1, 2 and 4 provide estimates of Equation 1.2, Column 2 from Equation 1.3. Table 1.2 shows the Habsburg × year coefficients. These coefficients are depicted in Figure 1.3. Stars denote statistical significance on the level of one (***), five (**) and ten (*) percent. Robust standard errors used.

Figure 1.4
Estimation with Austria-East - year fixed effects only
Austria-East coefficient plots

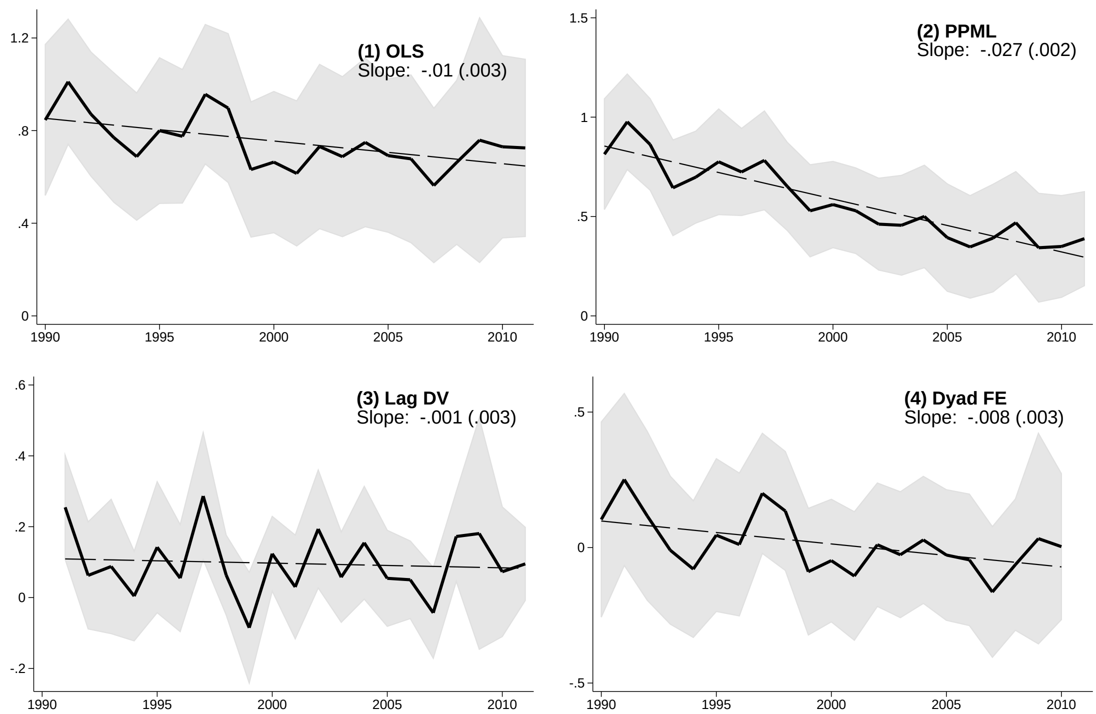

Notes: Coefficients of the $(Aus \times East) \times year$ interaction term in Equation 1.2 and Equation 1.3 with 95 percent confidence intervals. Line of best fit with slope and s.e. are also recorded.

Figure 1.4 displays the Austria-East by year interaction terms from an estimation with Austria-East coefficients. These results show a statistically significant effect in 1990 which declines linearly and monotonically in both OLS and PPML estimation techniques. The other two techniques show no significant results. Once we add controls for the Habsburg $\times$ year coefficients, this trend becomes insignificant in our preferred specification. A weak downward slope remains only in the PPML specification, statistically insignificant from zero, see Figure 1.5. These graphs suggest that Austria-East does not play a pronounced role once we control for Habsburg membership.

In Table 1.3 we proceed to estimate Equations 1.2 and 1.3 from above with only coefficients for Habsburg membership. As expected, distance negatively impacts

FIGURE 1.5
Joint estimation with Austria-East dummies and Habsburg - year fixed effects
Austria-East coefficient plots

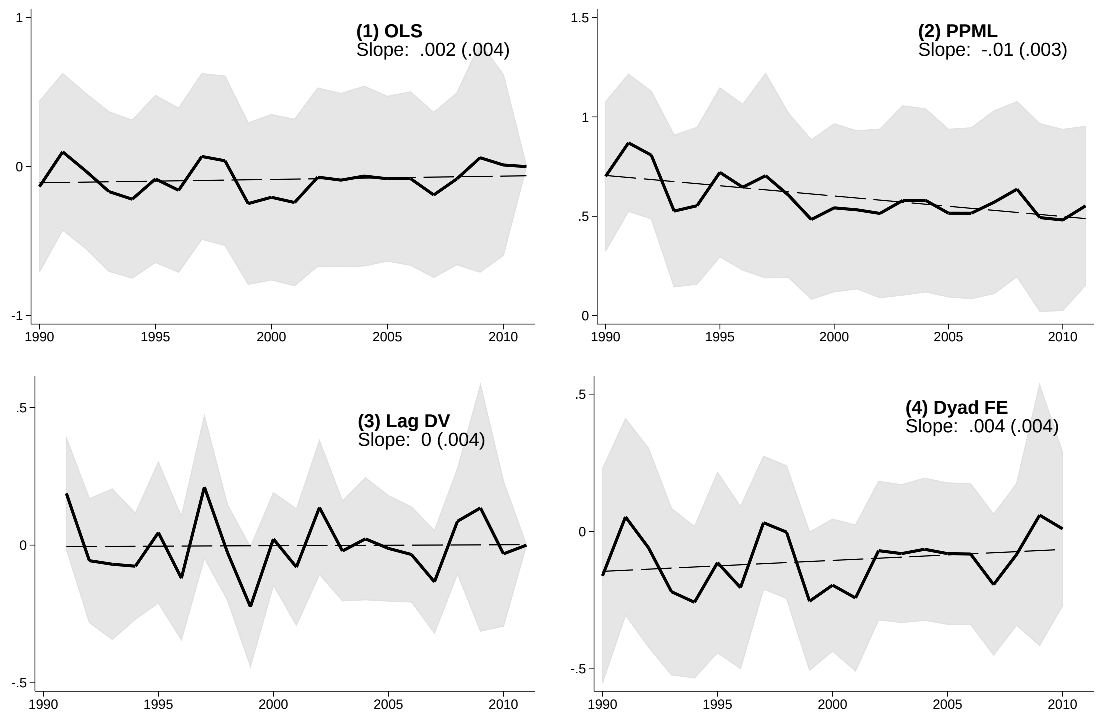

Notes: Coefficients of the $(Aus \times East)_{int}$ interaction term in Equation 1.2 and Equation 1.3 with 95 percent confidence intervals. Line of best fit with slope and s.e. are also recorded.

trade in all specifications where we can include this control variable. The displayed time varying dyadic effects tend to show the expected sign, but coefficients vary across specifications. The latter is expected as these specifications differ in many respects, for example, the PPML code is written to be estimated using levels rather than natural logarithms on the left hand side variable. Silva and Tenreyro (2006) also find a significantly smaller effect of geographical distance. Some of the coefficients show unexpected signs such as negative coefficients for common currency and 'Both EU'. This might reflect that some wealthy economies such as Norway and Switzerland are not part of EU and Eurozone. The PPML coefficient of distance exactly corresponds with that of HMR.

One concern about these results might be that the opening of the trade relations between East and West might be dynamic, increasing or decreasing, in the first years after the opening of the Iron Curtain because of various reasons other than the decline of historic and cultural ties. For example, the installation or reuse of transport infrastructure might suggest a dynamic trade relationship between an eastern and a western country or the slow establishment of personal exchange and interaction. In both these examples we would expect an increasing relationship, but there may be others. To mitigate concerns that such effects drive our results we run a placebo exercise in which we estimate 'Habsburg' effects on a relationship other than Habsburg, for which we do not expect the same decay of cultural ties. We choose Germany as the placebo country, which shares the language with Austria, and also a direct border with many eastern countries. When we estimate the trading relationship with Germany instead of Austria being the 'Habsburg' country west of the curtain, we do not find significant relationships. These results are reported in Appendix A.2, and in this table we use the same specification as applied in Tables 1.3 and 1.2. We also report results for similar placebo exercises using Switzerland, the Netherlands, Belgium-Luxembourg and Italy as alternative placebo countries, and we find no strong trend for either of these countries with the exception of a moderate decrease in Italy, which was partly Habsburg. We interpret this finding to cast doubt on the relevance of other dynamic effects shaping initial trade relationships.

Appendix A.2 demonstrates robustness of these results for different estimation strategies, additional control variables, different choices for the Habsburg definition, aggregation of countries, how to deal with missing and zero data, adding internal trade flows, and different treatment of standard errors. We find generally that this main trend is strongly robust to modifications of this type.

1.5 Product level results

In this section we shed more light on the mechanism driving our main result by studying various product categories separately. In Figure 1.6 we report the main OLS specification for each of the two-digit HS product codes except for services for which no BACI data are available. In 13 of the 15 plots the trend is downward sloping, and in 10 the downward trend is significant at 5% level of significance. The graph is upward sloping for *animal products* and *skins and leather*, both of which are small industries, accounting for 0.7 and 0.6 percent of all exports in Europe in 2000 respectively. This graph shows that our main results of a strong initial Habsburg surplus that weakens over time is not driven by a few industries, but is observable for most industry groups individually, to a varying degree however. The strongest effects in magnitude are found for machinery, foodstuff and miscellaneous. The general trend within most groups implies that industry composition changes alone cannot account for the observation of that effect.

If our results are driven by an instinct of going back to where things had been before the wars, we might expect some correlation across industries from the monarchy to trade in the 1990s. We next run our main regression separately for products traded predominantly in the monarchy and other products. Given that the product space changed considerably materially and in terms of classification over the course of these 50 years, we conduct the match on a broad level. Eddie (1989) characterizes the dual monarchy as a marriage of wheat and textiles. Good (1984) lists as main traded items in the monarchy from 1884 to 1913 food and beverages, crops, sugar, flour crops, sugar and flour originating in Hungary and industrial raw materials, textiles, machinery, and manufactured products originating in Austria. Following these classifications we classify the industries *foodstuff*, *machinery*, and *textiles* as main industries traded in the monarchy.

Figure 1.6

Main specification by HS2 product groups, in increasing numerical score: animal prods (1-5), vegetable prods (6-15), foodstuff (16-24), mineral prods (25-27), chemical industries (28-38), plastics (39-40), skins and leather (41-43), wood prods (44-49), textiles (50-63), footwear (64-67), stone and glass (68-71), metals (72-83), machinery (84-85), transportation (86-89), miscellaneous (90-97).

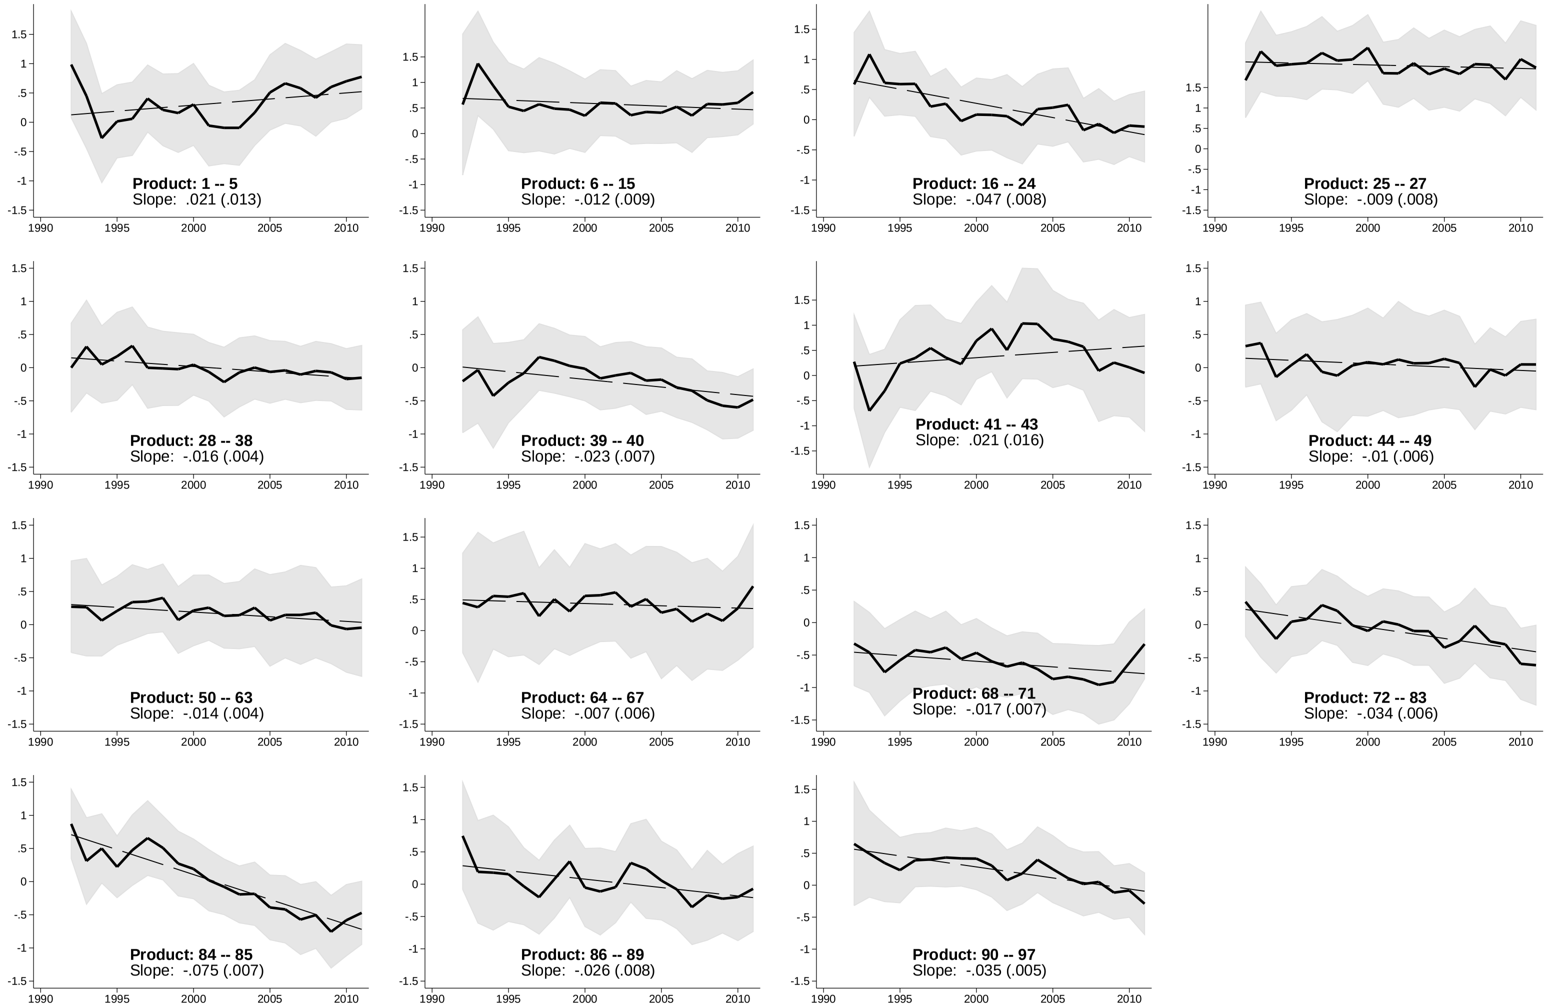

FIGURE 1.7
Main goods traded in the Monarchy and other goods

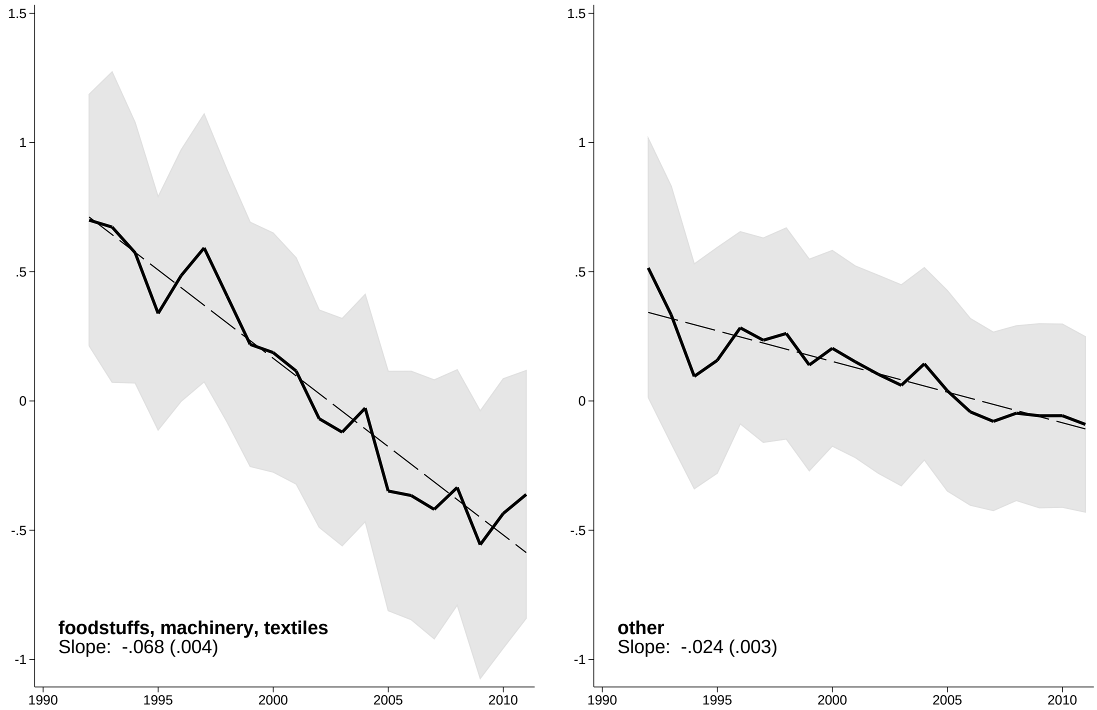

Notes: Coefficients of the Habsburg by year interaction term H_{int} in Equation 1.2 and Equation 1.3 with 95 percent confidence intervals. The left hand side panel shows the main goods traded in the monarchy, the right panel all the other goods.

In Figure 1.7 we run our main OLS regression separately for these Habsburg industries and the others. We find that both product classes display a significant, monotonic downward slope, which is not surprising given that we find the downward slope for most individual HS2 product categories. The initial trade bonus for the Habsburg traded goods is almost double that for the others, and the slope in the plot showing the Habsburg traded goods is also 2.8 times larger. The surplus trade becomes insignificant in both cases in the 2000s. This is consistent with the interpretation that cultural memory plays a part in generating the initial Habsburg surplus.

We next study the effect by heterogeneous and homogeneous products, following the standard classification by Rauch (1999). We merge the classification at the

Figure 1.8
By degree of heterogeneity and transport costs.

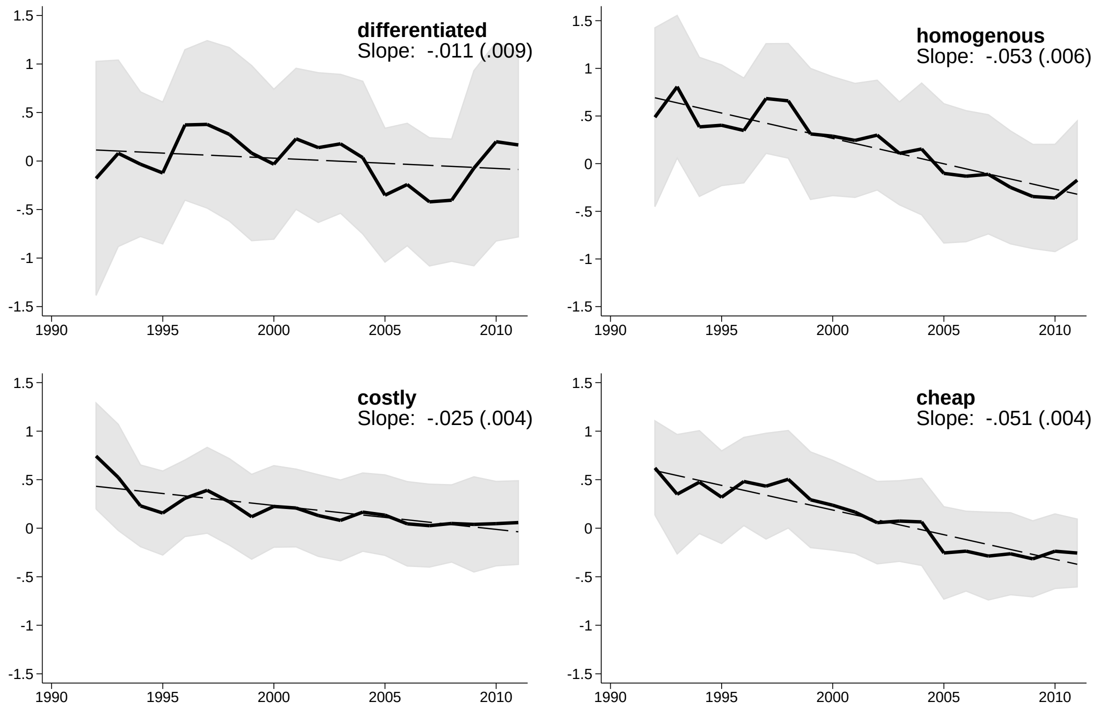

Notes: Coefficients of the Habsburg by year interaction term H_{int} in Equation 1.2 and Equation 1.3 with 95 percent confidence intervals. The top two panels compare trade of differentiated and homogeneous goods. The bottom two panels consider the different transport costs.

level of HS4, keeping only matched trade flows. These are fifteen percent of total trade flows. We think that the Habsburg bonus disappears over time as Europe adjusts to the new trading environment and converges to the new optimum. This suggests that initial deviations from the optimum, which here happen to coincide with the gravity framework, were not the first best choice. We would expect to find that the Habsburg bonuses are thus stronger for homogeneous goods, for which search costs and the costs of not using the optimum product are smaller, and thus the temptation to follow an intuitive heuristic when buying greater. As can be seen in the top panels of Figure 1.8, indeed we find the bonus is stronger initially and falls more rapidly for the homogeneous goods, while there is not such a clear pattern for the differentiated products.

If transport infrastructure surviving from the monarchy was an important driver of our findings, we should expect to see a stronger effect for goods easier to transport. To measure this effect, we obtain data on unit values from the CEPII TUV dataset. This dataset gives Free on Board (FoB) unit values per ton for each HS6 product. If, in line with some literature, we assume that the costs to ship a ton of any good are fairly similar, then inverse unit value data can serve as a proxy for transport costs as the ratio of transport costs per value transported would be smaller. Using this proxy we compare above and below median goods separately in the bottom two panels of Figure 1.8. The panel of 'costly' goods refers to above median transport cost goods while 'cheap' refers to below median ones. The standard pattern emerges, and the initial surplus trade is similar in both specifications. If there was a difference, it would be that the goods that are harder to transport adjust earlier. An explanation for this earlier drop may be that for these goods the costs of a suboptimal country to import from are higher, so adjustment may be quicker. In any case, this difference is not very strong, and coefficients rest firmly within the confidence intervals of the other graph in both cases.

1.6 Discussion of estimates

We consider a number of possible explanations why the countries of the monarchy trade more with each other in the first years after the collapse of the Iron Curtain.

First, this result might just be a consequence of a misspecification of the gravity equation. A highly structural approach of the kind we employ is easily prone to introduce noise when looking at specific bilateral trade volumes. If, for example, we overestimated the distance between Austria and the eastern countries, the residuals for these bilateral observations in a standard gravity model would be positive.[22] Or

[22] Given the location of Vienna in the east of Austria we actually underestimate the distance relative to the harmonic mean suggested in Rauch (2016).

there might be some natural geographic advantage that facilitates trade between these countries, and this reason might have brought about both the Monarchy before 1918 and the surplus trade after 1989. However, explanations and examples of this type could cast doubt on the existence of a static Habsburg surplus trade. What we observe is a trade bonus that declines linearly and monotonically over time, and it does so robustly across a number of very different estimation methods. This dynamic result is hard to explain as a simple statistical property of miss-specification or measurement error. If it was a purely mechanical specification error, our placebo exercise, that replaces Austria with Germany, would be prone to suffer from the same problem. We further verify that our main specification is robust to the use of different measures of distance, such as the distance between the most populated city and two measures of weighted distances. Our numerous robustness checks which vary estimation strategy, aggregation of countries and control variables should also help to address this concern.

Second, this difference might have to do with better existing transport infrastructure dating back to the times of the monarchy. However, most of this infrastructure was unused and laid bare during the Cold War and was derelict by 1989. The main rail lines connecting Austria with the East were abandoned; for example, in 1945 the track connecting Bratislava and Vienna, the Pressburger Bahn, the rail to the Czech Republic via Laa an der Thaya and the connection via Fratres-Slavonice were abandoned. All these lines were not revived until today. Transcontinental connections such as Vienna-Hamburg or Vienna-Berlin have switched permanently to run via Passau instead of Prague. There is also evidence that reconstruction and construction of new networks was slow after 1990, as in Hungary "there were no significant changes in the lengths of the linear transport network in the first half of the 1990s" (Erdösi, 1999). Further, even if a degenerated rail line provides a strong advantage to trade, we would not expect this surplus to contribute immediately given the time it takes to renovate such a network. Thus we would expect a slight rise of the Habsburg bonus in the first years, as this infrastructure is brought back to full capacity. In the product level section we do not find a big difference between products that are cheap or expensive to transport, which should also address this concern.

Third, this trade bonus might just reflect the specific history of bilateral developments after 1989 that are unconnected to history. Austria might have had a starting advantage, after all it was between Austria and Hungary that the Iron Curtain first opened. While it is true that the Iron Curtain was symbolically opened first between Austria and Hungary[23], things moved rapidly after that. The first symbolic opening on August 19^{th} 1989 was less than three months before the opening of borders within Germany on November 9^{th}. The first time Germans could flee was on September 10^{th} and 11^{th}. Most of the people who fled in the two months before the broader opening were East Germans. Thus the head start was neither long, nor specifically beneficial to the Austrian economy.

Fourth, it may be that language barriers are initially favourable for bilateral trade from Austria to the East, given that a higher fraction of citizens in the eastern countries still speak German than in other European countries. This explanation is similar to the interpretation we favour, however the placebo exercise using Germany and Switzerland suggest that the German language cannot explain this surplus trade and in fact does not seem to contribute to its decline.

Fifth, there could be cultural factors other than the monarchy that help to foster trust between the countries that we call Habsburg countries. It might be, for instance, that Austria's political neutrality helped to win the trust of eastern trading partners. This, however, should predict a general increase in trade for Austria with all eastern countries, rather than the selected members of the former monarchy, and would be absorbed by the interactions of Austria with all of Eastern

[23]Curiously enough in the presence of the would-have-been-emperor Otto von Habsburg.

Europe that we include. Further, we would not expect this or similar effects to decline over time, as, in spite of the monarchy, Austria's political neutrality persists.[24] The placebo exercise using Switzerland may also help to address this concern.

Sixth, there may be historical legacies and cultural forces that foster trust between these countries. For example, the surnames of the Austrian and Czech prime ministers at the first official state visit between Austria and the Czech Republic after 1990 provide a suggestive anecdote: Vranitzky is a typical Czech surname[25] while Klaus is a German first name. The cultural proximity of the Habsburg countries is also present in the Eurovision voting data by Felbermayr and Toubal (2010).[26] Historic, cultural and genetic similarities establish trust which in turn supports trade relationships. The monarchy was also the last memory of a functioning state before the wars and communism for many people in the East, and there may have been the impulse to return to what worked last when the chance appeared. This is the explanation that we favour. Why should this trade bonus deteriorate relative to other countries over time? The answer might partially be found in HMR as these factors are part of trading capital, and like other forms of trading capital they tend to deteriorate over time. In this particular case, as other countries of Western Europe establish relationships based on trust with the East, the Austrian advantage disappears as countries reorient themselves towards the new geopolitical reality. At the same time the last inhabitants on both sides of the Iron Curtain who personally remember the monarchy died in the two decades after 1990, which further may contribute to the weakening importance of the monarchy in culture.

[24]Despite joining the EU and the Euro, neutrality remains an important part of the Austrian political identity, and is a core element of its constitution and political identity.

[25]It means in Czech from the town of Vranice.

[26]In the Felbermayr and Toubal (2010) data available from Toubal's website we compute the mean Eurovision score given from country i to j and from j to i for each year and country pair. We define Habsburg as the countries in their dataset that we count as part of the monarchy in our main measure, these are Austria, Bosnia and Herzegovina, Croatia, Hungary, Yugoslavia and Slovenia. Conditional on time fixed effects these Habsburg countries have a score that is 0.048 higher than the mean of the sample, a difference that is significant at the 5% level of significance.

This explanation is consistent with our observation that the effect does not hold for the Eastern countries once we control for the Habsburg effect and is stronger in magnitude and significance for Habsburg alone than for all the countries of the east. This finding is consistent with the examples of Habsburg nostalgia mentioned at the end of the history section.

To compare these findings to HMR we conduct a few simple calculations using our estimates. HMR write that on average trade remains 31 percent higher after 60 years following their OLS specification which they obtain by exponentiating the surplus trade effect and subtracting one. Using this same methodology and the numbers provided in their paper, this implies that colonial relationships lead to a trade boost of 350 percent in the year of colonial break up. We can use our estimates directly to produce equivalent estimates. Following column 1 in Table 1.2 our corresponding numbers are surplus trade of 69 percent in year zero and 21 percent in year 10. We assume for mathematical convenience and sake of simplicity that the decay is linear. This assumption is consistent with the graphs provided by HMR, and by our own Figure 1.3, and implies a negative slope of 5.3 for the decay of trading capital, and 4.8 for the decay of the cultural part of it.[27] We can conclude that the decay of the cultural component of trading capital is 10 percent slower than the decay of all trading capital. This comparison does not require us to specify the start year of the decay.

Remarks on the estimated share of the stock of trading capital that is cultural are less precise as we do not know which year we should use as the equivalent year for colonial break up of the Habsburg monarchy. 1989 is not the end of the colonial relationship. In fact, we do not know the end we should use in our example, as we do not know if the heavy involvement of the Soviets in the East sped up cultural

[27] As an additional robustness check, we repeat our analysis including a year trend and Habsburg $\times$ year interaction term. This is a more parametric analysis compared to our main specification as it forces the slope to be linear. We find a statistically significant negative slope on the interaction term in all specifications.

memory loss, or froze it compared to a situation in the free market. Our analysis of trade flows before 1990, provided in Appendix A.2, does not suggest a decline before 1990. We can estimate the year in which the stock of cultural trading capital is exhausted, which is when the curves in Figure 1.3 become zero: around 2010. If we assume that the Soviet Union worked as a freezer of cultural capital and count the years 1918-1945 and 1990-2010 as years of decay we end up with an expected boost of 225.6 percent in year zero, compared to 350 percent implied in HMR, which would amount to 65 percent. Assuming that after the Iron Curtain fell people looked to the year before the wars and communism and that the decay was only for 20 years (1990-2010) we estimate the historical and cultural component. It amounts to 27 percent of trading capital. If we normalize the start year such that trading capital and its cultural component become zero at the same point in time, we estimate four fifths. We include this exercise as a natural comparison, but of course it is rather crude.

1.7 Conclusion

The countries of the former Austro-Hungarian monarchy trade substantially more after the fall of the Iron Curtain than a standard gravity model would predict. This initial Habsburg surplus trade is large, about four times the effect of a currency union. It deteriorates rapidly, in a monotonic and linear way, and disappears within one or two decades.

We suggest that the most likely explanation is that these forces relate to historical legacies and cultural memory parts of trading capital. These forces, established under Habsburg rule, seem to have survived over four decades and gave an initial trade boost which disappeared rapidly as countries arranged themselves with the new geopolitical circumstances. This is consistent with the following observations: (i) This surplus is found for the Habsburg countries, but not for placebo combinations of Austria-East or Eastern-Habsburg and Germany, Switzerland or the Netherlands, so it is very targeted to the area where we expect it. (ii) The effect is stronger for homogeneous goods as we would expect. Since substitution is less costly for homogeneous goods, our findings point to a preference not based on economic calculation. (iii) The effect is double for the main goods traded in the Habsburg Monarchy than for other goods which could point to some persistence of trading legacies. (iv) A number of alternative explanations, such as better infrastructure, can be ruled out. (v) This surplus trade coincides with a certain Habsburg nostalgia in the 1990s found among historians of that time.

Empires leave a lasting legacy that affects trade for decades to come. We conclude that a big part of this legacy seems to be neither physical capital nor institutional capital nor infrastructure, but is in fact some nostalgic attachment to the brand of the former empire, that we could call cultural capital.

Chapter 2

Impact of European Food Safety Border Inspections on Agri-Food Exports: Evidence from Chinese Firms

This chapter is based on joint work with Anne-Célia Disdier and Lionel Fontagné. It benefited from funding from the European Commission (EC)'s Seventh Framework Programme under Grant Agreement no. 613504 (PRONTO). The views expressed in this publication do not necessarily reflect the views of the EC. We thank Daniel Baumgarten, Matthieu Crozet, Carsten Eckel, Peter Egger, Gabriel Felbermayr, Lisandra Flach, Neil Henshaw, Franziska Hünnekes, Sébastien Jean, Anna Koukal, Chen Li, Sandra Poncet, Ferdinand Rauch, Navid Sabet, Alexander Sandkamp, Georg Schaur, Oliver Temple, Florian Unger and participants at ETSG 2015, WIIW 2015, the University of Munich seminar, the CEPII seminar, the PSE-Paris 1 Workshop on "Chinese firms in a globalized world", and the CEPR workshop "Quantifying Non-Tariff Barriers to Trade and Investment" for helpful suggestions.

2.1 Introduction

Trade liberalization has driven tariffs down. For example, the average tariff applied to Chinese agri-food exports to the European Union (EU) dropped to a low of 13 percent in 2007. However, access to the European market remains difficult because individual exporters are required to meet regulatory standards and face procedural obstacles. Non-tariff measures (NTMs) may act as substantial barriers in the decision to export because they potentially increase the cost of exporting. This problem applies particularly to agri-food products due to stringent sanitary and phyto-sanitary (SPS) regulations in most developed markets.[1] Exporters from developing countries often hold a comparative advantage in these products and tend to be over-represented in sectors heavily affected by border rejections. Consequently they are most likely to struggle with meeting stringent sanitary standards due to inadequate traceability, poor storage, limited access to certification bodies etc. (Essaji, 2008). While European standards are not designed to discriminate against imported goods, exporters in poor countries may be driven out of exporting completely.

If shipments do not comply with regulations, NTMs introduce an element of uncertainty related to possible border rejections. While the cost of meeting a standard is usually certain, there remains the risk of rejection at the importer's border.[2] The risk of rejection is determined by the variation in the quality of the exported products and the stringency of controls at the border. The former can be reduced by investment in quality or controls prior to shipment. In this paper, we are concerned with the latter: the impact of the stringency of border controls on imports. These

[1] Sanitary risk refers to food-borne human illness and animal diseases, and phyto-sanitary risk refers to risks from plant pests and transmission of diseases. In the literature sanitary measures are interchangeably referred to as health regulations or food hygiene regulations.

[2] The cost of meeting a standard is certain for exporters producing a good with their own inputs. If an exporter sources his inputs from many different suppliers, it may be difficult for him to assess the cost of meeting a standard.

are observable by the exporter but likely endogenous to past rejections which signal a high variance in the quality of the exported products. This is where externalities among exporters from the same country, region, or both emerge. Part of the cost of being rejected may be borne by competitors from the same exporting country. A spell of rejections can ultimately lead to an outright ban of a product from a particular origin. In some cases, negative externalities induced by rejections may have a product rather than a product-country dimension. However, our data suggest that such cases are rare. Most rejections have a product-country dimension and are due to production methods, climatic conditions, or both affecting a given country.

In this chapter we study the impact of the risk of rejection at the European border on Chinese agri-food exporters.[3] We find that exporters are more likely to exit the European market if the product they export has been affected by a rejection in previous years. At the same time, rejections favour the entry of new firms. Thus, border rejections increase turnover at the extensive margin of trade. Furthermore, the impact is heterogeneous across firms. Small firms are affected more strongly than big firms by this turnover. At the intensive margin, surviving firms increase their exports after a spell of border rejections. This suggests a re-allocation effect towards big and productive exporters.

This chapter's contribution to the literature is threefold. Firstly, whilst details on the occurrence of regulations gives evidence on *de jure* NTMs, knowledge about rejections sheds light on their *de facto* trade impact. Border rejections represent an example of NTMs where regulations are actually enforced. It follows that our NTM measure can be considered a *de facto* barrier for exporters.[4] Food sanitary standards have become an important policy concern in the EU making this market

[3] We do not investigate the effects of European rejections on exports to non-European markets. See Baylis et al. (2011) for an example of diversion effects for seafood products. Our research only concerns rejections from the EU as we do not have data on rejections elsewhere in the world.

[4] For additional evidence on the importance of distinguishing between *de jure* and *de facto* institutions see, for example, Acemoglu and Robinson (2006).

particularly sensitive to the issue at stake. Further, European standards are often more restrictive than international ones.

Secondly, to the best of our knowledge, this is the first study to look at the effect of SPS measures on firm-level exports from a large and significant developing economy, namely China. Since its accession to the WTO in 2001, China's impressive trade growth has accelerated further. Arguably, China is the world's most dynamic and important exporter. At the same time, anecdotal evidence suggests that Chinese agri-food exporters are struggling to meet sanitary standards.[5] Our dataset covers the universe of Chinese agri-food exports. It permits us to study the effect of rejections at the extensive and intensive margin of trade and to pay explicit attention to the role of firm heterogeneity. Theory suggests that large and more productive firms are likely to react differently to NTMs than small firms.

Thirdly, we use a rarely exploited dataset of rejections to measure the trade-impeding impact of SPS regulations at the European border.[6] The Rapid Alert System for Food and Feed (RASFF) database records all European border rejections of shipments due to sanitary concerns. Among other information, it includes the origin of the rejected shipment and a product description. We manually match the product descriptions in RASFF with HS codes at the 4-digit level of disaggregation. Although we cannot identify individual exporters that have been rejected, we merge the firm-level data with the RASFF rejections at the product and year dimension. The resulting dataset permits us to analyse the impact of border rejections on firms' export decisions.

---

[5] Frequent scandals, press articles, and anecdotes have documented the problems among Chinese exporters to meet sanitary standards. For example, German newspaper *Der Spiegel* reports: *"In recent years, China has become a major food supplier to Europe. But the low-cost goods are grown in an environment rife with pesticides and antibiotics, disproportionately cited for contamination and subject to an inspection regime full of holes."* (17/10/2012)

[6] The exception, again, is Jaud et al. (2013). However, they treat the data in a totally different way.

Enforcing sanitary standards is difficult, especially for imports from developing countries. Most agri-food imports have passed through multiple middlemen before reaching supermarket shelves. This makes it extremely difficult to trace their origins. Regulatory agencies conduct spot checks, but inspections are not random. Certain countries, firms, or products may be subject to special focus. Similarly, repeated controls are not random if custom officials expect large variations in quality from one shipment to the next. Even if one assumes an equal distribution of quality failures across countries and random inspections, shipments from large countries will be targeted more frequently by inspections if controls disregard the origin of the products. Chinese exporters thus present an interesting case study. They face considerable uncertainty concerning the probability of successful entry and costs involved in exporting. They could well be targeted by custom officials, who maximize their chance of identifying a fraudulent shipment.

The remainder of this chapter is organized as follows: Section 2.2 reviews the related literature and provides additional motivation for the research question. Section 2.3 presents the data on border rejections and Chinese firm-level exports. Section 2.4 describes the empirical strategy. Section 2.5 reports the estimation results and robustness checks. Finally, Section 2.6 concludes.

2.2 Related literature

Frontiers in research on NTMs

Non-tariff measures (NTMs) have attracted a great deal of attention in the recent trade literature. The two main issues highlighted are information sources and trade restrictiveness. Most of the research focuses on agri-food products and related sanitary measures because these are the primary drivers of safety and traceability concerns in an international trade context. Seminal contributions to this literature are Kee et al. (2009) and Disdier et al. (2008). The former measures the trade restrictiveness of NTMs by computing tariff equivalents. The latter sheds light on the magnified impact of NTMs on developing countries. All these studies face a common dilemma. On the one hand, one can use indirect evidence on border protection in a gravity equation setting. This risks capturing much more than NTMs. On the other hand, one can use direct *de jure* evidence on the presence of NTMs. This approach has the drawback that data is often outdated and incomplete.[7]

However, two much more important issues must be considered. First, not all NTMs are actual barriers to trade. This issue casts doubt on the validity of systematic assessment of their trade reducing impact. Second, not all exporters are affected equally by an NTM. This highlights the importance of studying the impact of these measures at the micro level as in Fontagné et al. (2015). Hence, this chapter combines information on rejections – a measure identified as obstacle to trade – with firm-level export data. This allows us to explore the impact of NTMs on individual exports in terms of the uncertainty introduced.

Uncertainty and export flows

Uncertainty in relation to trade costs has been addressed from two perspectives. Firstly, uncertainty is an impediment to trade from the *exporting* country perspective. Red tape or deficient infrastructures can generate uncertainty about delivery dates and the quality of the batch delivered (Nordas and Piermartini, 2004). Using data on internal transport costs of 24 sub-Saharan countries, Freund and Rocha (2011) demonstrate that uncertainty from inland transit times reduces export values. An extra day of time uncertainty reduces export values by 13 percent. Using a heterogeneous-firms model, Handley (2014) shows that trade policy uncertainty delays the entry of exporters into new markets. He argues that uncertainty about

[7] See Chen and Novy (2012) on the distinction between direct and indirect approaches.

future tariffs creates a real option value for waiting. Binding tariffs reduce such uncertainty. Osnago et al. (2015) illustrate the effect of trade policy uncertainty at the product level. A one percent reduction in the difference between bound and applied tariffs increases exports by one percent. Feng et al. (2014) measure the uneven impact of uncertainty on exporters in China. They study the US market in the years surrounding China's WTO accession and find that a reduction in tariff uncertainty induces a reallocation among Chinese exporters. Their work also indicates that firm entry and exit increases.

Secondly, uncertainty is an impediment to trade from the *importing* country's perspective. The importer cares about the quality, safety, or both of the product, which is typically unobservable. For goods traded repeatedly, reputation from a given origin may overcome the problem. For example, the consumer (importer) must be able to precisely identify the identity of the producer (exporter) (Shapiro, 1983). If the exporter's identity is unknown, the challenge to distinguish between safe and unsafe goods is more difficult. This case applies particularly to commercial relationships in international trade and with developing countries. In such cases, it is conceivable that the importer forms his expectation about the quality of a product on the exporting country's total record of quality problems. In our case, he or she obtains information from border rejections. It follows that individual exporters suffer from the problems encountered by other exporters of the same good from the same country. These information externalities can be accommodated or magnified by minimum quality standards or origin labelling (Falvey, 1989). Since information externalities are not internalized by the individual exporter, the quality provided by a large country with many firms tends to be low – a collective reputation problem. Mcquade et al. (2012) propose a theory related to these effects and argue that it fits the Chinese case well.

In this chapter, we apply our data to the issues of reputation and uncertainty raised in the literature. If the importer cannot distinguish between 'safe' and 'unsafe' trading partners, or if it is too costly to acquire the necessary information, we expect negative spillovers among Chinese exporters of the same good following a spell of border rejections.

Finally, the issue of spotting shipments which fail to comply with regulations relates to the broader literature on optimal auditing and the associated discrimination bias. The literature refers to statistical discrimination in a situation where officers target a specific group in order to maximize successful searches (Becker, 1957). For example, Knowles et al. (2001) use information on outcomes to disentangle racial prejudice from such statistical discrimination. While related to our research question, statistical discrimination is a theme with implications beyond the scope of this chapter. We do not have information on the frequency of controls, but rather solely on the incidence of rejections. Hence, we can neither assess whether Chinese firms are over-represented in controls, nor whether the rate of rejection of shipments is equal across groups.

Uncertainty component of NTM-related barriers

Somewhat surprisingly, the uncertainty component of NTM-related barriers has been mostly overlooked in the literature on NTMs and border inspections. To the best of our knowledge, there are four main papers that provide econometric investigations of the impact of import rejections on agri-food trade but none uses firm-level export data. Three of these papers deal with rejections conducted by the United States (US), while the fourth examines European rejections.

First, Baylis et al. (2009) investigate whether exporters learn from import rejections and whether these are influenced by political economy concerns. Using monthly rejections by country and product from 1998-2004, they find that new exporters are less affected by rejections than are experienced ones. This suggests that inspections are not random but are instead targeted at exporters, who have been identified as unsafe. Furthermore, rejections are not only driven by safety concerns but also by domestic political concerns.

Second, Jouanjean et al. (2015) focus on reputation. Using a sample of US rejections at the country-product dimension from 1998-2008, they highlight neighbour and sector reputation effects. If a product from a neighbouring country is rejected in the previous year, the probability of a country experiencing at least one border rejection of the same product this year more than doubles. At the sector level, the probability of a rejection increases by 62% if a related product from the same country is rejected in the preceding year.

Third, Grundke and Moser (2014) consider to what extent border rejections deter entry into the US. Using a gravity equation approach, they show that the cost of rejections at the US border falls heavily on developing countries. They use EU rejections as an instrument in part of their analysis. Grundke and Moser (2014) focus their argument on the demand for protection in the US and stricter enforcement of NTMs. Like the two papers above, they do not explicitly refer to uncertainty as a trade barrier.

Fourth, Jaud et al. (2013) study the effect of European rejections on aggregate trade flows. They document that the EU increases the number of countries it sources agri-food imports from, but that import volumes are concentrated among a small number of exporting countries. They conclude that entrants first start small. Later incumbent exporters, who have proved safe, grab most of the EU market share. Although the paper does not mention uncertainty in the import market, the mechanism they refer to is clearly linked to the mechanism studied in this chapter.

2.3 Data and descriptive statistics

2.3.1 Data

Although products subject to sanitary requirements experience systematic controls before shipment in the exporting country, controls at the border of the importing country ensure fairness of the process and retain the possibility of recognizing problems related to transportation. If a problem is identified, the shipment is likely to be rejected. We combine information on rejections of agri-food shipments at the European border with Chinese firm level export data. This allows us to measure the impact of uncertainty and regulations on firms' export decisions. We cannot identify individual exporters that have been rejected. Hence, we use the incidence of rejections as the unit of measurement of the rejection variable. Further, we merge the firm-level data with the RASFF rejections at the product and year dimension.

Food alerts and border rejections

The Rapid Alert System for Food and Feed (RASFF) consists of a cross-border information exchange system on emergency sanitary measures in the European Economic Area (EEA).[8] RASFF members must notify the European Commission about any serious health risk deriving from food or feed. Starting from its creation in 1979, all notifications are publicly available via the RASFF portal.

To construct our dataset, we record all notifications by RASFF member states over the period 1979-2011 and make several cleaning decisions:

- First, we keep notifications over the entire period 1979-2011 even if our firm-level data cover a shorter period in order to exploit the variation in notifications over time and their cumulated effect on trade flows.

[8]EEA includes the EU27 countries plus Iceland, Liechtenstein, and Norway.

Table 2.1
RASFF members

Since 1995		from 2004		from 2007
Austria	Italy	Cyprus	Lithuania	Bulgaria
Belgium	Liechtenstein*	Czech Republic	Malta	Romania
Denmark	Luxembourg	Estonia	Poland	
Finland	Netherlands	Hungary	Slovenia	
France	Norway*	Latvia	Slovakia	
Germany	Portugal			
Greece	Spain			
Iceland*	Sweden			
Ireland	United Kingdom			

Notes: * not EU, but EEA members

- Over our sample period, two rounds of RASFF membership enlargements occurred, both of which we account for. The list of RASFF members is reported in Table 2.1.[9]

- We treat the RASFF border as the relevant location for observing notifications and consider all notifications by RASFF members regarding non-RASFF countries. We ignore notifications concerning products originating from other RASFF countries.

- Since we are concerned with rejections due to SPS concerns, we restrict our analysis to agri-food products; i.e. products belonging to chapters 01-24 of the HS classification.

- Some shipments may be initially rejected but allowed entry into the RASFF market after some improvements. For example, entry may be allowed after the exporter has made changes to the product labelling. However, the majority of inspected shipments declared 'unsafe' are permanently rejected entry into the RASFF market. Using information available on the RASFF portal, we can identify whether or not entry was ultimately rejected.[10] Since we

[9]We exclude Switzerland which from 2009 is included in RASFF border controls of products of animal origin but not in other types of controls.

[10]We use information on *border rejections* for the period 2008-2011. This is reported on the RASFF portal and refers to consignments that have failed entry to the RASFF market, and are

are interested in *de facto* restrictive rejections, we retain only observations related to permanent import rejections.

- If a rejection specifies two origin countries, we split the observation into two: one for each origin.

Taking these factors into account leaves us with a total of 14,860 rejections for the period 1979-2011. Among these, 1,690 rejections are related to Chinese shipments.

The RASFF portal contains information on products only in verbal form. We code the rejection data at the HS 4-digit level – the most disaggregated level at which we can identify rejections. We provide a detailed description of the applied methodology in Appendix B.1. Using this approach, we are able to match 89 percent of all rejections with an HS4 code (13,241 out of 14,860), and 91 percent of Chinese rejections (1,537 out of 1,690).

Unfortunately, the RASFF portal neither provides the quantity or value of rejected products, nor the name of the exporting firms. Therefore, we use the incidence of rejections as measure for the rejection variable.

Chinese exports at the firm-level

Chinese customs data provide information on exports by firm, HS6 product, destination and year.[11] Our dataset covers the universe of Chinese agri-food exports for the period 2000-2011. Thus it avoids stratification or sampling issues or both; as such it is preferable to surveys often used in the literature.

Further, our dataset identifies whether the firm is a wholesaler. We use this information to restrict our attention to non-wholesalers. While intermediaries play an

---

not allowed to enter through another border post. Before 2008, this precise information on border rejections is not available. We exploit information on *notifications* and on the *action taken* by RASFF authorities to identify a border rejections. This change in rejections' identification before and after 2008 does not affect our estimation results.

[11] We thank Sandra Poncet for providing the data.

important role in trade, we want to focus on the direct decisions of firms. Intermediaries may display different export behaviour and may react less strongly to border rejections.

We aggregate all exports by firm-destination-year at the HS4 level. This corresponds to the level to which we are able to code the RASFF border rejections. It is possible that some firms might export different HS6 products within one HS4 sector. To address this concern, we verify that the large majority of HS4-firm observations also uniquely identify an HS6 shipment (see Table B.2 in Appendix B.2). Even among multi-HS4 product firms around 70 percent of HS4 sectors include only a single HS6 product.[12]

[12]Econometric estimations conducted only on firms exporting a single HS6 product within an HS4 sector do not provide results significantly different from the ones obtained with the entire sample of firms. Table B.2 shows also that the majority of firms are present in only one HS4 sector. Therefore, in our sample the impact of spillovers within firms and across HS4 sectors is likely to be small.

2.3.2 Descriptive statistics

Table 2.2 reports the number of Chinese exporters present in all world markets, and in the RASFF market. In the sample period, between 24 percent and 32 percent of Chinese exporters are present in the RASFF market. The number of active exporters rose between 2001 and 2007 and again after 2009, with a small drop in export activity during the 2008-2009 financial crisis. The sample of products exported over time is relatively stable, with a decrease after 2007. Contrary to the number of exporters, no further increase is observed at the end of the crisis. Many exporters to the RASFF market are single-product firms. On average, firms export 1.6 products to the RASFF market and the median is equal to one. Figure 2.1 plots Chinese agri-food exports over the sample period. Total exports and flows to the RASFF market are separately reported. In line with the growth in the number of exporting firms, exports tend also to increase over the period (except in 2009).

Figure 2.2 provides statistics related to RASFF rejections for all shipments regardless of origin. A significant rise in the number of RASFF rejections between 2000 and 2003 is depicted in Panel 1. This increase primarily reflects growing attention to sanitary risks and the increased application of the RASFF system by its members. The increase in rejections in 2003 is likely linked to Central and Eastern European countries harmonizing their regulations before their accession to the EU in 2004. Starting in 2003 the annual number of rejections oscillates between 1,000 and 1,500. The decreases in 2006 and 2007 are neither driven by changes in RASFF membership, nor the moving EU border. Panel 2 highlights that China, our country of interest, is among the countries most affected by RASFF rejections.[13]

[13] Turkey and Iran are ranked among the top rejected origin countries. Mycotoxins are a well known issue of Turkish exports of pistachios and dried figs, and Iranian pistachios. All Iranian exports of pistachios are double checked for freedom from mycotoxins.

Table 2.2
Chinese firms: basic descriptive statistics

		2001	2003	2005	2007	2009	2011
World agri-food exports							
Nb. of firms		7,340	8,834	12,321	12,259	11,314	11,604
Nb. of HS4 products		192	195	196	192	185	185
Agri-food exports to RASFF market							
Nb. of firms		1,800	2,083	3,176	3,604	3,548	3,730
Nb. of HS4 products		137	135	150	151	140	136
Nb. of HS4 products per firm	mean	1.68	1.57	1.64	1.68	1.61	1.59
	median	1	1	1	1	1	1

Notes: These statistics exclude wholesalers.

Figure 2.1
Chinese agri-food exports between 2000-2011 (in logs)

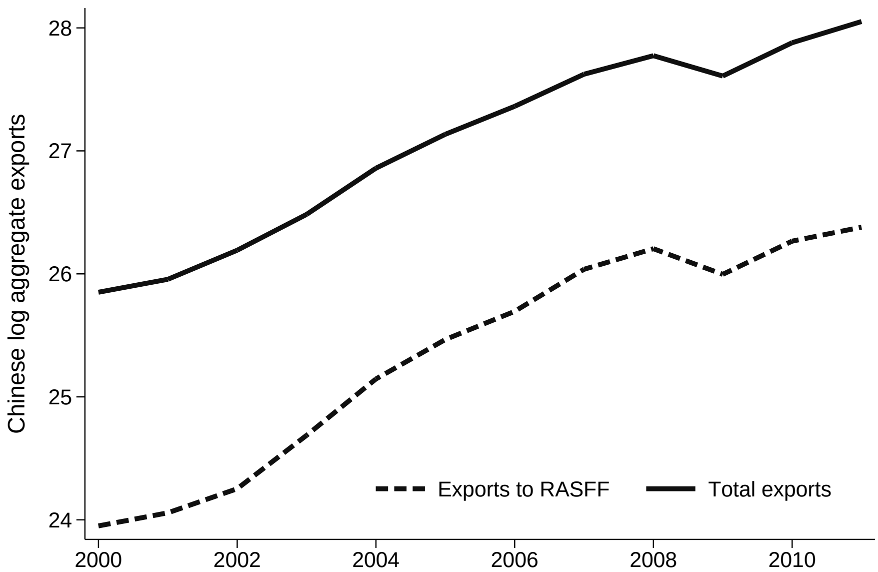

Notes: These statistics exclude wholesalers. For clarity, the statistics are presented every two years.

Figure 2.2
RASFF rejections on all shipments

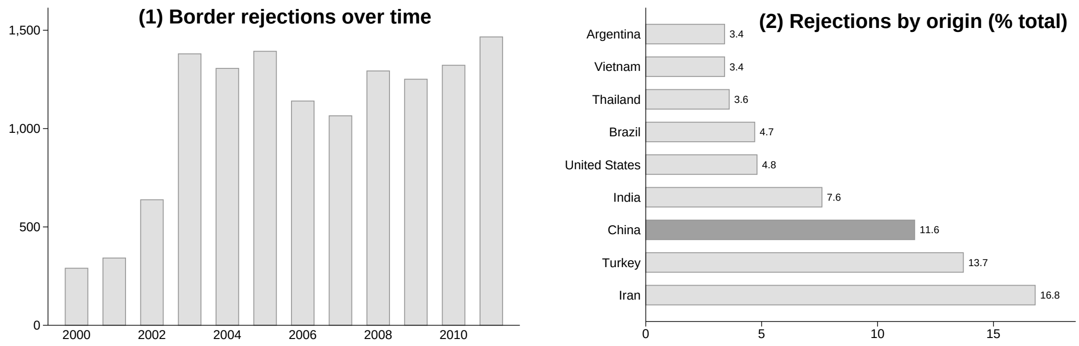

Figure 2.3
RASFF rejections on Chinese shipments

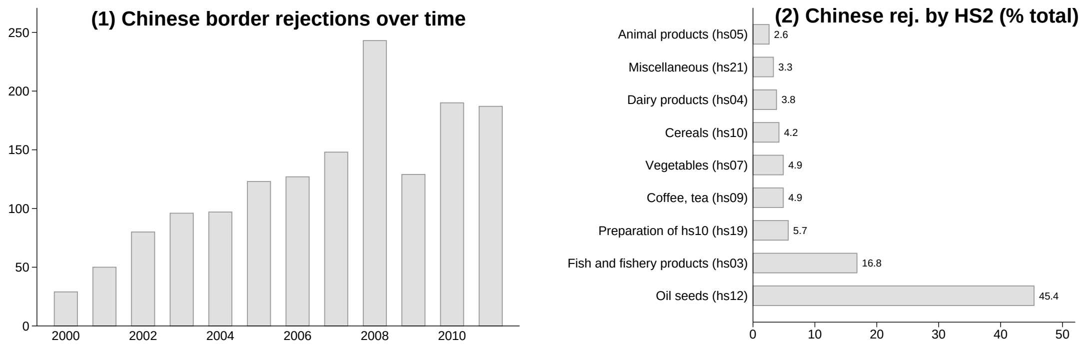

Figure 2.4
Correlation between current and lagged RASFF rejections (in logs)

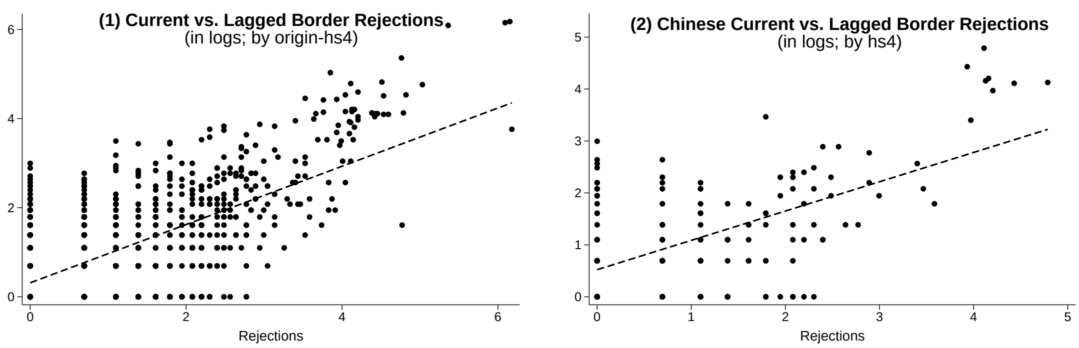

Notes: y-axis: lagged rejections (in $t-1$); x-axis: current rejections (in t).

Figure 2.3 reports the number of RASFF rejections affecting Chinese shipments (Panel 1) and the main HS2 sectors affected by rejections (Panel 2). Rejections of Chinese shipments increase over time with a dip in 2009 related to the crisis. This suggests a positive correlation between Chinese exports, depicted in Figure 2.1, and Chinese rejections at the RASFF border. In addition, we observe a strong increase in the number of rejections in 2008. This increase could be driven by a diversion of Chinese exports from countries strongly hit by the economic crisis to the EU. Firms exporting to countries with lower standards may try to export to the EU if demand in the countries with lower standards decreases. If their products do not satisfy EU requirements, this could result in an increase in rejections. An alternative explanation is related to protectionism. At the beginning of the 2008-2009 crisis, inspections and rejections were potentially used to protect European producers from Chinese competition. In our empirical analysis, we include sector-time fixed effects to control for this increasing trend. Panel 2 shows that oil seeds (HS12) and fish and fishery products (HS03) are the Chinese sectors most affected by rejections. Jointly they account for more than 60 percent of all rejections. The high share of sector HS12 in Chinese rejections relates to mycotoxin problems in peanuts (HS1202).

Figure 2.4 shows whether hysteresis is driving RASFF rejections. It compares current and lagged rejections at the same country-HS4 product dimension for all countries (Panel 1) and for China (Panel 2). It provides descriptive evidence of positive correlations. It also highlights that inspections, and therefore rejections, are not random but driven largely by past rejections. In the following section, we present more rigorous evidence to this initial analysis.

2.4 Empirical strategy

We investigate the trade impact of RASFF border rejections on Chinese firms. As discussed above, border inspections and possible rejections create some uncertainty and may impact exporters. Furthermore, this impact is likely to be heterogeneous across firms because not all shipments are inspected and inspections are not random. Certain firms or products, which present a higher safety risks, tend to attract particular scrutiny. In addition, some exporters are more able than others to invest in maintaining the quality of their products or in controls prior to shipment. Especially the biggest and most productive firms can likely afford to reduce their risk of rejections in this way.

An apparent limitation of our data is that we cannot directly identify the shipments and exporters affected by a RASFF rejection. The 'natural' model to be estimated would be:

$$\begin{aligned} y_{i,s,t} &= \alpha + \beta_{d,1} \text{ rejection}^1_{i,s,t-1} + \beta_{s,1} \text{ rejection}^2_{i,s,t-1} \\ &\quad + \mu_i + \phi_{HS2,t} + \epsilon_{i,s,t}, \end{aligned} \tag{2.1}$$

where i refers to the firm, s to the HS4-digit product category and t to the year. $rejection^1_{i,s,t-1}$ is a dummy=1 if firm i had a border rejection in product s in period $t-1$, and 0 otherwise; and $rejection^2_{i,s,t-1}$ indicates if another firm had a border rejection in period $t-1$, and 0 otherwise. Thus, $\beta_{d,1}$ would measure the *direct* effect of having a border rejection and $\beta_{s,1}$ would measure any *spillovers* to other firms. Consequently, $(\beta_{d,1} + \beta_{s,1})$ would measure the direct and spillover effects for a firm that had a border rejection and other rejections which occurred in the same HS4 category.

However, two issues raise obstacles on that natural route. Firstly, we lack the information on the affected Chinese exporter. We can only observe the overall effect of a rejection on the Chinese exports of a given product, i.e. our rejection variable is a combination of $reject^1$ and $reject^2$. Hence, our coefficient of interest measures the effect of a rejection of a particular product on *all* exporters of that product. Our estimated effect, thus, combines the direct effect of rejections as well as the indirect effect on Chinese competitors of that same product.

Secondly, although a rejected shipment is not present in EU import statistics it may be *present* in the Chinese customs data if it has passed through Chinese customs. It follows that the impact on Chinese exporters may not be observed in the current period. Hence, we choose to consider the incidence of rejections in $t-1$.

Against this background, we follow the empirical strategy suggested by Fontagné et al. (2015) and estimate the following equation:

$$\begin{aligned} y_{i,s,t} \; &= \alpha + \beta_1 \text{ rejection}_{s,t-1} \\ &+ \; \beta_2 \; ln(\text{size})_{i,t-1} \; + \; \beta_3 \; (\text{rejection}_{s,t-1} \times ln(\text{size})_{i,t-1}) \\ &+ \; \mu_i \; + \; \phi_{HS2,t} \; + \; \epsilon_{i,s,t}, \end{aligned} \tag{2.2}$$

where i refers to the firm, s to the HS4-digit product category, and t to the year.

We introduce HS2 sector-year ($\phi_{HS2,t}$) and firm (μ_i) fixed effects to control for unobserved heterogeneity. Sector-year fixed effects control for business cycles and import-demand shocks at the sector level. Industry fixed effects also capture the fact that rejections may be more frequent in industries where EU food safety standards are particularly stringent or in industries where shipments occur many times over the course of a year or both. Firm fixed effects control for time-invariant characteristics specific to a firm such as productivity or average size.

We do not cluster the standard errors. Our main variable of interest is the interaction term between rejections and firm size. This variable varies at the firm-HS4-year level, negating the need to cluster.[14]

As explained above, the RASFF border is the relevant location for our study. Since RASFF countries exchange information on rejections, a rejected product will not be able to enter the market via another RASFF border. Therefore, we do not consider export flows to each RASFF country separately, but aggregate exports to all RASFF countries. Thus, the RASFF market as a whole is the only destination in our analysis. This aggregation presents another advantage. A product could be rejected by a country which is not its final destination. However, Chinese customs data report only final destinations. This divergence between the final destination and the country of rejection could bias the results of an analysis conducted at the country level. Aggregation at the RASFF market level addresses this issue.

Considering the RASFF market as a whole does not allow us to properly control for tariff protection. However, in our analysis, this is not a major issue. All importing countries (except Iceland, Liechtenstein, and Norway) are part of the EU and apply the same common external tariffs. Therefore, tariffs are almost invariant across RASFF countries. Also, the tariffs imposed by RASFF countries on Chinese products did not vary significantly between 2000 and 2011, and a large part of any variation is captured by the set of sector-year fixed effects included in our estimations below. Therefore, the absence of a control for tariffs does not bias our results.

We define three dependent variables, $y_{i,s,t}$:

- A dummy for exit that equals 1 if the firm exports the HS4 product to the RASFF market in $t-1$ but not in t, and 0 otherwise. The counterfactual is

[14] We conduct a robustness check with clustered standard errors in section 2.5.3; the results remain unchanged.

firms that export a given HS4 to RASFF countries in $t-1$ and also in t. We disregard re-entry in later periods;[15]

- A dummy for entry that equals 1 if the firm exports the HS4 product to the RASFF market in t but not in $t-1$, and 0 otherwise. As in Javorcik (2004) and Nabokin (2014), this requires inflating the dataset, since we need to account for the counterfactual firms that could have chosen to enter, but chose not to. We inflate the dataset to include all Chinese exporters as potential entrants that at some point in the sample period export the HS4 product. Hence, the counterfactual is firms that do not enter the market; i.e. do not export a given HS4 to RASFF countries in $t-1$ or in t.

The entry and exit variables capture the extensive margin of trade at the firm-HS4 dimension. They are not analogous. As highlighted by the counterfactual, exit is conditional on the firm being active in $t-1$, while entry is conditional on not exporting in $t-1$.

- The value of the export flows for the intensive margin – specifically, the value exported by the firm to the RASFF market for a given HS4 product in year t. We focus on survivors. That are those firms that are already present in $t-1$ and continue to export in year t. In other words, we do not consider firms that start to export in year t.

Our set of explanatory variables includes border rejections and firm characteristics. We consider two different measures for border rejections. Our rejection measures ($rejection_{s,t-1}$) are:

- A dummy for past rejections that equals 1 if at least one shipment from China of that particular HS4 product was rejected at the RASFF border

[15] Recall that we focus on the RASFF market only and do not consider exports to non-RASFF countries. Therefore, a firm may exit the European market but may continue to export to non-RASFF countries.

in $t-1$, and 0 otherwise. Essaji (2008) suggests using lagged rejections as internal instruments; i.e. before actual exports in t;

- The cumulated number of past rejections from China for that HS4 product. This number is computed simply as the sum of Chinese shipments of that particular HS4 product which were rejected in the past; i.e. from 1979 until year $t-1$.

The trade literature highlights that firms' export performance is heterogeneous and driven largely by their productivity (Melitz, 2003). Unfortunately, Chinese customs data do not provide details on firms' characteristics such as productivity, employment, or total sales. Thus, to control for firm heterogeneity and its impact on export performance, we refer to firm size, defined as the natural log of their total agri-food exports in $t-1$, $\ln(\text{size})_{i,t-1}$.[16] As shown in the literature, export values are a good proxy for firm size, and big exporters are usually more efficient and more productive (Mayer and Ottaviano, 2008). For ease of interpretation, we center firm size around the median size of all firms in that year.

To capture heterogeneous effects on the impact of rejections across firms, we interact both our rejection variables (the dummy and cumulated number) with firm size.

We estimate all equations by ordinary least squares (OLS). The extensive margin dependent variables are dichotomous in nature. However, we prefer the linear probability model (LPM) to non-linear models such as logit or probit since LPM avoids the incidental parameter problem in the presence of the large number of fixed effects we employ. Besides, the LPM model provides good estimates of the partial effects on the response probability near the centre of the distribution of the explanatory variables' vector (Wooldridge, 2010).[17]

[16] Firm size is computed as $\log(1+$ total agri-food exports in $t-1)$. This approach allows us to keep brand new firms for which lagged size is equal to zero in our entry estimations.

[17] The LPM model is often used in the trade literature.[18] An alternative approach consists in using a random effects probit model. However in this model, the unobservable random variable

We conduct a series of robustness checks with wholesalers in section 2.5.3; our main conclusions remain unchanged.

2.5 Results

2.5.1 Extensive margin of trade

The first three columns in Table 2.3 present the impact of Chinese rejections on the exit of Chinese firms from the RASFF market. In columns 1 and 2, rejections are measured using a dummy that is set to 1 if at least one shipment of the same HS4 was rejected in $t-1$. We investigate exit in year t. Column 3 reports the cumulated number of past rejections of Chinese shipments for that HS4 over time until $t-1$. In all columns we control for firm size. Columns 2 and 3 also include an interaction term between firm size and past rejections. The results suggest that when we control for heterogeneity in the impact of rejections across firms, past rejections increase the probability of exit of Chinese firms from the RASFF market. According to column 3, past rejections raise the probability of exit by 4.8 percent. In addition, exit affects small and less productive firms more than bigger and more productive ones; the estimated coefficient of the interaction term between firm size and rejections is negative. We find that small firms tend to exit more, regardless of past border rejections.

Columns 4-6 in Table 2.3 report the impact of Chinese rejections on the entry of Chinese firms into the RASFF market. The estimations include the same explanatory variables as in columns 1-3. We find that rejections tend to favour the entry of new firms. The estimated coefficients on both rejection measures (dummy and cumulated number) are positive and significant. The magnitude of the effects is between 0.8 percent (column 3) and 1.1 percent (column 2) depending on the

should have a normal distribution and be independent from the observable variables, which is a strong assumption (Wooldridge, 2010).

measure used for rejections. Also, it seems that rejections promote the entry of small firms more than big firms; the estimated coefficient of the interaction term between firm size and rejections is negative and significant. Finally, regardless of past rejections, big and productive firms enter the RASFF market more easily than small ones. If we compare the estimated coefficients of the exit and entry probability, we find that past rejections have a much stronger impact on firm exit than on firm entry.

Our results are in line with Jaud et al. (2013), who find that sanitary risk increases the diversification of European imports at the extensive margin. Here, we observe turnover among Chinese firms exporting to the RASFF market. Past rejections increase both the exit of Chinese exporters and the entry of new ones. Furthermore, the effect on both exit and entry is stronger for small firms.

The last column in Table 2.3 does not examine exit and entry probabilities, but aggregates the observations at the HS4 sector-year level and instead considers the log number of Chinese firms exporting to the RASFF market. Interestingly, the estimated coefficient of the cumulated number of past rejections is negative and significant, suggesting that exit tends to dominate entry. Border rejections reduce the total number of Chinese firms exporting to the RASFF market. Also, the number of small firms shows a bigger decrease compared to big firms, and the estimated coefficient of the interaction with firm size is positive and strongly significant. As expected, the presence, in the past, of big firms in the market has a negative effect on the number of firms currently in the market.

Table 2.3
Extensive-margin estimations

	Exit from the RASFF market in year t			Entry in the RASFF market in year t			Log number of firms
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dummy = 1 if at least one rejection in $t-1$	-0.026***	0.124***		0.006**	0.011***		
	(0.007)	(0.031)		(0.003)	(0.003)		
Dummy for rejection in $t-1$ × Firm size		-0.012***			-0.002***		
		(0.002)			(0.0004)		
Cumulated nb. of past rejections until $t-1$			0.048***			0.008***	-0.138*
			(0.012)			(0.001)	(0.072)
Cum. nb. past rejections × Firm size			-0.005***			-0.001***	0.017***
			(0.001)			(0.0001)	(0.006)
Firm size	-0.0047***	-0.043***	-0.041***	0.013***	0.013***	0.014***	-0.010**
	(0.002)	(0.002)	(0.002)	(0.0002)	(0.0002)	(0.0003)	(0.005)
Observations	49220	49220	49220	178951	178951	178951	1542
R^2	0.391	0.391	0.392	0.081	0.081	0.082	0.951

Note: Fixed effects for firms and HS2-year in columns 1-6 and for HS4 sectors and HS2-year in column 7 (not reported). Robust s.e. in parentheses. Firm size defined as log of firm's total agri-food exports in $t-1$. ***/**/* indicate significance at the 1%/5%/10% level.

2.5.2 Intensive margin of trade

Table 2.4 reports results on the intensive margin of trade. Columns 1-3 focus on the value of exports to the RASFF market by survivors, that is firms present in years $t-1$ and t. Our results highlight three main facts. First, and independent of border rejections, bigger firms tend to survive and increase their exports to the RASFF market. Second, on average firms that continue exporting products hit by rejections neither increase nor decrease their exports to the RASFF market. The two variables, the dummy and cumulated number of past rejections, have no significant impact on the export values in columns 2 and 3. Third, some heterogeneity is observable across firms, and the results for the interaction terms between past rejections and firm size suggest that big and more productive incumbent firms increase their exports to the RASFF market in the year(s) following a rejection. Therefore, large firms do benefit from the exit of small exporters consecutive with a rejection.

Column 4 investigates the impact of border rejections on the quantity exported by incumbents, while column 5 examines the price – measured as the unit value – of the products exported by these firms.[19] The heterogeneous effect of past rejections across firms remains positive but is less significant ($p < 0.05$ for quantity and $p < 0.10$ for price). In terms of magnitude, the effect on price is smaller than the effect on quantity. Finally, regardless of past rejections, firm size has no impact on price. These results suggest that big and productive incumbent firms increase the quantity exported to the RASFF market, and perhaps to a lesser extent, the product price.

Our results at the intensive margin show a concentration of Chinese exports among big and productive exporters. The effect is stronger for products hit by past re-

[19] Some prices exhibit extreme values. We exclude these outliers by deleting the top and bottom 1 percent of the price observations.

TABLE 2.4
Intensive-margin estimations

	Ln exports to the RASFF market in t (Surviving firms)				
	Value			Quantity	Unit value
	(1)	(2)	(3)	(4)	(5)
Dummy = 1 if at least one rejection in $t-1$	0.269***	-0.193			
	(0.031)	(0.165)			
Dummy for rejection in $t-1$ × Firm size		0.035***			
		(0.012)			
Cumulated nb. of past rejections until $t-1$			-0.010	-0.007	-0.002
			(0.067)	(0.070)	(0.027)
Cumulated nb. of past rejections × Firm size			0.014***	0.010**	0.003*
			(0.005)	(0.005)	(0.002)
Firm size	0.163***	0.151***	0.149***	0.141***	0.008*
	(0.010)	(0.011)	(0.012)	(0.012)	(0.005)
Observations	30999	30999	30999	30982	30486
R^2	0.623	0.623	0.625	0.651	0.788

Note: Fixed effects for firms and HS2-year in all estimations (not reported). Robust standard errors in parentheses. Firm size is defined as the log of firm's total agri-food exports in $t-1$. ***/**/* indicate significance at the 1%/5%/10% level.

jections. As above these results confirm Jaud et al. (2013), who also highlight concentration at the intensive margin, especially for risky products. When rejections are more frequent and cumulate, European importers concentrate their orders on large, and plausibly more reliable Chinese exporters, who increase their exports to the RASFF market.

Thus, we observe two effects: turnover of firms at the extensive margin of trade accompanied by a concentration at the intensive margin.

2.5.3 Robustness checks

In this section, we investigate the robustness of our results to alternative specifications and samples. We perform all the tests using our preferred estimations, i.e. those including the cumulated number of past rejections as a measure of border rejections and the interaction term between this rejection measure and firm size. We run three estimations in each case: one for the probability that the Chinese firms will exit the RASFF market, one for the probability of entry into that market and one for the intensive margin of trade.

First, we test whether our results change if the standard errors are clustered. As mentioned in Section 2.4, clustering is not mandatory in our case because our variable of interest, the interaction term between rejections and firm size, varies at the firm-HS4-year level. However as a robustness check, columns 1-3 in Table 2.5 include clusters defined at the HS4-year level. The results are not affected by their inclusion.

A second source of potential bias relates to churning flows and potential reverse causality. To check for this, we introduce in the estimation a measure of the mean length of HS4 flows exported to the RASFF market. We report the results in columns 4-6 of Table 2.5. This variable has a significant influence on both the extensive and intensive trade margins but its inclusion does not affect our previous conclusions.

Endogeneity may stem also from our focus on Chinese rejections and Chinese firms' exports. Potential bias is reduced by the use of lagged rejections. In addition, we replicate our main estimations using two alternative sets of rejections: (i) non-Chinese rejections, (ii) all rejections whatever the product origin; i.e. Chinese and non-Chinese. Table 2.6 reports the results. For the extensive margin of trade, the magnitude of the estimated coefficients is lower, but they have the same sign and level of significance as in Table 2.3. This suggests that rejections related

Table 2.5
Robustness: clustering and churning flows

	With clusters			Churning flows		
	Exit (1)	Entry (2)	IM (3)	Exit (4)	Entry (5)	IM (6)
Cum. nb. past rej. until $t-1$	0.048***	0.008***	-0.010	0.065***	0.007***	-0.107
	(0.013)	(0.001)	(0.085)	(0.012)	(0.001)	(0.066)
Cum. nb. past rej. × Firm size	-0.005***	-0.001***	0.014**	-0.006***	-0.001***	0.016***
	(0.001)	(0.0001)	(0.006)	(0.001)	(0.0001)	(0.005)
Firm size	-0.041***	0.014***	0.149***	-0.041***	0.014***	0.149***
	(0.002)	(0.0003)	(0.013)	(0.002)	(0.0003)	(0.012)
Mean length of flows				-0.271***	0.042***	1.566***
				(0.009)	(0.004)	(0.048)
Observations	49220	178951	30999	49220	178951	30999
R^2	0.392	0.082	0.625	0.405	0.042	0.640

Note: Fixed effects for firms and HS2-year in all estimations (not reported). Robust standard errors in parentheses. Firm size is defined as the log of firm's total agri-food exports in $t-1$. Columns 1-3: Standard errors in parentheses, clustered at HS4-year level. Columns 4-6: Regressions also include the mean length of flows. ***/**/* indicate significance at the 1%/5%/10% level.

Table 2.6
Robustness: Non-Chinese and entire sample of rejections

	Non-Chinese rejections			All rejections		
	Exit (1)	Entry (2)	IM (3)	Exit (4)	Entry (5)	IM (6)
Cum. nb. past rej. until $t-1$	0.024***	0.005***	0.116**	0.024***	0.005***	0.117**
	(0.008)	(0.001)	(0.049)	(0.008)	(0.001)	(0.047)
Cum. nb. past rej. × Firm size	-0.002***	-0.001***	-0.004	-0.002***	-0.001***	-0.004
	(0.001)	(0.0001)	(0.003)	(0.001)	(0.0001)	(0.003)
Firm size	-0.042***	0.014***	0.171***	-0.042***	0.015***	0.171***
	(0.002)	(0.0003)	(0.013)	(0.002)	(0.0003)	(0.013)
Observations	49220	178951	30999	49220	178951	30999
R^2	0.391	0.082	0.623	0.391	0.082	0.623

Note: Fixed effects for firms and HS2-year in all estimations (not reported). Robust standard errors in parentheses. Firm size is defined as the log of firm's total agri-food exports in $t-1$. Columns 1-3: Non-Chinese rejections. Columns 4-6: All rejections whatever the origin of the products. ***/**/* indicate significance at the 1%/5%/10% level.

to products imported into Europe from non-Chinese suppliers but also exported by Chinese firms shape the participation of Chinese firms. Rejections for a given product category increase the probability of additional controls on similar products from all origins. This restrains Chinese export participation, although the impact is weaker compared to the effect of rejections of Chinese products. Chinese exporters fear tighter controls on the type of products they export, even if these controls do not necessarily target their own flows. Results at the intensive margin of trade first confirm the expected market shares redistribution. Chinese firms substitute at least partially for competitors following rejection of *non-Chinese* products. The estimated coefficient of the cumulated number of past rejections becomes significant at the intensive margin. Second, we cannot reject the hypothesis that Chinese firms benefit equally, whatever their size, from this redistribution of market shares. The estimated coefficient of the interaction terms is not significant at the intensive margin, so the heterogeneous effect of past rejections disappears.

A potential issue raised by the previous estimations is the sensitivity of exported products to control. For instance, among oil seeds certain product categories (e.g. peanuts) are highly sensitive to mycotoxins and should be more often subjected to control at the RASFF borders. This outcome is not captured by our previous set of fixed effects. We control for the time-invariant characteristics of products by introducing HS4 fixed effects in addition to the HS2 fixed effects. This strategy allows us to disentangle the product-country vs. product-only dimensions related to inspections. The first three columns in Table 2.7 control for these unobservable product characteristics. At the extensive margin of trade, our previous results remain unchanged. At the intensive margin, however, the estimated coefficient of the border rejection variable becomes negative and significant. This suggests that firms export fewer HS4 products hit by rejections if we control for unobserved characteristics. Even for big firms the effect is negative; the sum of the coefficients of the cumulated number of past rejections and of the interaction term is negative.

Table 2.7
Robustness: HS4 unobservable characteristics and trade flows intensity

	HS4 characteristics			Trade flows intensity		
	Exit (1)	Entry (2)	IM (3)	Exit (4)	Entry (5)	IM (6)
Cum. nb. past rej. until $t-1$	0.078*** (0.013)	0.004** (0.002)	-0.159** (0.070)	0.254*** (0.082)	0.046*** (0.008)	0.077 (0.459)
Cum. nb. past rej. × Firm size	-0.005*** (0.001)	-0.001*** (0.0001)	0.013*** (0.005)	-0.025*** (0.006)	-0.007*** (0.001)	0.057** (0.025)
Firm size	-0.042*** (0.002)	0.014*** (0.0003)	0.152*** (0.011)	-0.043*** (0.002)	0.014*** (0.0003)	0.154*** (0.012)
Observations	49220	178951	30999	49169	177252	30987
R^2	0.409	0.083	0.655	0.391	0.082	0.624

Note: Fixed effects for firms and HS2-year in all estimations (not reported). Robust standard errors in parentheses. Firm size is defined as the log of firm's total agri-food exports in $t-1$. Columns 1-3: Regressions also include HS4 fixed effects (not reported). Columns 4-6: Cumulated number of past rejections weighted by the cumulated number of past export flows. ***/**/* indicate significance at the 1%/5%/10% level.

Our results suggest also that the negative externalities induced by border rejections have primarily a product-country rather than a product-only dimension. In other words, Chinese exporters are affected negatively by rejections which hit the same Chinese product as the one they export. This is in line with the results in columns 3 and 6 in Table 2.6. At the intensive margin, Chinese exporters seem to benefit from rejections affecting *non-Chinese* products but are negatively affected by rejections targeting *Chinese* products.

The last three columns in Table 2.7 account for the intensity of Chinese export flows to the RASFF market for each HS4 sector. The number of rejections of Chinese shipments varies across sectors (see Panel 2 of Figure 2.3). Part of this variation is due to the sanitary risk which of course may differ across products, but partly originates from the intensity of trade between China and RASFF countries. A sector characterized by many flows is likely – all else being equal – to encounter a higher number of rejections. To control for the intensity of trade, different weighting schemes can be used. For example, rejections could be weighed by trade volume or the number of export flows. Here, we weigh the cumulated number of past rejections by the cumulated number of past export flows. We compute this as the sum of the cumulated number of export flows by Chinese firms to RASFF countries within one HS4 sector over time. The results confirm, and even strengthen, our previous findings since the magnitude of estimated coefficients is larger than those reported in Tables 2.3 and 2.4.

Next, we test whether our results are sensitive to the sample of firms considered in the estimations.[20] First, we exclude exporters that only serve the RASFF market for a short period. To do so, we compute the number of years of presence of each Chinese firm exporting to the RASFF market. We restrict our sample to firms where the number of years of presence is above the median. The first three columns in Table 2.8 present the results of these estimations. The sample restriction has no impact on our previous conclusions. The three last columns in Table 2.8 add wholesalers to the sample of firms. So far, our analysis has excluded wholesalers in order to examine active firm export decisions. However, wholesalers represent a non-negligible number of Chinese exporters. In fact, their inclusion in the sample has almost no impact on the estimated coefficients, and the previous findings remain valid.

Firms exporting to other non-EU OECD markets may be more successful in passing RASFF inspections.[21] These markets also impose stringent safety regulations, and conduct inspections. Therefore, exporters to these markets are more likely to sell safe products and to have higher productivity. This may help them deal with inspections, their related costs, and uncertainty. Table 2.9 distinguishes between firms exporting to at least one OECD market outside the RASFF market in $t-1$ vs. other firms. We investigates whether rejections have different trade effects on

[20] Unfortunately, information on ownership is missing for many firms. Therefore, we cannot test whether rejections have a differentiated impact on foreign, private, and state-owned firms.

[21] Non-EU OECD countries are Australia, Canada, Japan, New Zealand, Switzerland, South Korea, and the US.

Table 2.8
Robustness: firms' number of years of presence and wholesalers

	Above median nb. year of presence			With wholesalers		
	Exit	Entry	IM	Exit	Entry	IM
	(1)	(2)	(3)	(4)	(5)	(6)
Cum. nb. past rej. until $t-1$	0.050***	0.009***	-0.010	0.043***	0.008***	-0.039
	(0.013)	(0.001)	(0.067)	(0.009)	(0.001)	(0.054)
Cum. nb. past rej. × Firm size	-0.005***	-0.001***	0.014***	-0.004***	-0.002***	0.016***
	(0.001)	(0.0001)	(0.005)	(0.001)	(0.0001)	(0.004)
Firm size	-0.041***	0.014***	0.149***	-0.037***	0.015***	0.128***
	(0.002)	(0.0003)	(0.012)	(0.002)	(0.0002)	(0.009)
Observations	45469	133977	30999	88858	352192	51998
R^2	0.293	0.094	0.523	0.363	0.067	0.595

Note: Fixed effects for firms and HS2-year in all estimations (not reported). Robust standard errors in parentheses. Firm size is defined as the log of firm's total agri-food exports in $t-1$. Columns 1-3: Firms with a number of years of presence above the median. Columns 4-6: With wholesalers. ***/**/* indicate significance at the 1%/5%/10% level.

Table 2.9
Robustness: OECD presence in $t-1$

	Exit		Entry		IM	
	No	Yes	No	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)
Cum. nb. past rej. until $t-1$	0.089***	0.030**	0.004***	0.006**	-0.150	0.001
	(0.030)	(0.014)	(0.001)	(0.003)	(0.149)	(0.079)
Cum. nb. past rej. × Firm size	-0.008***	-0.004***	-0.001***	-0.001**	0.024**	0.013**
	(0.002)	(0.001)	(0.0003)	(0.0002)	(0.011)	(0.006)
Firm size	-0.052***	-0.036***	0.013***	0.006***	0.089***	0.153***
	(0.006)	(0.003)	(0.001)	(0.001)	(0.029)	(0.014)
Observations	11528	37692	126216	52735	6213	24786
R^2	0.603	0.366	0.118	0.187	0.796	0.604

Note: Fixed effects for firms and HS2-year in all estimations (not reported). Robust standard errors in parentheses. Firm size is defined as the log of firm's total agri-food exports in $t-1$. Columns 1, 3, and 5: Firms not exporting to at least one OECD market (other than the RASFF market) in $t-1$. Columns 2, 4, and 6: Firms exporting to at least one OECD market (other than the RASFF market) in $t-1$. ***/**/* indicate significance at the 1%/5%/10% level.

these two groups of firms. First, we observe that our previous conclusions, namely the diversification at the extensive margin and concentration at the intensive one, are accurate for both groups of firms. However, there are some differences in the magnitude of the estimated coefficients. Firms already exporting to another OECD market in $t-1$ are less likely to exit from the RASFF market due to border rejections. In addition, this effect is magnified for big and productive firms. Columns 3 and 4 indicate that entry to the RASFF market induced by rejections is slightly stronger for firms already exporting to at least one other OECD market. At the intensive trade margin, productive incumbent firms exporting to OECD markets in $t-1$ are also more likely to increase their exports to the RASFF market in t compared to other firms (columns 5 and 6).

2.6 Conclusion

In this chapter, we study whether a rise in uncertainty related to the risk of border rejections affects imports from a large developing economy. NTMs may act as substantial barriers in the decision to export because they potentially increase the cost of exporting. If border rejections result in an increased likelihood of inspection, a series of import rejections could induce negative spillovers for competitors from the same origin, the same product or both.

Our results show that Chinese exporters of agri-food products are more likely to exit the European market if the product they export has been rejected in previous years. At the same time, rejections favour the entry of new firms. This highlights turnover effects at the extensive margin of trade. At the intensive margin, border rejections increase the exports of surviving firms; i.e. a re-allocation effect. Furthermore, the microeconomic impact of the risk of rejection is heterogeneous across firms. Turnover at the extensive margin mainly concerns small firms, while concentration at the intensive margin benefits big firms more. Overall, the number of exporting firms tends to decrease, but total exports of the surviving firms increases. Our results confirm the key role played by uncertainty, and that big and more productive firms are more resilient than small ones to the risk of border rejections.

Our results contribute to the literature on firm heterogeneity and trade. We provide a more nuanced understanding of the impact of *de facto* restrictive regulations on exporting firms. Furthermore, given the importance of food safety and importers' emphasis on sourcing from reliable producers, our results suggest that policy makers and law enforcers should adopt a comprehensive approach and pay attention to individual firms rather than focusing on entire sectors.

Chapter 3

Striking Evidence? Demand Persistence for Inter-City Buses from German Railway Strikes

Many thanks to MeinFernbus Flixbus GmbH (especially Tina Bosler and Anna Humpert) for providing me with the necessary data. This chapter has been greatly improved by the comments and suggestions of numerous colleagues and seminar participants including but not limited to Daniel Baumgarten, Maximilian Breu, Carsten Crede, Fabian Dehos, Carsten Eckel, Florian Englmaier, Gabriel Felbermayr, Lisandra Flach, Lionel Fontagné, Lucas Goodman, Neil Henshaw, Julian Hinz, Anna Koukal, Chen Li, Lukas Mergele, Takeshi Murooka, Ferdinand Rauch, Alexander Ross, Alexander Sandkamp, Klaus Schmidt, Monika Schnitzer, Daniel Wissmann as well as seminar participants from the MGSE Colloquium and IO and Trade Seminar in Munich and Paris, the RGS doctoral conference in Bochum, and doctoral seminar in Bonn-Mehlem. The MeinFernbus Flixbus Team, Martin Rammensee of Busliniensuche.de, Andreas Oswald of Fernbusse.de, and Rüdiger Knobel of Simplex Mobility provided important doses of reality and data support. I thank Oscar Violante and Fabian Gräf for excellent data assistance.

3.1 Introduction

3.1.1 Motivation and Outline

Does a temporary shock such as a strike boost one's competitors' demand? Can such shocks have lasting demand effects? I analyse these general questions in the context of the German railway strikes of 2014-2015. The strikes forced travellers to use alternative transport modes. For many travellers this was their first encounter with inter-city buses – a newly liberalized market.[1] Such a shock – in introducing new customers to the railway's key rival – has the potential to result in new, long-term customers for buses who otherwise would have routinely stayed with rail.[2] A German newspaper article suggested that *"the young bus market could benefit sustainably from the strike"*.[3] To the best of my knowledge, this study is the first to present systematic evidence of these qualitative accounts.

This chapter combines three novel and extremely rich datasets: detailed booking data provided by Germany's largest bus provider MeinFernbus (MFB), emergency timetables published by Deutsche Bahn (German Rail; hereafter referred to as DB) during the strikes, and a web-crawled dataset of all rail itineraries. Using this data, I study the adjustments of travellers to inter-city buses during the strike, and test for demand persistence. The German railway strikes of 2014-2015 provide several desirable features for a quasi-natural experiment setting. Competition from buses

[1] The market was liberalized by law as of January 2013. Previously the Passenger Transportation Act only permitted inter-city bus services if the state-owned railway company was unable to provide an acceptable service. Dürr et al. (2015) provide more details on the liberalization.

[2] Inter-city buses are defined as regularly scheduled services exceeding a distance of 50km. In the literature they are often interchangeably referred to as 'inter-urban' or 'long-distance' buses.

[3] Full relevant excerpt: *"(...) The young bus market could benefit sustainably from the strike. (...) Due to the strikes business travellers are compelled to try the bus and then use it again (...) The number of repeat bookings climbs."* (url: *http://www.faz.net/aktuell/wirtschaft/wer-vom-bahnstreik-profitiert-mietwagen-und-fernbusse-13603674.html*; 20/05/2015) Other anecdotal evidence is provided by Spiegel magazine who suggested that *"(...) the structural change will accelerate in the German domestic inter-city market."* (url: *http://www.spiegel.de/wirtschaft/service/bahn-streik-fernbus-unternehmen-profitieren-von-gdl-ausstand-a-1001003.html*; 05/11/2014)

played no role in the exposure across routes, the occurrence, or the timing of the railway strikes. Furthermore, this was the first German railway strike in which buses – a viable alternative – were available.

My empirical strategy consists of two steps. Firstly, I test which routes were primarily affected *during* the rail strike. While the exposure of *rail routes* to the strike can be deduced from the emergency timetables, the exposure of *bus routes* is not ex-ante clear to the researcher. On the one hand, it is not clear how well travellers were informed about variations in DB service cancellations across routes. On the other hand, travellers may only switch if the bus service is a close enough substitute to rail. I demonstrate that the only channel driving MFB ticket sales during the strikes is the closeness of substitution, measured by the bus travel time. Travellers switched to buses even on routes with little or no rail service cancellations. This suggests that they were not well informed about their exposure to the rail strike or had no trust in DB's ability to implement the emergency timetables. I show that the effect of the rail strikes was largest on routes with a short absolute bus travel time.

Secondly, I estimate the effects of the strikes on ticket sales *after* DB operations returned to normal; i.e. whether there was a persistent effect. In a difference-in-differences setting, I use the channel identified in the first step to define treatment and control group. More precisely, I compare the change in the number of customers between high (short bus travel time) and low (long bus travel time) strike-exposed routes. Although the common trend assumption does not seem to be completely tenable in the given context, my results point to a persistent effect on the ticket sales for inter-city buses on the affected routes. I follow the methodology of Nunn and Qian (2011), who employ a similar strategy in a different setting.[4] They estimate period-specific treatment effects for the *pre-period* in order to com-

[4] Nunn and Qian (2011) study the impact of the introduction of the potato from the Americas on Old World population growth and urbanization.

pare these to the post-treatment coefficients. Following their methodology, my results also remain largely unaltered to a number of alternative specifications and robustness checks.

This chapter proceeds as follows: The remainder of this section reviews the related literature and discusses several features of the railway strikes in 2014-2015. Section 3.2 introduces the datasets and provides new descriptive statistics on the inter-city bus market. Section 3.3 introduces potential transmission channels and tests which bus routes were most affected *during* the rail strike. Section 3.4 uses the results from the previous section to test for demand persistence *after* the strike. Section 3.5.2 reports robustness tests. Finally, Section 3.6 concludes.

3.1.2 Related Literature

The literature on the subject of rail strikes and their effects on traveller behaviour is surprisingly sparse. Bauernschuster et al. (2015) and Van Exel and Rietveld (2001) provide overviews. Often inference relies on survey data and the effect of the strike is studied retrospectively. While strikes occur on a regular basis, they are not easily anticipated and may not last long enough to formulate an appropriate research design. To the best of my knowledge, the only notable exception is Larcom et al. (2015), whose contribution is closely related to this chapter. They show that the 2014 London underground strike resulted in about 5 percent of commuters permanently changing their commuting route. They suggest that individuals under-experiment in normal times. Public transport strikes are often considered to be highly economically damaging (Kennan, 1986). Larcom et al. (2015) and this chapter highlight an unintended and potentially positive channel, which is often overlooked in the literature: if the rail strike revealed information, it may have been welfare improving. Some customers, who were forced to experiment with buses, discovered that their previous choice was not optimal.

This chapter contributes to this strand of the literature in two ways. Firstly, I study inter-modal switching across transport modes for inter-city transport – a less-frequent travel decision than daily commuting. The frequency of the travel decision might matter, as suggested by the behavioural economics literature on salience (Chetty et al., 2007). Secondly, in comparison to Larcom et al. (2015) the longer post-strike period allows me to better understand the short- and medium-term impacts of any effect.

This chapter supplements the classic literature relating to the way in which individuals decide between alternatives. There is a large and long-standing debate on rational decision-making (Weitzman, 1979; Morgan and Manning, 1985) and constraints such as search costs (Baumol and Quandt, 1964) or information asymmetries. My results cannot be reconciled with the classical economic assumption of perfectly informed and rational consumers. After all, bus services were available before the strikes and the availability of internet bookings – the primary booking channel – remedy some of the search costs. Porter (1991) argues that exogenous shocks may help individuals find their optimal choice by triggering a period of experimentation. The underlying idea of experimentation due to exogenously-imposed constraints, such as the non-availability of rail services, applies to the setting in this chapter.

Furthermore, learning could explain a permanent increase in demand for bus services. Travellers may learn about the service and quality of buses by actually testing and experiencing them. Foster et al. (2012) link the importance of consumer learning to plant growth. Alternatively, consumers may be pushed out of previous habits or update their beliefs on the relative quality of the two goods. In addition, they may have changed their perception about the reliability of rail, or they may have obtained new information from increased media coverage of inter-city buses during the strikes. Coates and Harrison (2005) find a negative impact of labour disputes over player salaries on future game attendance in Major League Baseball in the US. Their results point to additional potential mechanisms at play: retaliatory motives and damage to the brand. While related to my research question, the precise mechanism at work is a question for future research beyond the scope of this chapter. This chapter's main contribution is to show the demand effects on inter-city buses during the rail strike, and to test the effect's persistence.

Finally, this chapter is among the first of a small but growing body of literature which studies the German market for inter-city buses. The German market for buses was liberalized with the explicit intent of increasing inter-modal competition. New liberalizations are currently under consideration in several other European economies. Thus, the primary concern of this literature has been to study the impact of the market liberalization of German buses on rail ticket prices and services. Böckers et al. (2015) find that the effect on the DB network was larger at the periphery of the network.[5] Bataille and Steinmetz (2013) and Hirschhausen et al. (2008), provide theoretical models on the effect of the liberalization. These studies of inter-modal competition relate to a slightly older literature on the entry of low-cost airlines into Germany in the early 2000s (Friebel and Niffka, 2009). Dürr et al. (2015) study competition within the inter-city bus market, and estimate the price effect of a recent large merger of MeinFernbus and Flixbus.[6] Neither of these studies considers the effect of the recent German railway strikes. Further, the studies rely on data from online price comparison websites which usually provide few time-series observations. Given the uniqueness and level of detail of the booking dataset, the descriptives presented in this chapter contribute to a much improved insight into this young market and its dynamics.

[5] See also Evangelinos et al. (2015).

[6] See Gagnepain et al. (2011) for a more general review of bus market competition.

3.1.3 The German railway strikes of 2014-2015

The locomotive drivers' union (Gewerkschaft Deutscher Lokomotivführer; hereafter referred to as GDL) is a relatively small but powerful union, and has a long history of disputes with DB. The 2014-2015 negotiations, however, constituted the most ferocious industrial action in the history of DB. Two factors contributed to the ferocity of the dispute: GDL was in a power struggle with a rival union, and new legislation was under review which threatened GDL's right to represent service personnel in future wage negotiations. Between September 2014 and May 2015 the dispute resulted in nine strike waves and 22 days affected by strikes – 354 hours of service disruptions. Because of the importance of the rail network to the economy, the dispute was followed extremely closely by both the German media and the public.[7]

In the 2014-2015 labour dispute, there were nine strike waves as specified in Table 3.1. I study the effects of the three major waves in 2014 (strikes 4-6; bold in Table 3.1), and disregard all strikes in 2015, because they coincide with the merger of MFB and rival competitor Flixbus in January 2015. In addition, I disregard minor warning strikes, as they only lasted a few hours and were announced with many days advance warning. My data suggest that the strikes were too short to have any measurable impact on the bus market. Customers could re-arrange their travel plans within the rail network at little cost.

The 2014-2015 strikes display several desirable features for an ideal quasi-natural experiment. Firstly, the timing of the strikes was arguably exogenous. Strikes result from a breakdown of negotiations, the exact timing of which is unpredictable as negotiations often collapse quickly and unexpectedly. Once negotiations have broken down, the exact timing of a strike is still not clear. It could be delayed by

[7]This chapter is concerned with passenger transport. Note, however, that the railway strikes affected both passenger and freight services by DB.

Table 3.1
Dates and duration of railway strike waves in 2014-2015

Nr.	Strike Begin:	—	Strike End:	Duration (in hours):
1	Mon. 01/09/2014, 18:00	—	Mon. 01/09/2014, 21:00	3*
2	Sat. 06/09/2014, 06:00	—	Sat. 06/09/2014, 09:00	3*
3	Tue. 07.10.2014, 21:00	—	Wed. 08.10.2014, 06:00	9*
4	**Wed. 15/10/2014, 14:00**	**—**	**Thu. 16/10/2014, 04:00**	**14**
5	**Sat. 18/10/2014, 02:00**	**—**	**Mon. 20/10/2014, 04:00**	**50**
6	**Thu. 06/11/2014, 02:00**	**—**	**Sat. 08/11/2014, 18:00**	**64**
7	Wed. 22/04/2015, 02:00	—	Thu. 23/07/2015, 21:00	43
8	Tue. 05/05/2015, 02:00	—	Sun. 10/05/2015, 09:00	127
9	Wed. 20./05/2015, 02:00	—	Thu. 21./05/2015, 19:00	41

Notes: Bold rows indicate waves studied in this chapter. Strikes in 2015 are disregarded, because they coincide with the merger of MFB and rival competitor Flixbus in January 2015. * indicates warning strikes. Warning strikes are ignored, because they only lasted a few hours and were announced with many days' advance warning.

days, weeks or months if the parties are hopeful of making progress or political pressure is exerted. The trade union centrally decides to go on strike after consulting its members. Importantly, there is no evidence to suggest that competition from buses played any role in the occurrence, timing or length of the strikes. The strikes can be considered an exogenous positive demand shock to the German bus market. Having reached a decision, GDL usually announced strikes at short notice to maximize their impact. Each strike was typically announced only two days in advance.[8] Delaying or rescheduling a trip in anticipation of a strike was not possible for the majority of travellers. Consumers were directly affected. Figure 3.1 provides a detailed timeline of the two distinct weeks in which the rail strikes took place, which I cover in this chapter. It outlines the short pre-announcement period before each wave and the length of the strike.

[8]In the empirical exercise, I drop the two departure days before and after each strike wave to remove any anticipatory effects (see Ashenfelter and Card, 1985). The descriptives presented in Section 3.2.2 suggest that anticipatory effects are negligible.

FIGURE 3.1
Timeline of rail strike in weeks October 13-20 and November 03-10, 2014

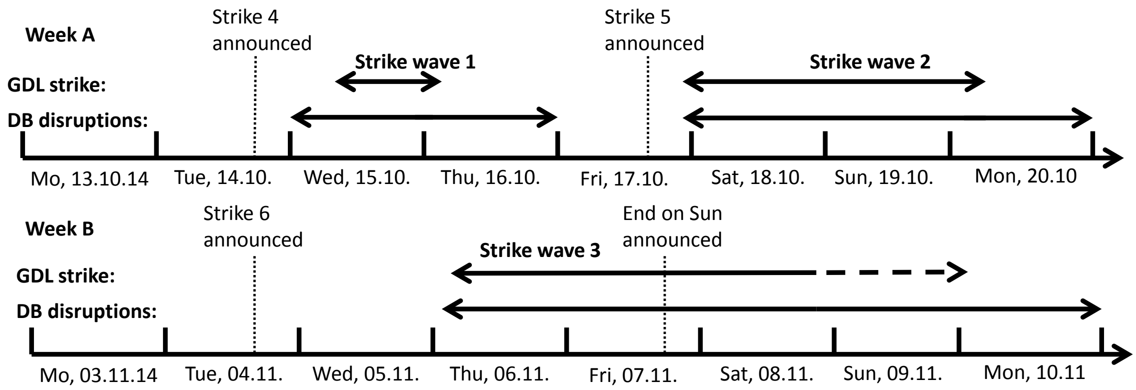

Notes: DB disruptions start before the first strike wave because DB adopted its emergency timetables with the beginning of the departure day to minimize the overall impact of the strike. DB disruptions lasted beyond the duration of each strike wave as it took time to return to normal timetable operations. Furthermore, the third rail strike wave in week B was ended prematurely on Saturday, although it had initially been announced to last until Sunday (as indicated by the dashed line). Following public pressure, the GDL announced it would return to work on Sunday November 9^{th} to allow travellers to reach the anniversary festivities of the Fall of the Berlin Wall around the country. Strikes 4-6 refer to Table 3.1. Throughout this paper I refer to the strikes as waves 1-3.

Secondly, GDL called for a strike nationwide. However, neither did GDL shut the network down entirely, nor were rail routes exposed to the same degree. GDL membership strength is weaker in West Germany, because many West German train drivers have civil servant status – a relic of DB's historical status as a state company.[9] The emergency timetables operated during the rail strike reflect the varying power of GDL across Germany. The change of service frequency specified in the emergency timetables was exogenous to the bus market: DB did not strategically focus rail services on routes which were under particular threat of competition from buses. The emergency timetables were the same in all strike waves in 2014-2015 and they are almost identical to those employed by DB in the last railway strikes of 2007-2008; i.e. long before the liberalization of the inter-city bus market in 2013.[10] Finally, DB made no attempt to employ locomotive drivers

[9]German civil servants have by law no right to strike or unionise.
[10]A direct comparison of emergency timetables in 2007-2008 and 2014-2015 is difficult because normal DB timetables have changed substantially. However, *rail lines* have changed little. Over

outside their usual geographic area of deployment for fear that they might be unable to return at the end of the day. While the exact rationale for offering some services over others is unclear, the geographic variation in strike exposure mirrors GDL membership, not the inter-city bus network. I discuss the transmission of the rail strike on bus routes in Section 3.3 below.

Thirdly, excluding those under focus, the last major rail strikes date back to 2007-2008, but the market for inter-city buses was not liberalized until 2013. In the 2014-2015 labour dispute, inter-city buses – a clearly defined rail substitute – were a viable alternative for the first time. Car and airline services were, of course, available in previous strikes. The inter-city bus market not only received substantial media coverage during the strikes but also attracted many travellers who had never travelled via inter-city buses before. For example, in an April 2014 survey prior to the strike, only 12 percent of young Germans indicated that they had used the newly available bus services (YouGov, 2014). Among older age groups this percentage is likely to be even lower because the trade-off in accepting longer travel times and less convenience for cheaper fares typically appeals to younger customers.

Fourthly, switching between rail and bus can be done quickly and easily.[11] Tickets can be bought through price comparison websites via the internet or on the bus. Furthermore, bus departure terminals are located directly next to the rail station in most cities (Guihéry, 2015). Travellers could arrive at the rail station and easily transfer to inter-city buses when the implications of the rail strike became clear to them.

60 percent of rail lines had nearly the same fraction of service cancellations in 2007-2008 and 2014-2015.

[11] DB does not offer season passes on specific routes. It offers the *BahnCard* which grants fixed price reductions to card holders. *BahnCard* subscriptions can be cancelled annually. This may have locked travellers in to the services of DB, in which case any lasting effect beyond the strike would not be visible until the medium or long-term.

3.2 Data and descriptive statistics

3.2.1 Data

This chapter combines three novel and extremely rich datasets: detailed booking data provided by MFB, DB emergency timetables, and a dataset of all rail itineraries. The latter dataset is collected using a web-crawler linked to the website of a leading price comparison website – a collection approach rarely used in the economics literature. I combine the emergency timetables and travel itineraries to create a dataset of service cancellations and expected delays caused by the rail strike. I summarize key features of the data below. Given the novelty of the data, I document additional information on the construction of all variables in Appendix C.1.

MeinFernbus booking data

MFB is Germany's largest bus provider with a market share of roughly 50 percent during the sample period. In addition to being the key player in the German inter-city bus market, MFB's service quality as well as strategic use of local bus partners are representative of the entire inter-city bus industry.[12]

The dataset provided by MFB contains the universe of MFB ticket sales between any combination of 33 large German cities for departure dates from September 01^{st} to December 31^{st} 2014. Individuals who departed in the sample period, but who booked their ticket outside the sample period are also included. The original dataset contains about 1.7 million individual bookings. A booking observation includes detailed information on the bus service such as the route, price, departure date and time as well as information on the individual in form of an anonymized

[12]For example, free internet, luggage allowance, and leg-room are almost identical across the industry. See Dürr et al. (2015) for detailed introduction and comparison of players in the inter-city bus market.

e-mail address. The e-mail address identifies first-time and repeat bookings by an individual, and thus allows following a customer over time.

The key variable of interest is the natural logarithm of the number of tickets sold at the route and departure day level.[13] Thus, I aggregate the individual bookings at the route and departure day level – the unit of analysis in this paper.[14] A route is the combination of an origin- and destination- city, so different routes may be served by the same bus journey. For example, a bus ride from Munich to Berlin with a stop in Dresden serves three routes: Munich–Dresden, Munich–Berlin and Dresden–Berlin. I treat each route as an independent and separate market. This has the advantage that it captures travellers such as commuters who repeatedly travel. For these people I can calculate their precise exposure to the rail strikes. The drawback is that this definition does not capture travellers who return after the strike but travel on a different route.[15]

While rail strikes continued beyond the sample period to May 2015, I restrict the sample period to 2014. This is because MFB unexpectedly merged with rival bus provider Flixbus in January 2015. Any changes after this date may be driven by the effects of the merger and not the rail strike.

Figure 3.2 lists and maps all 33 cities in the sample. However, not all route combinations are served. Inter-city buses are not legally permitted to connect cities at less than 100km distance or where local train travel time does not exceed one hour. Some routes are only served on some weekdays or not served at all. I

[13]The dependent variable is computed as ln(1+tickets sold) at the route departure day level. This approach is common in the trade literature, and allows me to keep route-day observations with zero tickets sold (see Felbermayr and Kohler, 2006). In the dataset, zero observations only account for 0.3 percent of tickets sold and 7 percent of tickets sold to new customers. I confirm that my results are unaltered if I drop all zero observations.

[14]For clarity note that there are two time dimensions to each individual booking: the date of the booking and the date of the departure. I aggregate ticket sales to the route and departure date dimension. 95 percent of bus travellers arrive on the same date as they depart.

[15]In the later difference-in-differences analysis, strike-exposed customers, who return on a different (non-treated) route in later journeys, would bias the estimated effect downwards. Thus, the estimated effect could be interpreted as a lower bound to the true effect.

FIGURE 3.2
Map and list of German cities in the sample

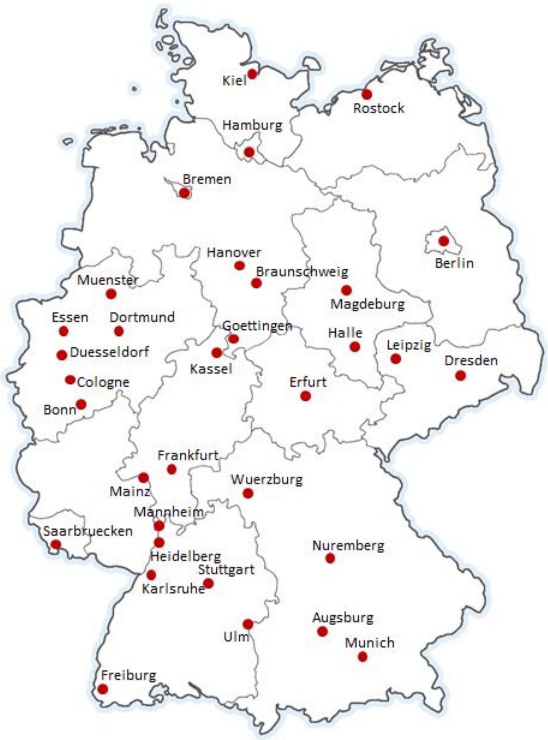

Cities:	
Augsburg	Heidelberg
Berlin	Karlsruhe
Bonn	Kassel
Braunschweig	Kiel
Bremen	Leipzig
Cologne	Mainz
Dortmund	Magdeburg
Dresden	Mannheim
Duesseldorf	Munich
Erfurt	Muenster
Essen	Nuremberg
Frankfurt (Main)	Rostock
Freiburg	Saarbruecken
Goettingen	Stuttgart
Hamburg	Ulm
Halle (Saale)	Wuerzburg
Hanover	

employ a strict definition of which routes to include in the dataset: I drop those routes on which the number of days in the sample in which no customer travels that route exceeds ten. I confirm that my results are not sensitive to this cutoff. Cutting the dataset in this way represents a trade-off between clarity and statistical power. Given the size of the dataset, however, this is not a major problem.

The final panel contains a cross-section of 312 routes and roughly 34,000 observations at the route and departure day level. The dataset is balanced in the sense that all routes are observed over the entire sample period and through all strike waves.

DB Emergency timetables and web-crawled itineraries

In addition to the MFB booking data, I construct a dataset of DB service cancellations and expected delays for each route during the rail strikes. This dataset combines emergency timetables provided by DB during the strikes and a dataset of all DB travel itineraries, which was collected using a web-crawler linked to the website of a leading price comparison website. The former provide data on normal frequency and the frequency during the strikes of all rail lines. The latter dataset includes all travel itineraries for the routes of the dataset during a complete week. A travel itinerary is defined as the specific departure times, stopovers and train numbers a traveller needs to take on a rail journey.

The DB emergency timetables list DB services at the *line level*. For example, ICE line 25 from Hamburg to Munich halved its operations from once every hour to once every two hours. However, actual travel itineraries are much more complex and often involve stopovers. A typical itinerary involves the use of multiple rail lines. Using actual itineraries takes into account that some DB routes are served through different paths in the rail network. Only the combination of emergency timetables and the travel itineraries allows me to construct the average exposure of each route to the rail strike. One data limitation remains, however: the DB emergency timetables do not include information on regional trains. I disregard routes where more than 10 percent of itineraries include the use of regional trains. This is not a major problem. Since the data focus on connections between the largest German cities, most itineraries include inter-city lines only.

To measure each route's exposure to the rail strike, I construct two variables: the fraction of cancelled rail departures during the strikes (*fraction services cancelled*) and the expected time delay (*additional travel time*). The expected additional travel time travellers have to incur to reach their destination is calculated as the time a traveller has to wait for the next train if their service is cancelled. On the one hand, this measure takes into account the typical stopovers involved on each route. On the other hand, I neither observe delays in the travel time due to unexpected stopovers, nor delays due to unexpected additional halts. Furthermore, actual waiting times may have differed substantially depending on the actual arrival of travellers at the rail station, which is unobserved. However, defining additional

Figure 3.3
Panel A: DB travel time normal vs. expected additional travel time.
Panel B: DB travel time normal vs. and fraction of services cancelled for each route during the rail strike

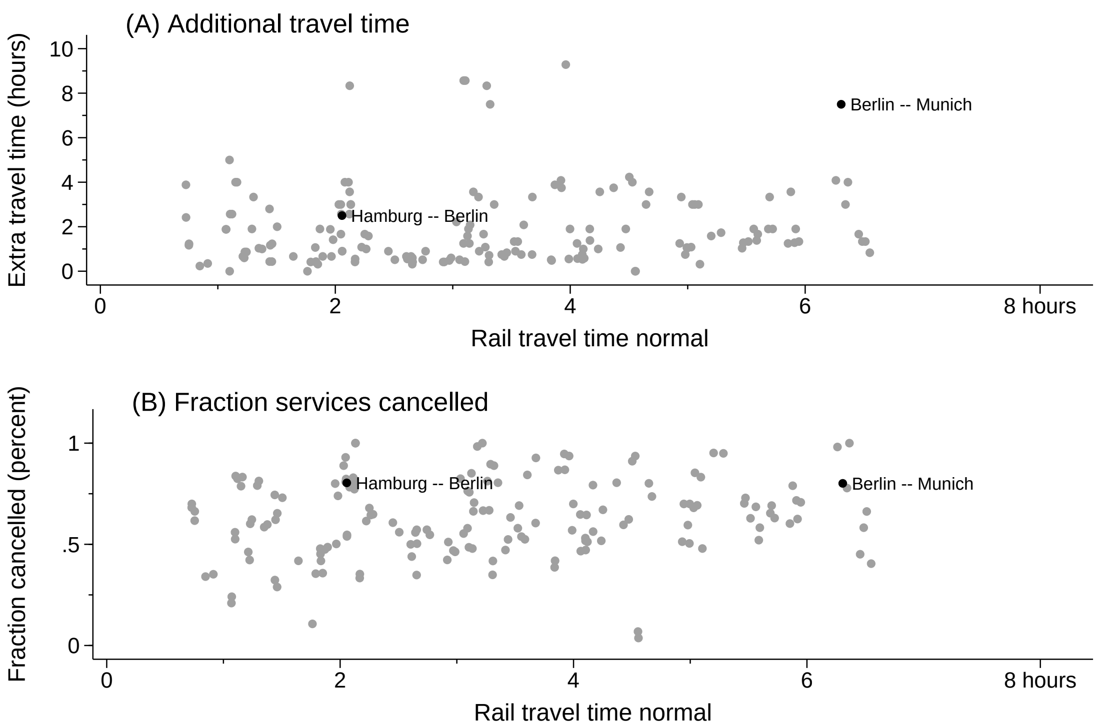

Notes: Datasource DB emergency timetables. Routes Munich–Berlin and Hamburg–Berlin are highlighted as examples.

travel time in this way has the advantage that it mirrors the structure of the emergency timetables, the primary source of information available to customers. Figure 3.3 plots the rail travel time in normal times against the additional travel time (Panel A) and the fraction of rail services cancelled (Panel B) for all routes. The routes Berlin–Munich and Hamburg–Berlin highlight the difference between the two measures. While both routes had almost identical service cancellations (about 75 percent), the time a customer had to wait for the next train was much longer for Berlin–Munich. This is because Hamburg–Berlin operated at a much higher frequency even in times of the strike. In addition, note that there is no visible systematic relationship between rail travel time and the strike-exposure measures.

3.2.2 Descriptive statistics

Before turning to the econometric analysis, I present some descriptive statistics. Given the novelty and level of detail of the dataset, they may be of more general interest. Additionally they highlight some important features of the data and clarify some selection choices I make for the empirical regression exercise below.

Figure 3.4 presents aggregate changes to MFB services over the sample period. For ease of interpretation I report weekly data.[16] Panel 1 plots the key variable of interest *ln ticket sales* for each departure. Sales peak during each strike wave as well as on national holidays such as October 3rd which in 2014 fell on a Friday, thus creating a long weekend. As expected, the increase in ticket sales is particularly pronounced for first-time customers (Panel 2). Panels 3-6 plot supply related descriptive statistics. Panels 3 and 4 display the negative trend in total capacity and departures over the sample period, reflecting the seasonality of public transport demand. Demand is weaker in winter and MFB reduced the frequencies of its services, especially on off-peak weekdays. Panel 4 indicates that MFB, despite the short time-frame of each strike announcement, was able to increase its capacity during the rail strikes. Panels 5 and 6 address the capacity utilization of MFB.[17] A concern might be that customers were not able to switch to inter-city buses during the rail strikes, because buses were operating at full capacity. If so, the number of people exposed to inter-city buses would be much lower, and the estimated effect on bus ticket sales should be considered a lower bound. As indicated in Panels 5 and 6 MFB buses have additional tickets available in more than 80 percent of

[16] Ticket sales on Friday and Sunday exceed weekday sales on Tuesdays and Wednesdays by a factor of almost two. Share of ticket sales per weekday for the dataset: Monday 13%, Tuesday 10%, Wednesday 10%, Thursday 12%, Friday 19%, Saturday 15%, Sunday 20%.

[17] Because a bus has multiple stops, the remaining capacity for each route does not correspond to the number of ticket sold for that route. For example, a bus that travels from Munich to Berlin via Dresden with 50 seats may be at capacity between Dresden and Berlin if 30 tickets were sold from Munich to Berlin and 20 from Dresden to Berlin. To address this issue, Panels 5 and 6 plot the bottleneck capacity: the remaining capacity for the section of the bus trip where the bus was most full.

FIGURE 3.4
Aggregate weekly descriptives on MFB ticket sales and supply

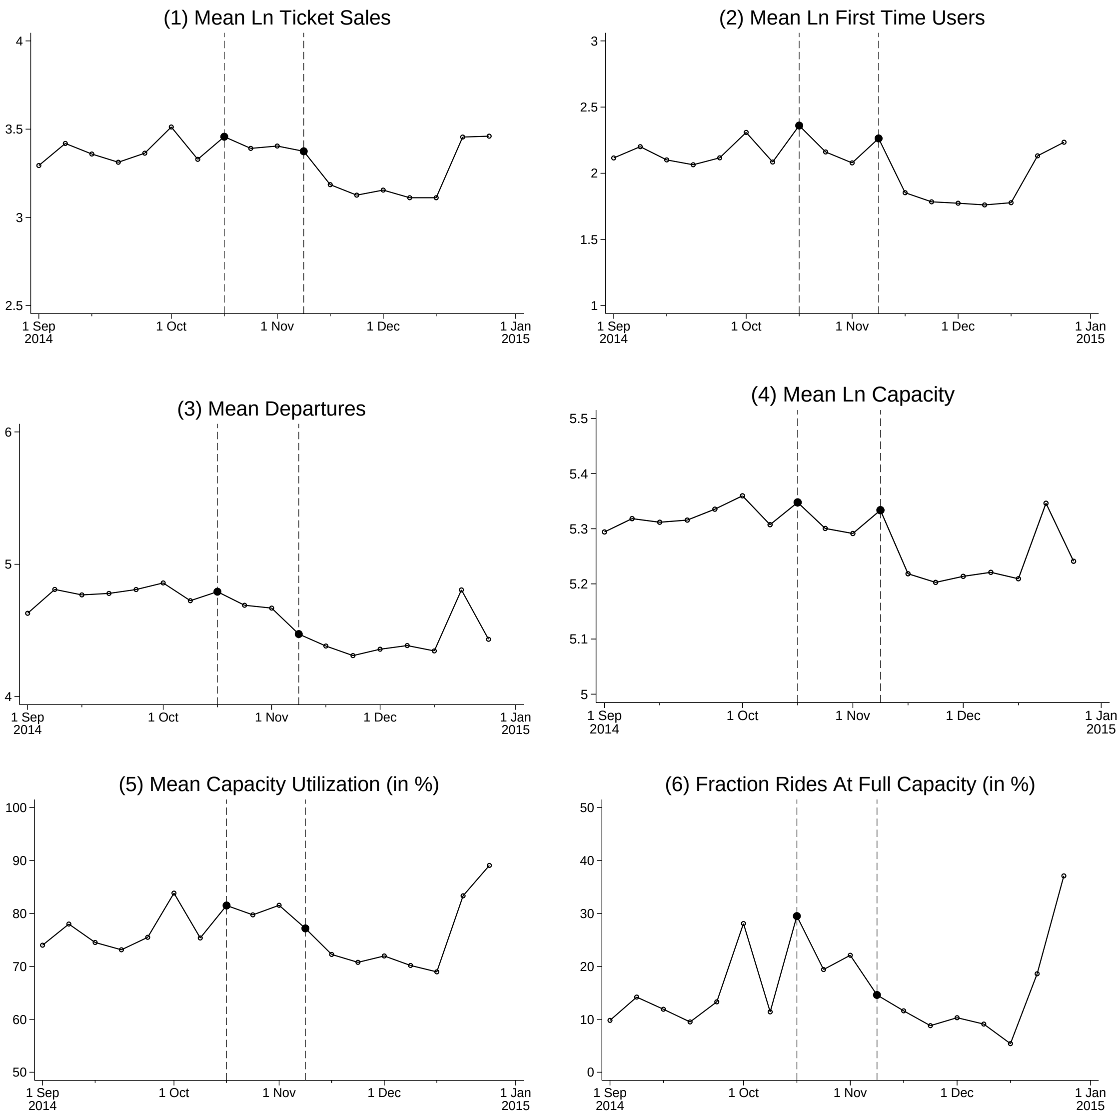

Notes: Sample at route - departure date dimension. Panels 1-4 report weekly averages over all routes. Panels 5 and 6 report averages for each bus journey. Panel 1 reports the average log number of total tickets. Panel 2 reports the log number of total tickets sold to first-time customers. Panel 3 reports the average daily departures per route. Panel 4 the daily capacity per route. Panels 5 and 6 report descriptives relating to the capacity utilization of MFB services. Vertical line and bold circles indicate weeks in which GDL was on strike.

Figure 3.5
Mean total ticket sales split by returning and new customers

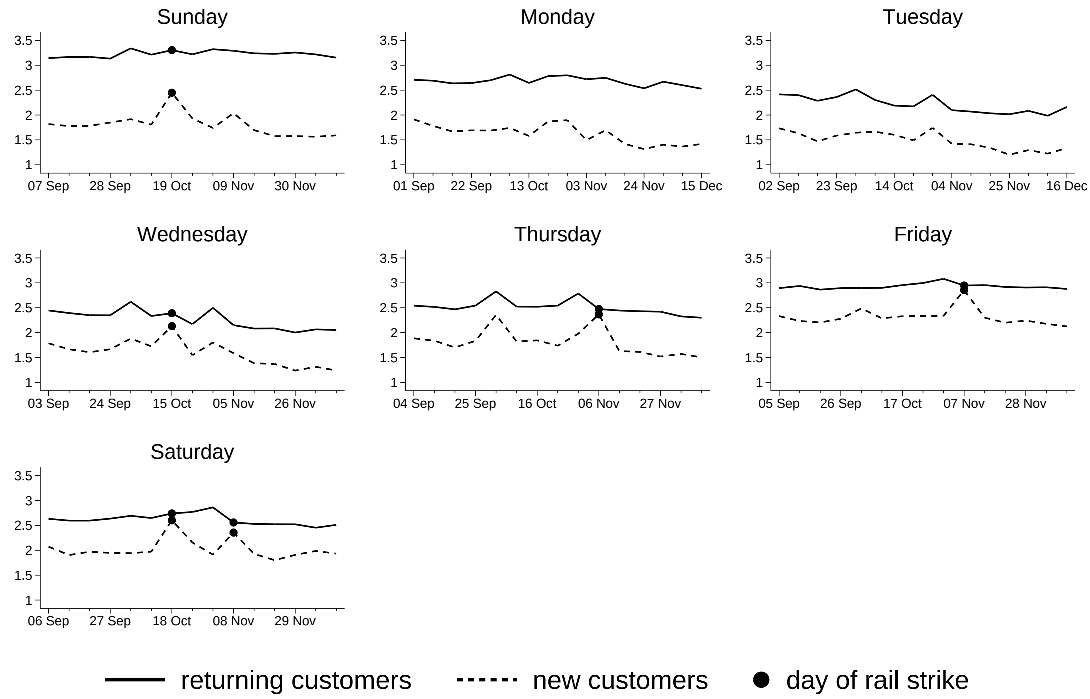

Notes: Data are split by weekday and bold circles indicate that the weekday was affected by a strike.

departures. This fraction does not increase substantially during the rail strike. Even if customers were faced with a fully-booked bus during the strike, there is a high probability that they could have successfully bought a ticket on the next bus.

The key takeaways from Figure 3.4 are twofold. Firstly, MFB ticket sales data display seasonality. To make sense of the effect of the strike, it is important to have an appropriate control group in the empirical analysis. Secondly, I drop the final two weeks of observations. Figure 3.5 displays how exceptional the Christmas travel period is. I do not want this seasonal shock to obscure my results. Cutting the dataset in this way represents a trade-off between clarity and statistical power. Given the size of my dataset, this is not a major problem. The remaining 36 post-strike departure days allow me to estimate the short- and medium-run effects of the rail strike.

Figure 3.6
Mean cumulative bookings for Friday departures

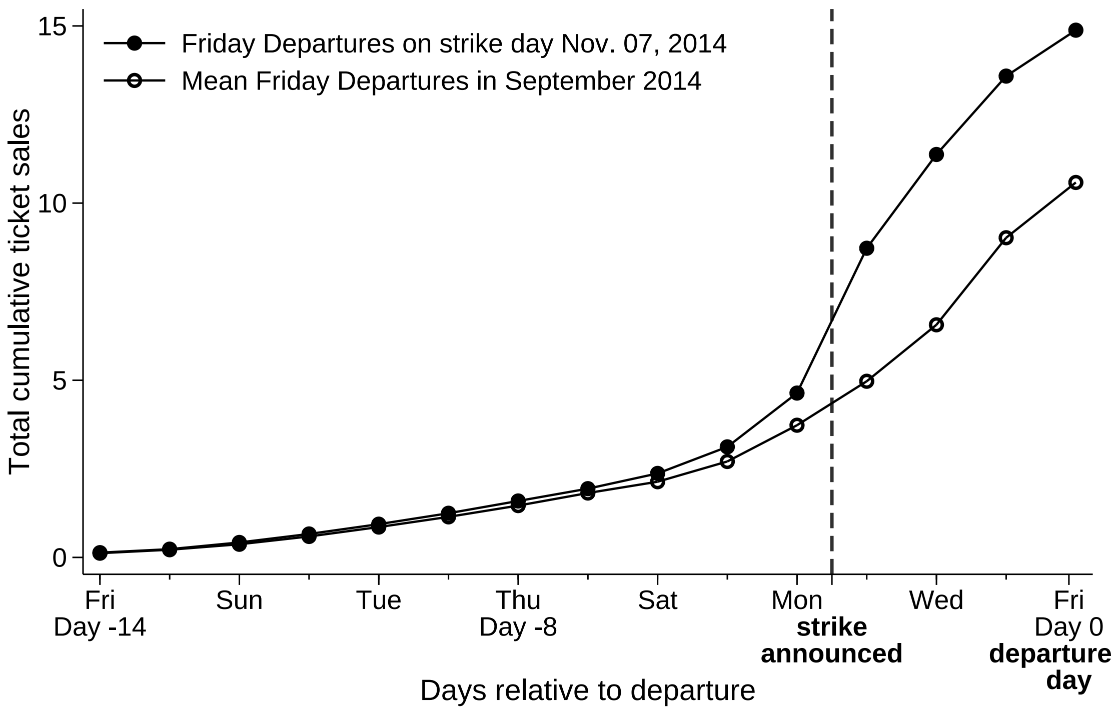

Notes: Data are split into bookings for Friday-departures in September, the month just preceding the rail strike, and bookings for departures for strike day November 07, 2014. The strike was announced 3 days prior to the strike (as indicated by the dashed line). Note that ticket sales are not in log scale here.

Figure 3.5 splits ticket sales into returning and new customers. The figure suggests the positive effect of the rail strikes on ticket sales during each strike wave. Sales during the strikes were almost exclusively driven by customers who had never previously travelled by inter-city buses. On average, 30 percent of bus passengers are first-time customers, and two thirds of these undertake at least one more booking in the future.

An additional concern may be that customers switched to buses for reasons unrelated to the strike. While my regression analysis controls for unobservable effects with fixed effects and indicators for observable events such as school holidays, there may have been unobserved parallel events that drove bus ticket sales during the rail strikes. To address this concern, Figure 3.6 compares cumulative bookings prior to departure for a day affected by railway strike with a typical booking curve. The dashed vertical line indicates the moment of the strike announcement for the third strike wave on November 07, 2014.[18] As is apparent, ticket sales only diverge from their usual trend after the rail strike was announced. The small sales departure from the usual trend before the announcement suggests that a few travellers booked bus tickets after negotiations had broken down, but before the strike was announced; i.e. very few travellers anticipated the strike. If travellers booked tickets for buses for departure days before the strike in anticipation, my results would be downward biased. While I cannot observe whether new bus customers switched from the railway, Figure 3.6 provides strong descriptive evidence that it was the rail strikes that drove the peak in ticket sales on the striking days.

3.3 Impact during the strike

3.3.1 Potential transmission channels

While the exposure of *rail routes* to the strikes can be deduced from the emergency timetables, the exposure of *bus routes* is not ex-ante clear to the researcher. In an ideal natural experiment rail and bus would be perfect substitutes, and customers would be perfectly informed about the exposure of their proposed route to the strike. They would experiment with buses only if affected by the strikes, and if inter-city buses were a reasonably attractive alternative. However, bus and rail services are neither perfect substitutes nor were customers perfectly informed about each route's exposure to the strike.

Thus, this section tests three potential channels that could determine the variation in exposure of the strike on inter-city buses *during* the rail strikes, and consequently the definition of the treatment group.

[18] See week B of Figure 3.1.

The transmission channels can be broadly categorized as follows. Firstly, bus and rail services are not perfect substitutes. The quality of bus and rail services differs both in observable characteristics, such as travel time, as well as unobservable characteristics, such as comfort. Relative and absolute travel time matter. For example, a trip from Hamburg to Berlin takes two hours by rail and three hours by bus while a trip from Munich to Berlin takes about six hours by rail, and only one hour more by bus despite the longer absolute travel time. It is unlikely that many travellers would have opted to take the bus on routes where the bus travel time significantly exceeds that of the railway. Instead, they may have simply cancelled their trip or opted for other transport modes such as cars or aircraft. Another quality characteristic is comfort. Despite offering free internet access, the comfort of travelling by bus is generally regarded to be lower than rail travel. In this case consumers may value additional travel time in a bus differently to additional travel time by rail. They may be unwilling to take the bus above a certain threshold travel time. Finally, bus and rail services differ in price. Buses are generally cheaper than DB services. It follows that it is unlikely that customers weren't able to afford to switch during the strike. Travellers, who had booked a rail ticket, could demand a refund during the strikes even if some later trains were available.

Secondly, travellers were not perfectly informed about emergency timetables and their exposure to the strike. They may have struggled to obtain the relevant information about their personal exposure to the rail strike. In addition to publishing detailed emergency timetables, DB operated a free hotline for customers. Given that rail strikes were announced with little notice, most travellers are likely to have purchased their ticket previously. Thus, they had strong incentives to inform themselves about delays and service cancellations relevant to their itinerary. However, it is unclear whether they were able to do so. It is indeed possible that travellers on all routes considered themselves to be affected by the strike. There is some anecdotal newspaper evidence which confirms this suspicion. It reports that some of the railways in operation during the strikes – instead of being overcrowded – were emptier than usual.[19] Moreover, travellers may not have trusted DB's ability to successfully implemented its emergency timetables. The ability to implement the emergency timetables often depended on the exact number of train drivers that would turn up (or not) on the strike day – the precise number of which was often uncertain until the last minute.

Thirdly, the effect of the strike on MFB ticket sales may be the result of a combination of service cancellations from the strikes and the closeness of substitution between the transport modes. Travellers may have switched to inter-city buses if their itinerary was significantly affected *and* inter-city buses were a sufficiently attractive alternative to DB services on their route.

Since it is not ex-ante clear which routes were affected during the rail strike, and which were not, I test each of these three potential transmission channels using a number of proxy variables specified below.

3.3.2 Specification

I restrict the dataset in three ways. Firstly, since the focus of this section is on the effect during the strike, I disregard the post-strike period so as not to condition results on post-strike outcomes. Secondly, I restrict the data to focus on ticket sales to first-time customers that booked in the final three days to departure.[20] This decision uses the level of detail of the MFB booking data and is motivated by the findings in the descriptives section: ticket sales to new customers give a clearer indication of the transmission channel during the strikes. Further, strike-related

[19] Source: manager-magazine (url: *http://www.manager-magazin.de/lifestyle/artikel/jeder-zweite-gueterzug-und-jeder-dritte-personenzug-faehrt-a-1001657.html*; 07/11/2014)

[20] Note that there are two time dimensions to each booking observation: the date of the booking and the date of the departure. Here I aggregate ticket sales to the route and departure date dimension if the ticket was booked in the final three days to departure. As outlined in Figure 3.6 this primarily captures booking after the announcement of the strikes by GDL.

bookings occurred primarily in the final days before departure; i.e. after GDL announced the precise timing of the strike. Thirdly, I disregard all ticket sales for departures two days before and after each strike. As outlined in Figure 3.1, there may be anticipatory effects and lagged treatments as DB services require time to return to normal operations. In addition, I disregard the intermediate fortnight between the second and third strike wave. It is not clear whether there would be a treatment effect between the strike waves in my sample.

My baseline regression takes the following form:

$$\text{ln ticket sales}^{\text{new}}_{ijt} = \alpha_{ij} + \tau_t + X_{it} + X_{jt} + \delta\,(\text{channel}_{ij} \times \text{strike}_t) + \epsilon_{ijt} \tag{3.1}$$

where ij refers to a route from origin-city i to destination-city j, and t to the departure day. The dependent variable *ln ticket sales*$^{new}_{ijt}$ is defined as the log of tickets sold to new customers in the final three days to departure. α_{ij} and τ_t are route and departure day specific fixed effects respectively. The route fixed effects capture observed and unobserved differences that are constant over time such as distance. The time fixed effects capture the effects of observed and unobserved temporal factors common to all routes such as national holidays, MFB marketing campaigns, or seasonal fluctuations.

X_{it} and X_{jt} are vectors of city-departure date specific control variables: A dummy for public holidays, school holidays and dummies for other major events.[21] I list all control variables used in the regressions in Table 3.2. Each control variable is interacted with month and weekday indicators to capture more variation in the data. Finally, the specifications with controls include origin- and destination-specific linear time trends.

[21]Note that German school holidays vary at the state level. Thus, school holidays are not captured by the departure day fixed effects. Source: schulferien.org

In an additional specification, I include origin- and destination- departure day specific fixed effects, denoted γ_{it} and γ_{jt} respectively.[22] This is my preferred specification. Note that the inclusion of these route-specific fixed effects nests a complete set of origin and destination specific fixed effects. Furthermore, these strong fixed effects make the inclusion of the departure day fixed effects and the control variables redundant.

ϵ_{ijt} is the error term. Using a difference-in-differences strategy with many years, I have to worry about serial correlation at the group level. Conventional standard errors may severely understate the true standard errors (Bertrand et al., 2004). To address potential serial correlation within routes and time correlation, I cluster standard errors by route throughout the paper.

$(channel_{ij} \times strike_t)$ is the interaction term of interest. On the one hand, $strike_t$ is a vector of indicators for each strike wave studied in this chapter. As discussed in the background section, I disregard minor warning strikes, as they only lasted a few hours and were announced with many days advance warning. Any impact of these earlier warning strikes would bias my results downward. On the other hand, $channel_{ij}$ captures the different potential transmission channels.

To capture the effect of each potential transmission channel, I use proxy variables as follows. Firstly, I proxy the degree to which rail and bus services are substitutes using three variables: the *relative travel time difference* between rail and bus, *absolute travel time difference*, and *bus travel time.* Panels 1 and 2 of Figure 3.7 show that routes with a short bus travel time also show a small absolute bus travel time difference; i.e. both variables are strongly correlated. Thus, bus travel time captures the likelihood that, even if the absolute travel time difference is

[22]Note that the inclusion of origin-day and destination-day fixed effects mirrors the fixed effects typically used in the estimation of gravity trade models to address 'multilateral resistance' terms (Anderson and van Wincoop, 2003).

small, travellers regard buses as sufficiently comfortable only for bus routes below a certain threshold travel time.

Secondly, I measure the strike exposure using the two variables constructed from DB emergency timetables in the data section: the *fraction of services cancelled* and *additional travel time* that customers had to endure to reach their destination during the strikes. The latter explicitly takes into account the fact that some routes operated at a much higher frequency even in times of the strike.

Thirdly, both the closeness of substitution and the exposure to the rail strike could be the primary factors driving bus ticket sales during the strike. To capture this channel, I estimate a set of regressions with a triple interaction between the proxies of the above channels. The triple interaction takes the following form: $(channel_{ij}^{sub} \times channel_{ij}^{exp} \times strike_t)$, where $channel^{sub}$ are the variables from the substitution channel (*relative travel time*, *absolute travel time* and *bus travel time*), and $channel^{exp}$ includes the exposure channel variables (*fraction services cancelled* and *additional travel time*). This specification also includes the first-order interaction terms to distinguish the triple interaction term. Note that Equation 3.1 does not include the lower-order terms as they are captured by the route and departure day fixed effects.

I repeat separate regressions for each proxy variable. Moreover, I estimate each channel variable as a dummy indicating whether it is above/below the median value. This is to ease interpretation and to make the estimated regression coefficients for each proxy more easily comparable. Thus, the dummies for *relative travel time*, *absolute travel time* and *bus travel time* equal one if the route is shorter than the median. Likewise, the dummies for *fraction services cancelled* and *additional travel time* equal one if the fraction of cancellations or travel delay exceed the median value respectively. In the robustness section, I confirm that my results are unaltered to using continuous definitions for the treatment variables.

FIGURE 3.7
Treated and control routes for each channel variable

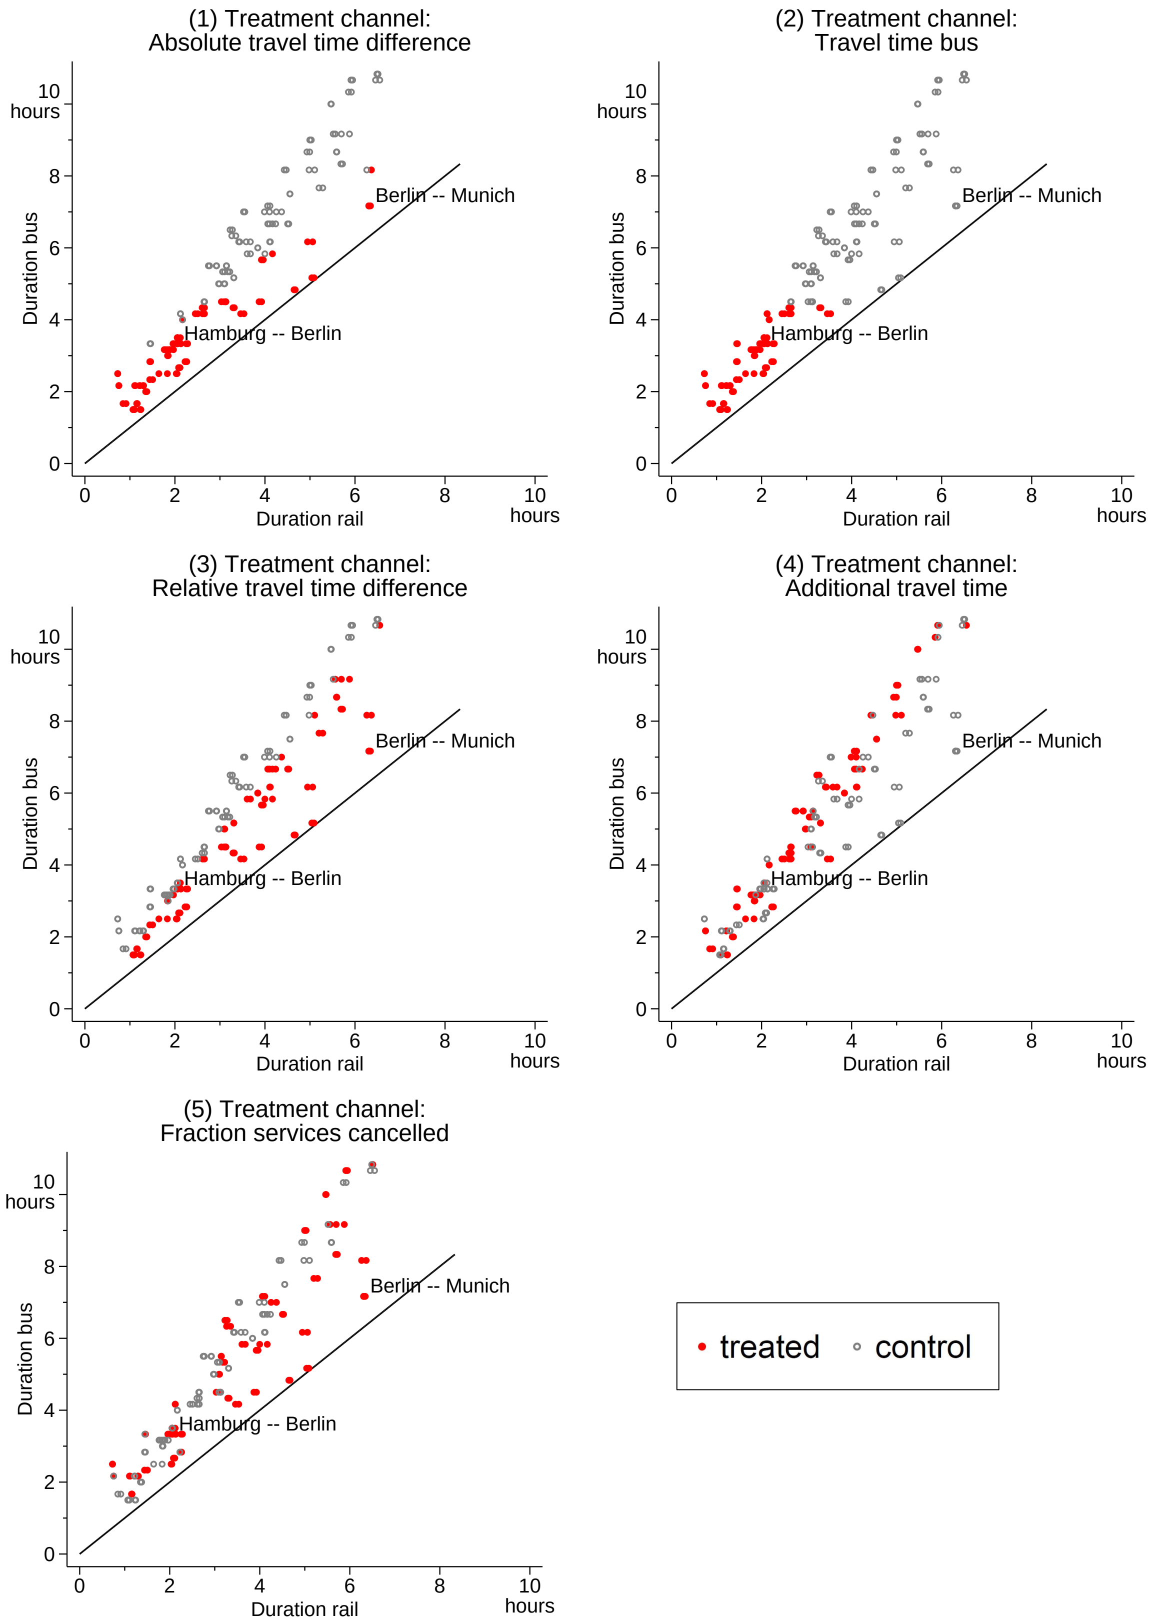

Notes: Each panels display scatter of routes in duration rail and duration bus space with 45 degree line. For each proxy transmission variable, Panels 1-5 indicates whether a route is part of the treatment or control group. *Relative travel time*, *absolute travel time* and *bus travel time* are treated if the route is shorter than the median. *Fraction services cancelled* and *additional travel time* are treated if the route is above the median value. See Table C.1 for specific variable definitions). Routes Hamburg–Munich and Munich–Berlin plotted as examples.

TABLE 3.2
Summary statistics for whole sample period

Variable:	N	Mean	Median	SD	Min	Max
Dependent variables:						
ln ticket sales$_{ijt}$	33,384	2.50	2.40	1.12	0	6
ln ticket sales$_{ijt}^{\text{new}}$	33,384	1.35	1.39	1.01	0	6
Proxy channel variables (channel$_{ij}$):						
Fraction services cancelled	17,762	0.63	0.63	0.19	0	1
Additional travel time	17,762	114.67	78.50	106.36	0	557
Relative travel time	17,762	1.64	1.64	0.34	1	4
Abs. travel time difference	17,762	116.99	109.89	67.53	5	285
Bus travel time	33,384	289.10	265.00	149.29	60	650
Control variables (X_{it} and X_{jt}):						
School holiday	33,384	0.30	0.00	0.46	0	1
Public holiday	33,384	0.04	0.00	0.20	0	1
Bundesliga (Division 1)	33,384	0.00	0.00	0.05	0	1
Bundesliga (Division 2)	33,384	0.00	0.00	0.02	0	1
Munich Oktoberfest	33,384	0.02	0.00	0.14	0	1
Stuttgart Wasen	33,384	0.02	0.00	0.14	0	1

Notes: Variables fraction services cancelled, additional travel time, relative travel time and absolute travel time difference have fewer observations because emergency time tables do not include information on regional trains. In addition, Table C.1 in Appendix C.1 provides definitions of all variables estimated in Equations 3.1 to 3.6.

Figure 3.7 displays how each channel variable divides routes into treatment and control. Routes are of course not clearly divided into treatment and control, but treatment is imprecise. A route which is classified as above the median for one of the channel variables is best thought of as being 'more treated' relative to a route below the median. Defining the treatment channel in this way has the drawback that my measure includes a number of 'false negatives' and leads to type II errors. Fricke (2015) demonstrates that in this case the estimated result will be biased downwards and could be interpreted as a lower bound to the true effect. Finally, Table 3.2 presents basic summary statistics (including the median) for the set of explanatory variables. In addition, Table C.1 in Appendix C.1 provides specific definitions of all variables estimated in Equations 3.1 to 3.6.

3.3.3 Results

Having employed this extensive combination of fixed effects and controls, the coefficient of interest indicates whether routes that were below (above) the median for one of the proposed channels differ significantly compared to routes above (below) the median. In total, I estimate Equation 3.1 in eleven regressions: a regression for each of the different proxy channel variables introduced above and triple interactions between the combination of closeness of substitution and exposure to rail strike proxies. Table 3.3 summarizes all regression results.

Based on the three transmission channels outlined above, I find no evidence for the exposure channel. The proxy variables measuring this channel, *additional travel time* and *fraction services cancelled*, yield no robust statistically significant effects during the strike. I do not find evidence for the third channel, the combination of exposure and closeness of substitution, either. None of the triple interaction terms between the proxies yield robust statistically significant coefficients. I move the regression tables C.2-C.10 to Appendix C.2 for space concerns. See rows 3-11 in Table 3.3 for a summary.

Table 3.5 reports the regression results for the proxy variable *absolute travel time difference* and Table 3.4 the results for the variable *bus travel time*. They are the only two channel variables which yield consistently robust and statistically significant coefficients. Thus, my results suggest that the primary channel driving MFB ticket sales during the strikes was the closeness of substitution as measured by the proxy variables *absolute travel time difference* and *absolute bus travel time*. This is surprising as it suggests that travellers switched to buses even on routes with little or no service cancellations. It follows that either they were not well informed about their exposure to the rail strike, or had no trust in DB's ability to implement the emergency timetables.

TABLE 3.3
Summary of regression results

Nr.	**Table Nr.**	**Transmission channel:**		**Strike wave 1:**	**Strike wave 2:**	**Strike wave 3:**
		1. Closeness of substitution:				
1	3.4	Bus travel time		✓	✓	✓
2	3.5	Absolute travel time difference		✓	✓	✓
3	C.2	Relative travel time difference		✗	✗	✗
		2. Exposure to rail strike:				
4	C.3	Additional travel time		✗	✗	✗
5	C.4	Fraction services cancelled		✗	✗	✗
		3. Combination of 1. and 2. (triple interactions):				
		Channel 1:	*Channel 2:*			
6	C.5	Bus travel time	Additional travel time	✗	✗	✗
7	C.6	Bus travel time	Fraction services cancelled	✗	✗	✗
8	C.7	Absolute travel time	Additional travel time	✗	✗	✗
9	C.8	Absolute travel time	Fraction services cancelled	✗	✗	✗
10	C.9	Relative travel time	Additional travel time	✗	✗	✗
11	C.10	Relative travel time	Fraction services cancelled	✗	✗	✗

Notes: Summary of regression results from Equation 3.1. Regression figures are reported in Appendix C.2. ✓ indicates positive and statistically significant coefficients at the 1% level for all combinations of fixed effects and controls reported in the regression table. ✗ otherwise. Please refer to Table C.1 for variable definitions.

Table 3.4
Transmission channel: bus travel time

	Dep. variable: ln ticket sales$_{ijt}^{new}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.- Day FE
Channel × Strike wave 1	0.257***	0.265***	0.249***	0.262***
	(0.0614)	(0.0613)	(0.0603)	(0.0673)
Channel × Strike wave 2	0.408***	0.389***	0.369***	0.302***
	(0.0516)	(0.0513)	(0.0512)	(0.0633)
Channel × Strike wave 3	0.488***	0.451***	0.426***	0.395***
	(0.0459)	(0.0456)	(0.0440)	(0.0514)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.748	0.754	0.757	0.816
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level. 166 clusters. ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day t specific fixed effects.

TABLE 3.5
Transmission channel: absolute travel time difference

	Dep. variable: ln ticket sales$_{ijt}^{new}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Channel × Strike wave 1	0.154**	0.170***	0.172***	0.148**
	(0.0605)	(0.0593)	(0.0578)	(0.0720)
Channel × Strike wave 2	0.218***	0.218***	0.220***	0.234***
	(0.0597)	(0.0582)	(0.0560)	(0.0667)
Channel × Strike wave 3	0.226***	0.188***	0.222***	0.308***
	(0.0521)	(0.0559)	(0.0512)	(0.0607)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.744	0.751	0.755	0.815
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level. 166 clusters. ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day t specific fixed effects.

As indicated in Figure 3.7 above, both *bus travel time* and *absolute travel time difference* are strongly correlated. One proxy variable may capture the effect of the other. Thus, I run an additional specification including both proxies simultaneously. This addresses whether travellers mainly disliked long bus travel times, or primarily cared about the travel time difference of the bus relative to rail, or both. Table 3.6 reports the results. I find that the absolute travel time difference proxy variable has no significant explanatory power in explaining ticket sales during the rail strikes once I control for the bus travel time. Thus, the primary factor explaining increased ticket sales for inter-city buses during the strikes is the length of the ride.

The magnitude of the effect during the strikes is large, but in line with expectations. Table 3.4 predicts that ticket sales to new customers in the final three days to the average route below the median bus travel time exceed ticket sales to the average route above the median by almost 50% in the third strike wave (column 1). The magnitude is similar but smaller for the other columns. As expected strike wave 1 yields the smallest coefficients as it fell on a Wednesday. Strike waves 2 and 3 fell on a weekend, whereby strike wave 3 was a longer strike.

Before using these findings to test whether the rail strike had an effect beyond the duration of the strike, I provide an additional test to confirm the results. I re-run the regression with *bus travel time* splitting the variable into 3-hour bins. The results are reported in Table 3.7. The table confirms the earlier result: the closer the substitution between bus and rail, the larger the effect during the rail strike. Column 1 of Table 3.7 suggests that routes connecting cities with a travel time below three hours observed almost twice as many bookings in the third strike wave than the longest routes in the sample. The estimated coefficients are similar in columns 2-4, where I include control variables and more demanding fixed effects.

Table 3.6
Transmission channel: Absolute travel time difference vs. bus travel time

	Dep. variable: ln ticket sales$^{new}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Absolute difference × Strike wave 1	-0.0219 (0.0793)	-0.0379 (0.0758)	-0.0460 (0.0756)	-0.166* (0.0886)
Absolute difference × Strike wave 2	-0.0926 (0.0665)	-0.0920 (0.0659)	-0.101 (0.0671)	-0.200*** (0.0765)
Absolute difference × Strike wave 3	-0.0740 (0.0634)	-0.0858 (0.0613)	-0.0954 (0.0608)	-0.103 (0.0650)
Duration bus × Strike wave 1	0.247*** (0.0677)	0.248*** (0.0664)	0.228*** (0.0665)	0.189** (0.0755)
Duration bus × Strike wave 2	0.365*** (0.0570)	0.346*** (0.0571)	0.323*** (0.0568)	0.215*** (0.0687)
Duration bus × Strike wave 3	0.453*** (0.0485)	0.411*** (0.0487)	0.382*** (0.0505)	0.351*** (0.0573)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.748	0.754	0.757	0.816
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level. 166 clusters. ***/**/* indicate significance at the 1%/5%/10% level. γ_{it} and γ_{jt} refer to specifications with origin- and destination-day t specific fixed effects.

Table 3.7
Transmission channel: Transmission channel: 3 hour bins for bus travel time

	Dep. variable: ln ticket sales$_{ijt}^{new}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Strike wave 1 × Duration 6 - 9 hours	0.0936 (0.133)	0.135 (0.130)	0.101 (0.126)	0.108 (0.143)
Strike wave 1 × Duration 3 - 6 hours	0.289** (0.126)	0.310** (0.125)	0.280** (0.120)	0.324** (0.143)
Strike wave 1 × Duration 0 - 3 hours	0.411*** (0.129)	0.472*** (0.131)	0.435*** (0.125)	0.472*** (0.154)
Strike wave 2 × Duration 6 - 9 hours	0.408*** (0.119)	0.405*** (0.115)	0.366*** (0.123)	0.345** (0.142)
Strike wave 2 × Duration 3 - 6 hours	0.556*** (0.117)	0.543*** (0.113)	0.508*** (0.120)	0.495*** (0.146)
Strike wave 2 × Duration 0 - 3 hours	0.942*** (0.119)	0.920*** (0.116)	0.878*** (0.122)	0.800*** (0.157)
Strike wave 3 × Duration 6 - 9 hours	0.411*** (0.131)	0.378*** (0.126)	0.326*** (0.122)	0.351*** (0.124)
Strike wave 3 × Duration 3 - 6 hours	0.702*** (0.129)	0.642*** (0.124)	0.606*** (0.120)	0.630*** (0.123)
Strike wave 3 × Duration 0 - 3 hours	0.980*** (0.131)	0.911*** (0.126)	0.861*** (0.123)	0.869*** (0.131)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15500	15500	15500	15300
R^2	0.750	0.756	0.759	0.817
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day t specific fixed effects.

3.4 Impact after the strike

As established in the previous section, it is primarily the closeness of substitution which increased demand *during* the rail strike. In this section, I test for any persistence of the effect *after* the rail strikes. Treatment and control groups are defined using the channel identified in the first step, namely the *bus travel time* proxy variable. As previously done, I code the treatment variable as a dummy equal one if the bus travel time of the route is below the median bus travel time.[23]

The post-strike regression takes the following form:

$$\begin{aligned}\ln \text{ticket sales}_{ijt} &= \alpha_{ij} + \tau_t + X_{it} + X_{jt} \\ &\quad + \delta_1 (\text{treated}_{ij} \times \text{strike}_t) + \delta_2 (\text{treated}_{ij} \times \text{post}_t) + \epsilon_{ijt}\end{aligned} \tag{3.2}$$

Equation 3.2 is very similar to Equation 3.1 in Section 3.3. I employ the same combination of specifications, control variables and fixed effects. The difference-in-differences (DD) methodology compares changes in the ticket sales of MFB between routes that differed in their closeness of substitution as measured by the absolute bus travel time.

However, the underlying data now also includes the post-strike period. I am interested in whether routes that were 'more treated' had significantly more customers beyond the strikes compared to the 'less treated' routes. Furthermore, the dependent variable *ln ticket sales*$_{ijt}$ is defined as the log total number of MFB customers. I no longer restrict it to new customers who booked during the final three days to departure, because I would like to investigate whether customers adjust their modal choice after their first experience of buses during the strike. The dependent variable now includes returning customers, some of whom travelled by bus for the first-time during the strike.

[23] See the robustness section for a continuous definition of the treatment variable.

TABLE 3.8
Impact after the strike – bus travel time

	Dep. variable: ln ticket sales			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.- Day FE
Treated × Strike wave 1	0.131***	0.136***	0.137***	0.127***
	(0.0462)	(0.0453)	(0.0440)	(0.0464)
Treated × Strike wave 2	0.291***	0.273***	0.273***	0.230***
	(0.0356)	(0.0348)	(0.0335)	(0.0389)
Treated × Strike wave 3	0.387***	0.359***	0.356***	0.322***
	(0.0377)	(0.0373)	(0.0344)	(0.0396)
Treated × Post	0.301***	0.284***	0.277***	0.282***
	(0.0228)	(0.0221)	(0.0196)	(0.0224)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	26832	26832	26832	26488
R^2	0.875	0.878	0.881	0.912
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level. 166 clusters. ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day t specific fixed effects.

$treated_{ij}$ indicates if a route was part of the treatment group, i.e. whether the bus travel time is shorter than the median. The interaction term ($treated_{ij} \times strike_t$) captures the effect during the strikes and should yield positive and statistically significant coefficients because this is how treatment was selected. The coefficient of ($treated_{ij} \times post_t$) then captures the treatment effect of interest: whether the treated group has significantly higher ticket sales *after* the rail strikes, that is after DB services returned back to normal operations.

Table 3.8 reports regression results for Equation 3.2. The table indicates that there was a statistically positive and significant effect beyond the duration of the rail strike. While the effect is significantly smaller in magnitude than the effect of treatment during the rail strike, it is remarkably persistent. Column 1 of Table 3.8 suggests that total ticket sales for the treated routes were almost 15 percent higher in the first strike wave, 30 percent higher in the second strike wave, and 40 percent higher in the third strike wave. Ticket sales were about 25 percent higher for the treated group after rail operations returned back to normal. Its magnitude is roughly the same once I include controls and different sets of fixed effects, and robust to a number of alternative specifications provided in the robustness section.

However, whether the effect can be interpreted causally depends on the identification assumption: would ticket sales for routes in the treatment group have changed the same during and after the railway strikes in the absence of a strike. I address this assumption below and present a number of robustness checks.

3.4.1 The common trend assumption

This chapter shares the typical advantages and disadvantages of a standard DD strategy. On the one hand, DD allows me to control for all time-invariant differences across routes as well as changes over time by including both route and time-period fixed effects. On the other hand, the DD identification hinges on the strong but easily stated assumption of a common trend: would treatment and control groups move in parallel in the absence of treatment? There may be time-varying confounding factors that are correlated with the treatment group.

To address whether the common trend assumption holds in this setting, I discuss a number of tests. Firstly, I use strong sets of fixed effects. My specification includes a number of time- and route-varying controls, as well as origin- and destination-specific linear trends. The different fixed effects capture any level effects such as distance or common seasonal variations. They also capture time-varying omitted

FIGURE 3.8
Mean log ticket sales split by treatment and control group

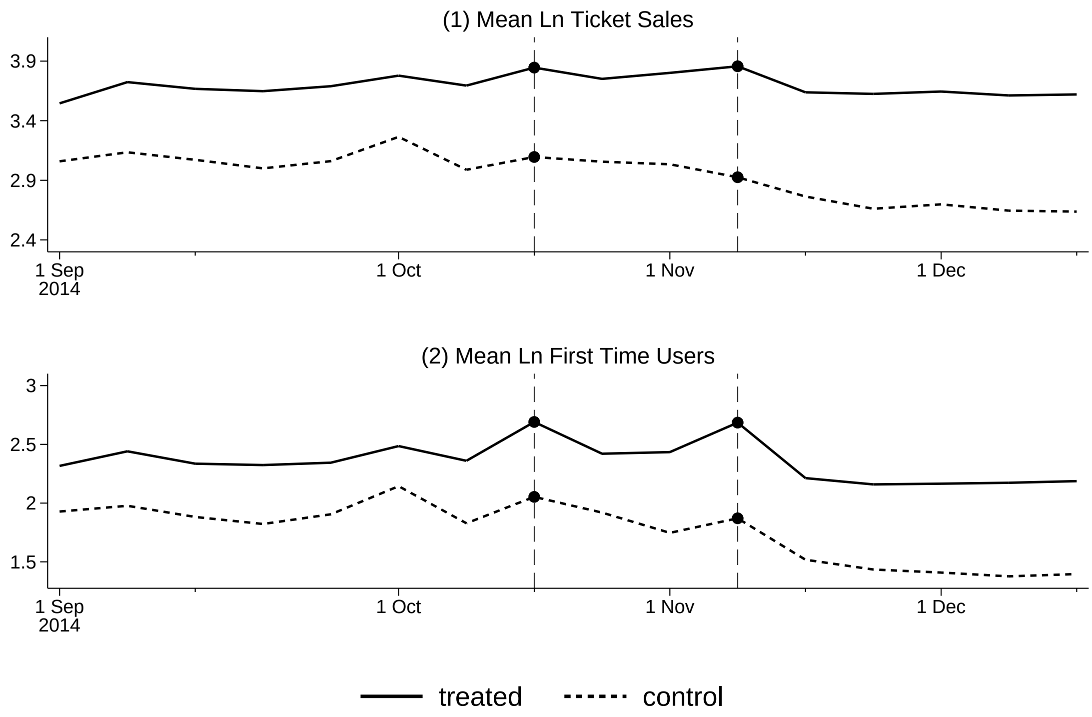

Notes: Sample period September 2014-January 2015. Vertical lines and bold circles indicate weeks and days, respectively, in which GDL was on strike.

variables such as MFB marketing expenditures. The origin- and destination-day fixed effects also capture possible linear trends. In addition, I estimate a specification with route-specific trends in Section 3.5.1 below. What remains are time-varying confounding factors that are correlated with the treatment groups.

Secondly, Figure 3.8 graphically compares the trend between the treatment and control groups for the mean log number of ticket sales to all and first-time customers. The common trend assumption meets the eyeball test. Before the rail strike, treatment and control group move remarkably in parallel. As expected, the treated group displays a visibly larger increase in sales during the strikes. The figure reports weekly averages, but a graph of daily ticket sales split by weekday yields the same result.

Thirdly, I re-estimate Equation 3.2 with pre-strike and post-strike treatment effects.[24] I report weekly coefficients to remove any weekday cyclicality. The estimated treatment effects for the pre-strike period act as a test for the common trend assumption. The pre-strike coefficients can be thought of as placebos. If trends are the same, the pre-strike coefficients should be constant and small in magnitude. If, however, pre-trends are present they would show up in the treatment group.

The specification takes the following form:

$$\begin{aligned} \ln \text{ticket sales}_{ijt} \ &= \alpha_{ij} \ + \ \tau_t \ + \ X_{it} \ + \ X_{jt} \\ &\quad + \ \delta_t \ (\text{treated}_{ij} \times \text{week}_t) \ + \ \epsilon_{ijt} \end{aligned} \tag{3.3}$$

where $week_t$ is a vector of week-fixed effects. The coefficients of interest, that is vector δ_t, must be measured relative to a baseline period. I normalize with respect to the first week of the sample which is standard in the literature. As above I run an additional specification with origin- and destination-departure day fixed effects. Unlike the previous specifications Equation 3.3 includes observations for the two days before and after each strike as well as the intermediate period between the second and third strike wave.

The plot of coefficients is reported in the main results section as Figure 3.9.[25] The coefficients report the correlation between the treated group (short bus routes) and the outcome of interest (log ticket sales) for each period. This has the additional advantage that I can evaluate the effect of the strikes over the course of the post-period: the week coefficients allow me to evaluate the effect at different elements of the post-period, as opposed to estimating an average effect only. It may take some time for the full effect to show up or for it to die out over time. The estimated

[24] Nunn and Qian (2011) and Autor (2003) provide good examples of estimating period-specific treatment effects in a difference-in-differences setting.

[25] Tables reporting coefficients of control variables and the exact coefficients are omitted for length but available upon request.

FIGURE 3.9
Coefficients of the ($\text{treated}_{ij} \times \text{week}_t$) interaction term in Equation 3.3 with 95 percent confidence intervals.

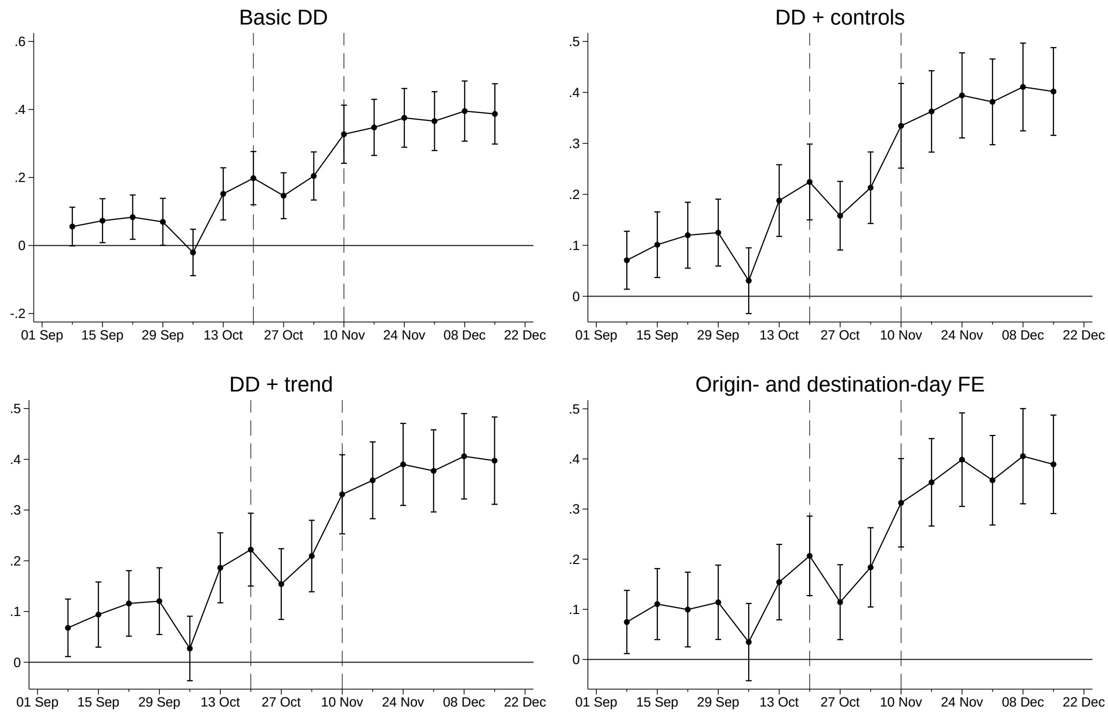

Notes: Dashed vertical line indicates weeks in which GDL was on strike. Standard errors clustered at the route level (166 clusters). Treatment variable: bus travel time.

weekly treatment coefficients are flexible in assessing the short- and medium-term effects.

The weekly treatment coefficients are reported in Figure 3.9. They display a remarkably persistent effect of the rail strike. There is a jump in the magnitude of the estimated treatment coefficients at the time of rail strikes. This jump in the magnitude of the estimated coefficients persists beyond the rail strikes until the end of the sample period. The post-strike treatment coefficients are constant around 0.4. Thus ticket sales to the treatment group are 40% higher than in the baseline period. The pattern of period-specific treatment coefficients is analogous to that of Nunn and Qian (2011). They also estimate period-specific treatment effects, and find coefficients that are constant and small in in the pre-period and increase in magnitude after treatment.

While there is a clear jump in the magnitude of the coefficients around the time of the strike, two issues cast doubt on the parallel trends assumption. Firstly, the magnitude of the treatment coefficients starts increasing too early, i.e. a week before the first two strike waves. This suggests that ticket sales for short routes already grew more strongly before the rail strike. Secondly, the post-strike coefficients are larger than the treatment coefficients during the strike, which is worrisome. This suggests that the common trend assumption is not completely tenable in the given context. If these different trends would simply reflect the heterogeneous effect of seasonality on short and long routes, and I had data from 2013, this problem may be addressed using a triple-difference-in-differences approach. However, even if these data were available the large changes in the inter-city bus market may not allow for an appropriate removal of seasonal effects.

3.5 Robustness

3.5.1 Route specific trends

Based on these results, this subsection estimates possible remedies. The possible violation of the common trend assumption suggests that there are factors which cause ticket sales to evolve differently on the control and treatment routes. For instance, there might be route-specific trends related to characteristics that affect ticket sales. I estimate two additional specifications with route-specific trends, $(\alpha_{ij} \times t)$. These capture any potential linear trend specific to each route. The regression takes the following form:

$$\begin{aligned} \ln \text{ticket sales}_{ijt} \ &= \alpha_{ij} \ + \ \tau_t \ + \ \gamma_{it} \ + \ \gamma_{jt} \ + \ X_{it} \ + \ X_{jt} \ + \ (\alpha_{ij} \times t) \\ &+ \ \delta_t \ (\text{treated}_{ij} \times \text{week}_t) \ + \ \epsilon_{ijt} \end{aligned} \tag{3.4}$$

The estimated coefficients of the ($\text{treated}_{ij} \times \text{week}_t$) interaction term in Equation 3.4 are plotted in Panels 1 and 2 of Figure 3.10. Because I cannot include route-clustered standard errors due to insufficient observations, I report robust standard errors. With this specification, the effect of the strikes will only be captured if there is a stark deviation from the trend (Angrist and Pischke, 2014). In this case, the common trend assumption does not appear to be violated.

A second specification with route-specific trends repeats the estimation using pre-period observations only following Repetto (2016): I estimate ϕ_{1ij} and ϕ_{2ij} using only data from the pre-strike period (September 01-October 14) in a quadratic trend model:

$$\ln \text{ticket sales}_{ijt} \;=\; \phi_{1ij}t \;+\; \phi_{2ij}t^2 \;+\; u_{ijt} \tag{3.5}$$

I then add the estimates for ϕ_{1ij} and ϕ_{2ij}, that is $\widehat{\phi_{1ij}}$ and $\widehat{\phi_{2ij}}$, back into the main specification. This method 'projects' pre-strike trends into the post-strike period:

$$\begin{aligned} \ln \text{ticket sales}_{ijt} \;=\; & \alpha_{ij} \;+\; \tau_t \;+\; \gamma_{it} \;+\; \gamma_{jt} \;+\; X_{it} \;+\; X_{jt} \\ & +\; \delta_{\phi_1}\,(\widehat{\phi_{1ij}} \times t) \;+\; \delta_{\phi_2}\,(\widehat{\phi_{2ij}} \times t^2) \\ & +\; \delta_t\,(\text{treated}_{ij} \times \text{week}_t) \;+\; \epsilon_{ijt} \end{aligned} \tag{3.6}$$

This specification controls for route-specific trends that were in place before the strikes and that may cause ticket sales patterns to be different across groups. I report results in Panels 3 and 4 of Figure 3.10. As above, I report results for both variables in the same figure, and only report coefficients for regressions including the complete set of control variables.

On the one hand, Panels 1 and 2 of Figure 3.10 report the specification with route-specific trends. I no longer find any statistically significant effect. However, this result may simply be due to the inability to include clustered standard errors into this specification. The route-specific pre-trends, on the other hand, confirm the earlier result. Although the common trend assumption does not appear to

Figure 3.10
Coefficients of the (treated$_{ij}$ × week$_t$) interaction term (route-specific trends)

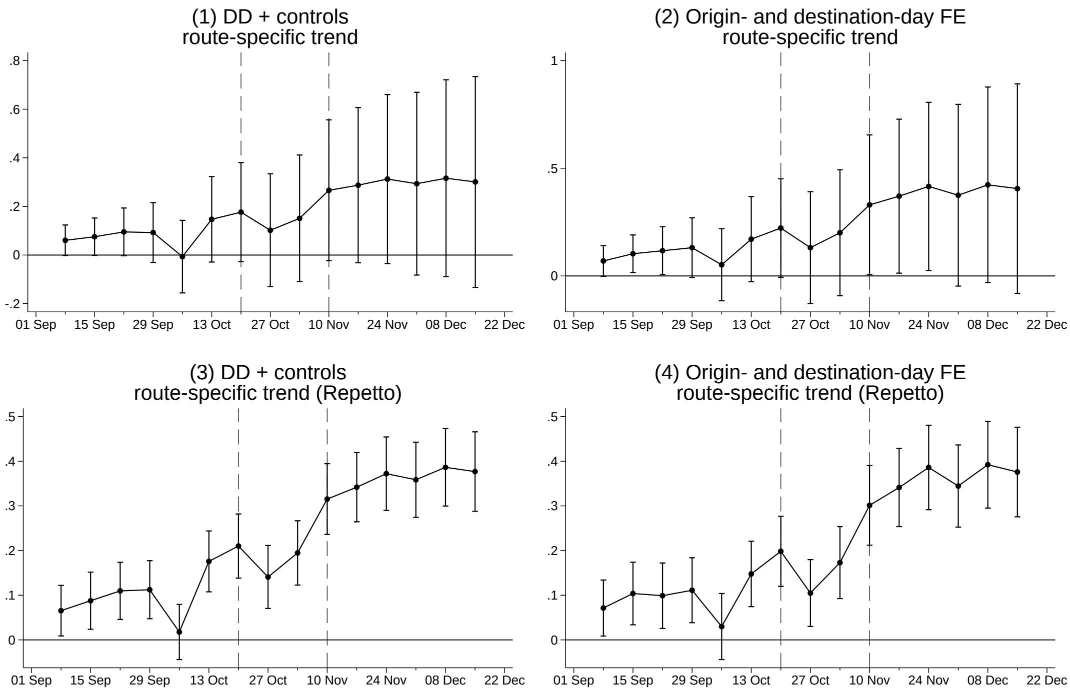

Notes: Dashed vertical line indicates weeks in which GDL was on strike. Panels 1 and 2 report coefficients from Equation 3.4 (route-specific trends, robust standard errors.) with 95 percent confidence intervals. Panels 3 and 4 report coefficients from Equation 3.6 (route-specific trends following Repetto, 2016) with 95 percent confidence intervals. Panel 1-2: robust standard errors. Panel 3-4: route clustered standard errors. 166 clusters.

be completely tenable in the given context, the lasting and remarkably persistent post-treatment effects for the treated routes is still visible.

3.5.2 Other robustness

In this subsection, I consider a host of additional factors, alternative specifications, and different definitions of the dataset to verify my previous results. For length, all regression tables are reported in Appendix C.3.

First, I conduct a robustness check with treatment defined as a continuous variable. A continuous 'treatment' is harder to interpret, but captures more variation in the channel variable. The regression results with the explanatory variable specified as the natural logarithm of bus travel time yield statistically significant coefficients equivalent to my previous results. In addition, I confirm that using variable *absolute travel time difference* for the post-strike regressions yields equivalent results.

Second, GDL membership rates are higher in East Germany because many train drivers in West Germany are civil servants. Travellers may not have been aware of the precise emergency timetables, and simply considered the effect of GDL strikes to be starker in East Germany. In that case, the relevant transmission channel would be to split routes into West- and East- Germany. Note that this specification does not allow for the inclusion of origin-day and destination-day fixed effects. I do not find that using this distinction explains MFB ticket sales during the strikes.

Third, I re-run my estimation with Berlin omitted from the sample. Berlin is special because inter-city buses were liberalized before 2013 – a historical relic from the Cold War division of Germany. My results are unaltered if I drop all routes to and from Berlin.

Fourth, I re-run my estimation of the effect *during* the strikes using ticket sales to all consumers as the dependent variable. While the estimated coefficients are lower, this change does not alter the previous results in a meaningful way. Bus travel time is the only factor that significantly explains MFB ticket sales during the strike. The same holds true if I do not drop the two days before and after each rail strike and include the intermediary week between the second and third strike wave.

Fifth, long routes are more likely to be served by aeroplanes. Customers may have switched to buses on routes with a short bus travel time because aircraft do not serve these routes. In this case, the short bus travel time would not be a proxy for closeness of substitution, but lack of other alternatives to rail. To address this concern, I show that my main results is insensitive to a re-run of

FIGURE 3.11
Mean fraction of tickets sold with a discount by treatment and control group

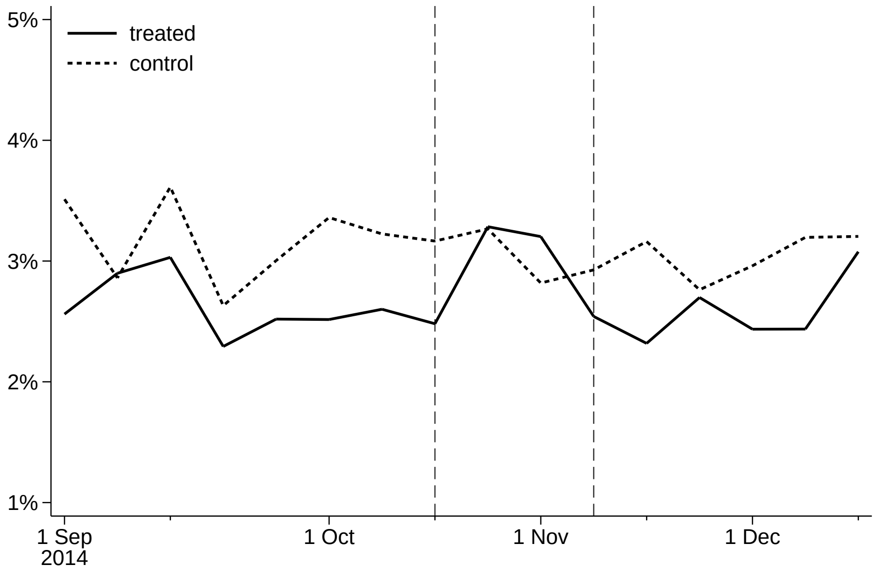

Notes: Vertical lines indicate weeks in which GDL was on strike.

my estimation where I restrict the sample to routes with no substantial national flight service.[26] During the sample period Germany's largest airline *Lufthansa* was also affected by strikes due to a labour dispute with its pilots. While an airline strike would primarily affect long bus routes, this robustness check also addresses spillover concerns from *Lufthansa* strikes.

Sixth, a concern might be the presence of unobserved marketing activity by MFB. A marketing campaign may have coincided with the rail strikes and targeted routes with a short bus travel time. While I do not have data on MFB's marketing budget, my dataset includes information on whether MFB sold a ticket at a discount. For example, MFB may have handed out vouchers or offered discounts via its mobile phone Application. Using discounts as a proxy for MFB marketing activity, Figure 3.11 plots the mean fraction of tickets that received a discount for each departure day split by treatment and control group. The fraction of tickets that receive a discount fluctuates between 2 and 4 percent in the sample period. Based on this

[26]To be precise, I drop the largest 10 bi-directional connections (20 routes) within Germany. This covers all city connections with an excess of 0.4 million annual passengers in 2016. Source: ADV Airport association

proxy measure, there is neither evidence that MFB increased its marketing activity in general, nor for the treatment group.

Seventh, an additional concern might be that travellers booked bus tickets after the November 2014 rail strike, because they were worried about potential future strikes. The rail strikes lasted beyond the strikes in 2014, and the labour dispute was only resolved after additional strike waves in April and May 2015. However, immediately after the strike wave in November GDL announced a temporary truce. It would refrain from industrial action until the new year. Even though some customers may have distrusted the truce, it is unlikely that increased bus ticket sales in this period are driven by the fear of new strikes.

Eighth, a further concern might be that many travellers are locked in to DB because they hold season passes. While DB does not offer season passes, it operates the *BahnCard* – a frequent traveller card granting fixed price reductions. More than half of all DB ticket sales receive discounts through the *BahnCard*.[27] Travellers may have waited for their *BahnCard* to expire before they switched to inter-city buses. While it is possible that any effects may not show up until later, my period-specific treatment effects suggest an immediate impact.

Finally, my dataset permits me to observe return ticket bookings. I confirm that my results are not sensitive to the inclusion of return tickets bought in a single booking session.

[27] Source: Welt.de (url: *https://www.welt.de/wirtschaft/article1069965/Die-Bahncard-hat-Verspaetung.html*; 31/07/2007)

3.6 Conclusion

This chapter exploits a novel and extremely rich dataset to investigate the effects of the 2014-2015 German railway strikes – the largest in German history – on the domestic demand for inter-city buses. The railway strikes provide a quasi-natural experiment setting to analyse the general question of whether a temporary shock can have lasting effects on one's competitors' demand.

I first test a number of potential transmission channels for inter-city bus demand, since it is not ex-ante clear which bus routes were affected *during* the rail strike. The results show that the only channel predicting peak ticket sales for MeinFernbus during the rail strikes is the closeness of substitution to the rail. Customers switched to inter-city buses if the absolute travel time difference was small or the absolute bus travel time was short. There is no evidence that travellers took into account the regional variation of the exposure to the rail strike, as measured by the fraction of cancellations and expected delay, in their decision on whether to switch to buses or not. Either they were not well informed about their exposure to the strikes as specified in the emergency timetables published by DB, or they may simply not have trusted DB's ability to implement the emergency timetables. In a second step, I use the channel identified in the first step, to test whether the strikes brought about lasting changes *after* DB services returned to normal operations. Although the common trend assumption does not seem to be completely tenable in the given context, my results still suggest a lasting effect on the ticket sales for inter-city buses on the affected routes. This result is robust to a number of alternative specifications, such as the inclusion of route-specific pre-trends.

The findings of this chapter open questions for future research. Given the history of interaction between GDL and DB, future rail strikes are very likely. Since the inter-city market for buses has consolidated substantially since 2014, future research may be able to remove seasonal effects and establish a stronger causal effect for the persistence of rail strikes on bus demand. Another intriguing avenue would be to uncover the potential mechanisms at play. These may range from new information asymmetries to retaliatory motives.

Appendix A

Appendix to Chapter 1

A.1 Data

The main source we rely on to obtain bilateral trade flows is the standard United Nations Commodity Trade Statistics Database (Comtrade). While a cleaned version of these data are available (Feenstra et al. 2005) we use the raw data as it gives us more years after 2000, up to 2011. We undertake some data cleaning ourselves, as described below. We verify that our main results are robust to using the Feenstra data up to 2000. We download aggregate trade data.[1] Our original sample of annual aggregate trade flow contains 32,386 observations reported as imports from 47 European economies over the period 1990 to 2011. The year 1990 marks the fall of the Iron Curtain and 2011 is the most recent year for which a full set of reported trade statistics is available. We use the 4-digit Standard International Trade Classification, revision 2, commodity code (SITC2) as it is the most detailed product classification for which the Comtrade database offers data spanning back to 1989, and it is the same as used by Feenstra et al. (2005). Individual

[1]Comtrade data are revised over time. The data described here were accessed on June 23, 2013 via the website http://comtrade.un.org.

observations are identified by origin-destination-year dimensions. Table A.1 lists all countries in the dataset.

The first problem we encounter is that of missing reported trade values. These are especially common in early years after a break-up or creation of an economy in the aftermath of the fall of the Iron Curtain. For example, Slovakia only starts reporting its trade flows in 1994, one year after the break-up of Czechoslovakia. Following the approach taken by Feenstra et al. (2005) we prefer importer reported statistics, assuming these are more accurate than those trade values reported as exports. Wherever possible we use exporter reported trade flows if the import reported trade flows is missing for a pair of countries. By this method we replace 2,293 missing observations in the total trade dataset - about ten percent of observations.

Within Comtrade, import reported data is valued CIF (cost, insurance and freight) and export reported data is valued FOB (free on board). FOB-type values include the transaction value of the goods and the value of services performed to deliver goods to the border of the exporting country. CIF in addition includes the value of the services performed to deliver the goods from the border of the exporting country to the border of the importing country. Following the methodology of HMR we correct this discrepancy by discounting CIF values by 10 percent. We compare the import and exported reported trade statistics whenever both reports are available. If we ignore all exporter and importer reported values that differ by a factor of greater than two either way, we find that reports valued as CIF exceed FOB reported values by a factor 1.12 on average, which confirms the HMR methodology.

For the product level regressions we rely on the BACI dataset from CEPII.[2] BACI provides bilateral values and quantities of exports at the HS 6-digit product disag-

[2] See Gauillaume and Zignago (2010).

gregation, for more than 200 countries. BACI data are available from 1995 only, we are grateful for CEPII to provide us with data from 1992. We cannot include 1990 and 1991 in our product level analysis. Services are not included in this dataset, and thus services is the HS 2-digit category that we do not include in the analysis. It is the only omitted category.

We use UN definitions (2013) to determine which countries to include as Europe. We start with all European countries, but undertake some aggregations to balance the data. Some of the nation break-ups following the fall of the Iron Curtain occur within key economies of the former Habsburg Empire. We prefer to work with a panel of stable country boundaries so that compositional differences do not drive our results. Fortunately these border changes consisted of splits in such a way that they can easily be mapped into larger units that remain stable over time. We aggregate trade flows to the smallest possible country which we can observe continually over the sample period. Table A.1 lists all country groups and years that merge/split and that we aggregate. After aggregating we drop within country trade (i.e. trade flows that were formerly reported as Czech Republic to Slovakia). Note that we only observe trade statistics from the Former Yugoslav Republic of Macedonia starting in 1993. Usually Comtrade country borders changes only occur at the beginning of a calendar year. There is one notable exception to this: both Serbia and Serbia-Montenegro report trade data in 2005. We keep and aggregate these observations within the same year as it might be due to Serbia-Montenegro breaking up at some point during the year, such that Serbia starts reporting its imports from some month when Serbia-Montenegro ceases to do so. Consequently, our measure of Yugoslavia contains reports from former Yugoslavia in 1990-1991, reports from four countries in 1992, five countries from 1992 to 2004, six countries in 2005 where both Serbia and Serbia-Montenegro report data, and six countries from 2006 and thereafter as Montenegro replaces Serbia-Montenegro. We drop a number of countries that belong to the former Soviet Union from the

Table A.1
List of European Economies and our aggregation method

Albania	Fmr Yugoslavia	Poland
Andorra*	France	Portugal
Austria	Germany	Rep. of Moldova**
Belarus**	Gibraltar*	Romania
Belgium***	Greece	Russian Federation**
Belgium-Luxembourg	Vatican City State*	San Marino*
Bosnia Herzegovina***	Hungary	Serbia***
Bulgaria	Iceland	Serbia and Montenegro***
Croatia	Ireland	Slovakia***
Czech Rep.***	Italy	Slovenia***
Czechoslovakia	Latvia**	Spain
Denmark	Lithuania**	Sweden
Estonia**	Luxembourg***	Switzerland
Faroe Isds*	Malta	TFYR of Macedonia***
Finland	Montenegro**	Ukraine**
Fmr Dem. Rep. of Germany***	Netherlands	United Kingdom
Fmr Fed. Rep. of Germany***	Norway	

Notes: Trade values estimated following the methodology of Feenstra et al. (2005). * Only appear as partner, not included as reporter country as trade and production data unreliable. Do not report trade statistics themselves. ** Former Soviet Union with changing borders. *** Aggregated with another country to balance the sample.

dataset (Belarus, Ukraine, Latvia, Lithuania and Estonia as well as the Russian Federation). With the dissolution of the Soviet Union these countries and the consequent political turmoil these economies only appear in the trade statistics two years after the beginning of the sample period (in 1992). We decide that the cost of introducing noise by including them is greater than the benefit of gaining some more observations, especially as these countries are not directly relevant for the question we study. We omit tiny countries such as the Vatican and the Faroe Islands, but include them as partner countries. Given these changes, the resulting panel of countries we work with is balanced throughout all the years we study.

We drop reported destinations that are designated ‘bunkers’ (UN code 837), ‘free zones’ (838), ‘special categories’ (839) and ‘areas not elsewhere specified (nes)’ (899). Moreover, we drop the highly incomplete observations reporting trade with San Marino, the Vatican, Andorra, Faroe Islands and Gibraltar. Table A.2

Table A.2
Aggregated Economies

Country	Years observed
Germany	
Germany	1991 - 2012
Fmr Dem. Rep. of Germany	1989 - 1990
Fmr Fed. Rep. of Germany	1989 - 1990
Czechoslovakia	
Czechoslovakia	1989 - 1992
Czech Rep.	1993 - 2012
Slovakia	1993 - 2012
Yugoslavia	
Fmr Yugoslavia	1989 - 1991
Slovenia	1992 - 2012
Bosnia Herzegovina	1992 - 2012
Croatia	1992 - 2012
TFYR Macedonia	1993 - 2012
Serbia and Montenegro	1992 - 2005
Serbia	2005 - 2012
Montenegro	2006 - 2012
Belgium-Luxembourg	
Belgium-Luxembourg	1989 - 1998
Belgium	1999 - 2012
Luxembourg	1999 - 2012

reports the elements by year for the countries that involve aggregation for our dataset.

We add a number of standard control variables, relying on standard sources. We obtain data on aggregate GDP and populations from the World Banks World Development Indicators (2013). We compute GDP per capita as GDP divided by population, both as reported by the UN. Following our methodology of aggregating trade flows, we derive GDP and population measures for Yugoslavia and Czechoslovakia as the sum of GDP and populations of the underlying countries. For example, Czechoslovakia's population is calculated as the sum of the Czech Republic's and Slovakian populations. GDP is measured in current US dollars (millions) and, in accordance with trade flows, not deflated. We obtain a number of gravity variables from the CEPII distance database used in Mayer and Zignago

(2005).[3] These include the country-specific variable landlocked as well as dyadic variables. These are common border, common (official) language, shared language spoken by at least 9 percent of the population, and distance. As measure of distance we use distance between capitals as it is a consistent measure we can apply to the aggregated economies. For example, we use Prague as the capital of Czechoslovakia throughout the sample period. The variables time difference, shared legal history, area and shared religion are from the gravity data set provided by HMR.[4] We also use this source to add time varying variables GATT/WTO membership, membership of RTAs (Regional Trade Agreements) and a common currency indicator. Since the HMR dataset only spans the years up to 2006, we update the time varying variables using data from the WTO.[5] Finally, we construct dummy variables for EU and Eurozone membership.[6] This latest source also allows us to generate a variable that indicates membership in the common currency.

[3] These data are available at http://www.cepii.fr/anglaisgraph/bdd/distances.htm (accessed 19/06/2013).

[4] These data are available at http://strategy.sauder.ubc.ca/head/sup (accessed 19/06/2013).

[5] Here we rely on two sources, http://www.wto.org/english/thewto_e for GATT/WTO membership and http://rtais.wto.org/UI/PublicPreDefRepByEIF.aspx for RTAs (both sites accessed 19/06/2013).

[6] We use the EU web site http://europa.eu/about-eu/countries/index_en.htm (accessed 10/07/2013)

A.2 Robustness and additional results

We verify that our results are robust to a number of alternative specifications and estimation methods. We omit the detailed numbers and figures for some of these robustness tests for reasons of space. Details on all robustness checks not displayed here are available upon request.

Habsburg definition

We define the Habsburg measure in different ways. We include all countries that are at least partly former Habsburg members, thus adding Italy, Poland, Romania, Serbia and Ukraine to the countries covered by the Habsburg fixed effects. The Habsburg coefficients remain fairly similar, yet become somewhat statistically weaker. This is as expected, given that this measure includes areas that were outside of the monarchy and thus should add more noise than signal. We run a separate regression including only Yugoslavia as an additional Habsburg member, and one in which we code Yugoslavia as being west of the Curtain. Yugoslavia is an ambiguous case given its unique history during the 20^{th} century. The monotonic downward slope is strongly robust to these specifications and variations thereof.

Panel estimation

We address the concerns brought forward by Anderson and Yotov (2012), that a disadvantage of pooling gravity data over consecutive years is that dependent and independent variables cannot fully adjust in a single year's time. We address this concern using the suggested methodology of keeping only intervals of 3 or 5 years. The downward slope in Panel 1 in Figure 3 of chapter 1 becomes -.038 (.004) when keeping only every third year from 1990, and -.034 (.002) when keeping only every

Figure A.1
Anderson-Yotov 5 year intervals

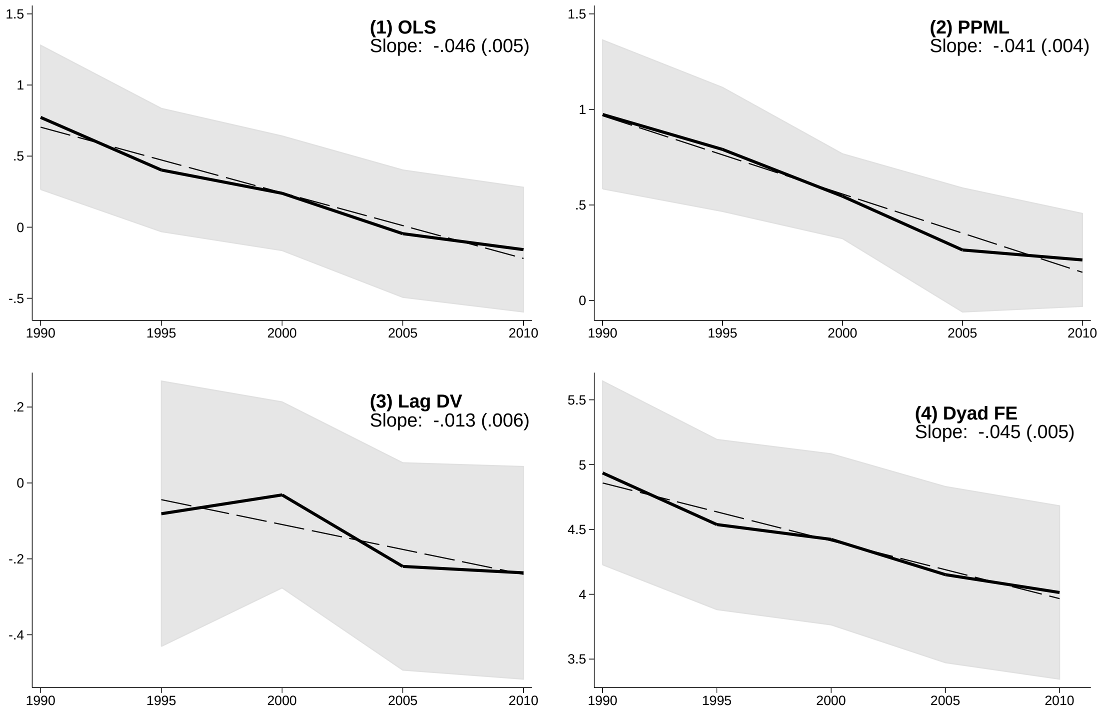

fifth year. Our findings seem not to be much changed by this adjustment. See also Figure A.1.

Additional control variables

Country pairs that are most affected by the fall of the Iron Curtain are those country pairs that shared an East-West border along the Iron Curtain, such as Austria-Czechoslovakia, Italy-Slovenia, Greece-Albania. To avoid a potential omitted variable bias we include time varying control variables for these country pairs. We continue to observe the downward sloping Habsburg coefficient with a fairly similar magnitude. We also include a dummy variable indicating that both countries are west of the Iron Curtain. The slope in our preferred estimation remains numerically at -0.044, and is not statistically significant from any of our estimated slopes at the 1% level of statistical significance. We also add measures of cultural

Figure A.2
Additional control variables: countries that share an east-west border (time varying), genetic distance, Eurovision voting, cultural and religious similarity.

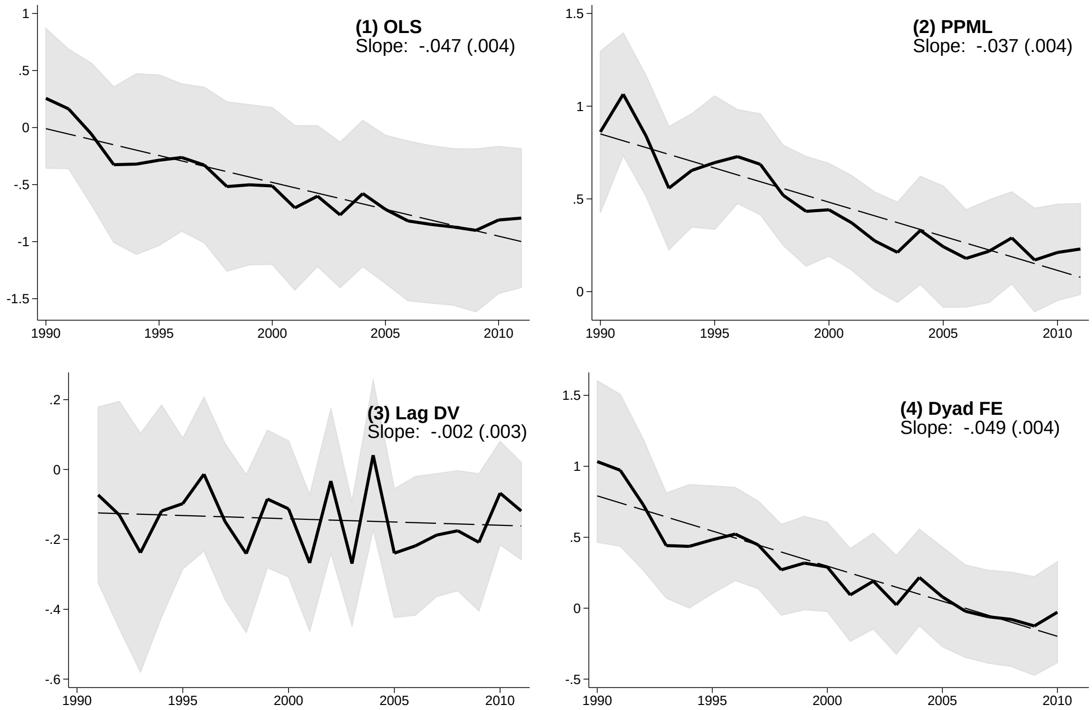

proximity. These are variables indicating Eurovision voting preference (data from Felbermayr and Toubal, 2010), genetic distance (data from Spolaore and Wacziarg, 2009) and cultural and religious similarity (data from Spolaore and Wacziarg 2015). The inclusion of these additional control variables does not change the slope in our main estimation, which remains at -0.047, close to the estimate in our main specification. These and other time invariant bilateral variables we may have omitted are covered by our specification that includes bilateral fixed effects. These results are also reported in Figure A.2.

Placebo

We run a couple of placebo exercises where we replace Austria by other Western European countries. These address the possibility that the opening of trade relations between East and West might be dynamic, increasing or decreasing, in the first years after the opening of the Iron Curtain because of various reasons other than the decline of historic and cultural ties. For example, the installation or reuse of transport infrastructure might suggest a dynamic trade relationship between an eastern and a western country, or the slow establishment of personal exchange and interaction. In both these examples we would expect an increasing relationship, but there may be others. To mitigate concerns that such effects drive our results we run a placebo exercise in which we estimate 'Habsburg' effects on a relationship other than Habsburg, for which we do not expect the same decay of cultural ties. We chose Germany as our preferred placebo country. It shares the language with Austria and also a direct border with many eastern countries. When we estimate the trading relationship with Germany instead of Austria being the 'Habsburg' country west of the curtain, we do not find significant relationships. These results are reported in Table A.3, and in this table we use the same specification as applied in Tables 2 and 3 of chapter 1. The PPML estimates display an increase of the effect for intermediate years, which may point to some form of catch up in the interim years. This effect however shows no monotonic trend in t and is not robust to the other specifications displayed. Most of the coefficients in Table A.3, including in the PPML specification are not statistically significant. We interpret this finding to cast doubt on the relevance of other dynamic effects shaping initial trade relationships. We also estimate further robustness tests using Switzerland, the Netherlands, Belgium-Luxembourg and Italy as placebo countries. Switzerland is similar to Austria in some cultural aspects and geographically close, yet does not share the Habsburg history. It also does not have a history of division and unification like Germany. The Netherlands and Belgium-Luxembourg are other countries similar in size and wealth to Austria. Italy is geographically close to both Austria and the Iron Curtain. As in the estimation with the German placebo, these estimations are exactly like our main estimation for Habsburg

Figure A.3
Further Placebo regressions.

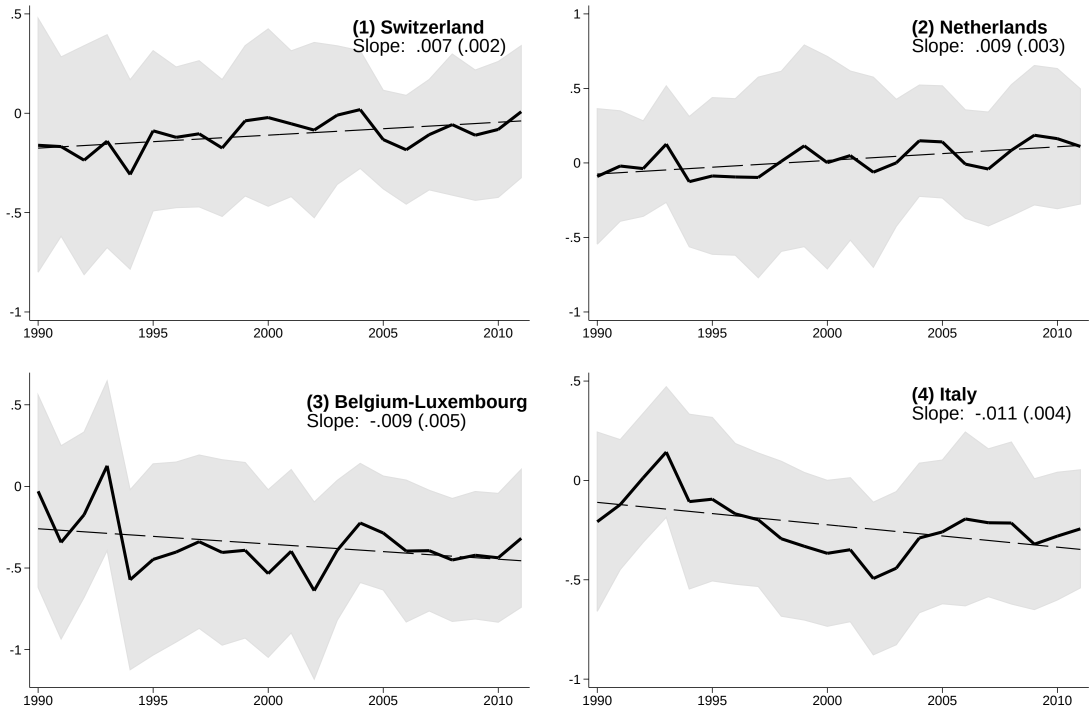

countries with the exception of replacing Austria by each of the placebo countries in turn. We display the main OLS estimation for these four countries in Figure A.3. The magnitudes of the slopes in all these four cases are much lower than for the main Habsburg specification. Switzerland and the Netherlands show a slightly positive trend with small magnitudes. Only Italy shows a negative slope, moderately significant and also small in magnitude. Italy is the only country of the four that partly had a Habsburg history herself, so the one placebo country where we might have expected a small negative slope.

Table A.3
Germany Placebo Coefficients

	(1) OLS	(2) PPML	(3) Lag DV	(4) Dyad FE
Dependent variable:	$\ln(x_{int})$	x_{int}	$\ln(x_{int})$	$\ln(x_{int})$
1990	-0.230	0.342		-0.130
	(0.375)	(0.225)		(0.238)
1991	-0.287	0.113	-0.213**	-0.278
	(0.285)	(0.213)	(0.0981)	(0.181)
1992	-0.140	0.196	0.0853	-0.0514
	(0.294)	(0.171)	(0.0944)	(0.175)
1993	0.106	0.431***	0.228***	0.186
	(0.286)	(0.167)	(0.0809)	(0.162)
1994	-0.158	0.358**	-0.227	-0.110
	(0.318)	(0.142)	(0.196)	(0.155)
1995	-0.0570	0.317*	0.108	-0.0191
	(0.346)	(0.180)	(0.0817)	(0.150)
1996	-0.0678	0.304*	-0.0319	-0.0151
	(0.307)	(0.184)	(0.0632)	(0.138)
1997	-0.00333	0.395**	-0.000351	0.0679
	(0.296)	(0.183)	(0.0804)	(0.132)
1998	-0.0299	0.490***	-0.0406	0.0433
	(0.291)	(0.177)	(0.0752)	(0.141)
1999	-0.00454	0.506***	0.0522	0.104
	(0.313)	(0.177)	(0.0796)	(0.137)
2000	-0.0777	0.416**	-0.0934	0.0192
	(0.330)	(0.178)	(0.0848)	(0.143)
2001	-0.0327	0.460***	0.0385	0.0688
	(0.305)	(0.170)	(0.0572)	(0.134)
2002	-0.0519	0.530***	-0.0353	0.0493
	(0.329)	(0.158)	(0.118)	(0.169)
2003	0.0254	0.544***	0.0483	0.133
	(0.274)	(0.144)	(0.0480)	(0.138)
2004	0.0509	0.462***	0.0112	0.160
	(0.263)	(0.159)	(0.0753)	(0.133)
2005	-0.0569	0.316*	-0.106	0.0521
	(0.281)	(0.189)	(0.0753)	(0.136)
2006	-0.115	0.268	-0.0585	-0.00521
	(0.310)	(0.184)	(0.0903)	(0.139)
2007	-0.145	0.214	-0.0530	-0.0417
	(0.287)	(0.175)	(0.0634)	(0.134)
2008	-0.183	0.154	-0.0743	-0.0802
	(0.288)	(0.172)	(0.0656)	(0.136)
2009	-0.156	0.0905	-0.00779	-0.0530
	(0.291)	(0.166)	(0.0813)	(0.143)
2010	-0.147	0.0673	-0.0296	-0.0469
	(0.291)	(0.166)	(0.0813)	(0.143)
2011	-0.102	0.102	0.0114	
	(0.323)	(0.170)	(0.103)	

Notes: Placebo exercise: Habsburg coefficients with Germany instead of Austria. Stars denote statistical significance on the level of one (***), five (**) and ten (*) percent. Robust standard errors used.

Figure A.4
Aggregate Eastern countries.

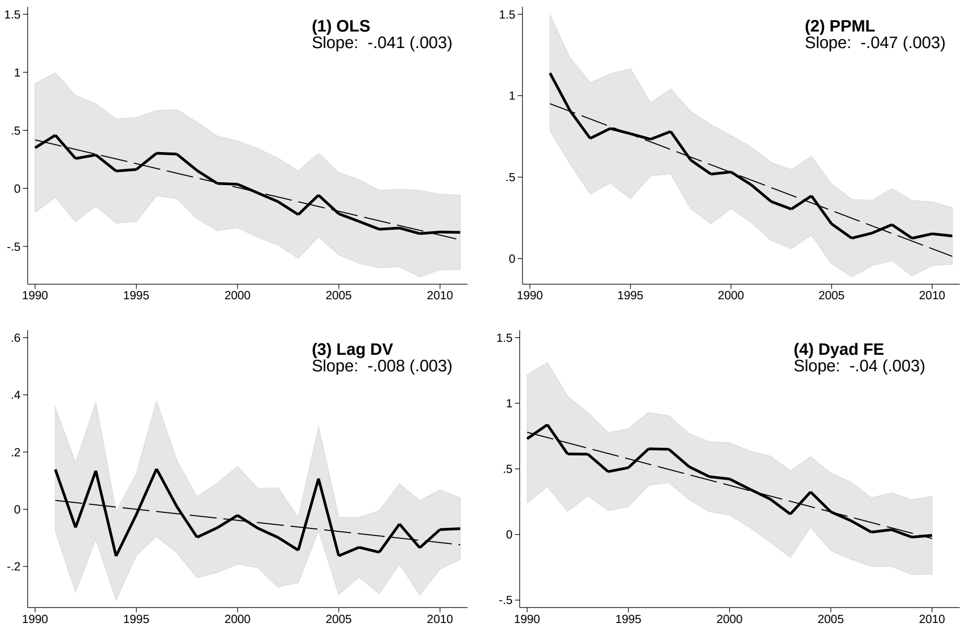

Aggregation

We test the sensitivity of our results to the choices of aggregating countries we make. We go to both extremes, by creating the most disaggregated and the most aggregated unit we can. In the most aggregated version we add all countries east of the Iron Curtain that were part of Habsburg into one observation, such that the dummy of interest becomes the bilateral flow between one Eastern and one Western aggregate. Despite the small sample of treatment in this robustness, which is just one bilateral trade flow between Austrian and the Eastern Aggregate, a strong, significant downward slope remains, although somewhat smaller in absolute magnitude than in the main specification (see Figure A.4). In this Figure we do not include Yugoslavia. When we do include Yugoslavia in the eastern aggregate we find a very similar picture, with a strong, significant downward slope of -0.033 (0.003) in our main specification.

Figure A.5
Disaggregate Eastern countries.

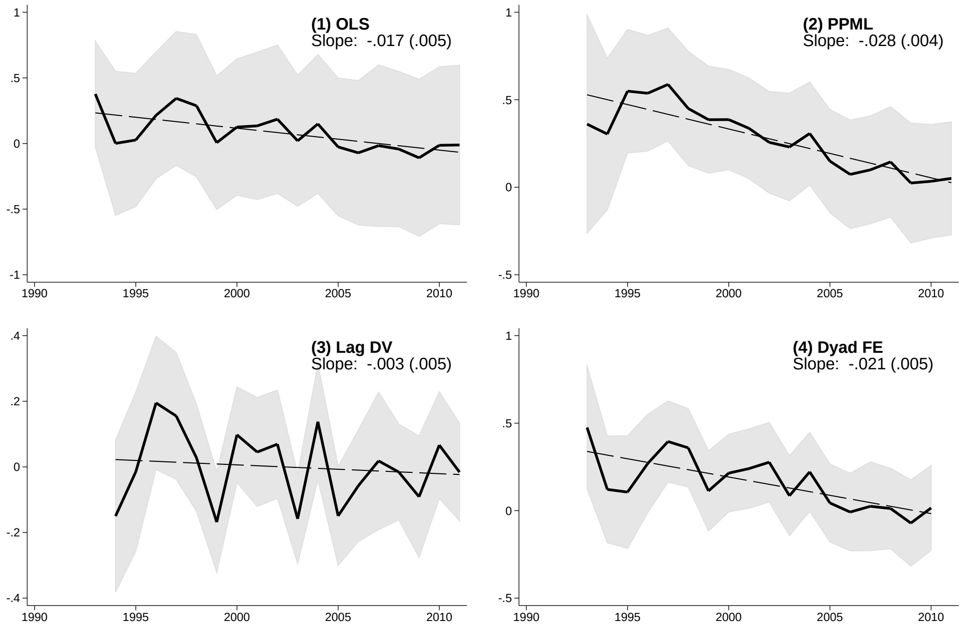

In Figure A.5 we repeat the main table for the most disaggregated version of countries, which splits Yugoslavia and Czechoslovakia into its components today from the moment of separation. This exercise is only possible for the years from 1993 onwards, when countries were separated. The downward slope remains strong and negative in all four specifications. The turbulent history of the countries of former Yugoslavia and the corresponding big shocks to their trade relationships are likely to contribute to the increased noise apparent in this graph compared to our main specification.

Pre-1990 trend

The timing of the surplus trade is mainly observable after 1990, when the countries fully integrate in the European market and rich trade datasets from the standard sources become available. It is interesting for our conclusion however to see what the pre-1990 trend looks like. A concern might be that Austria's special surplus that we observe in the trade volumes with Eastern Europe had been large and

Figure A.6
Pre-1990 trade (Datasource: Barbierei, 2002)

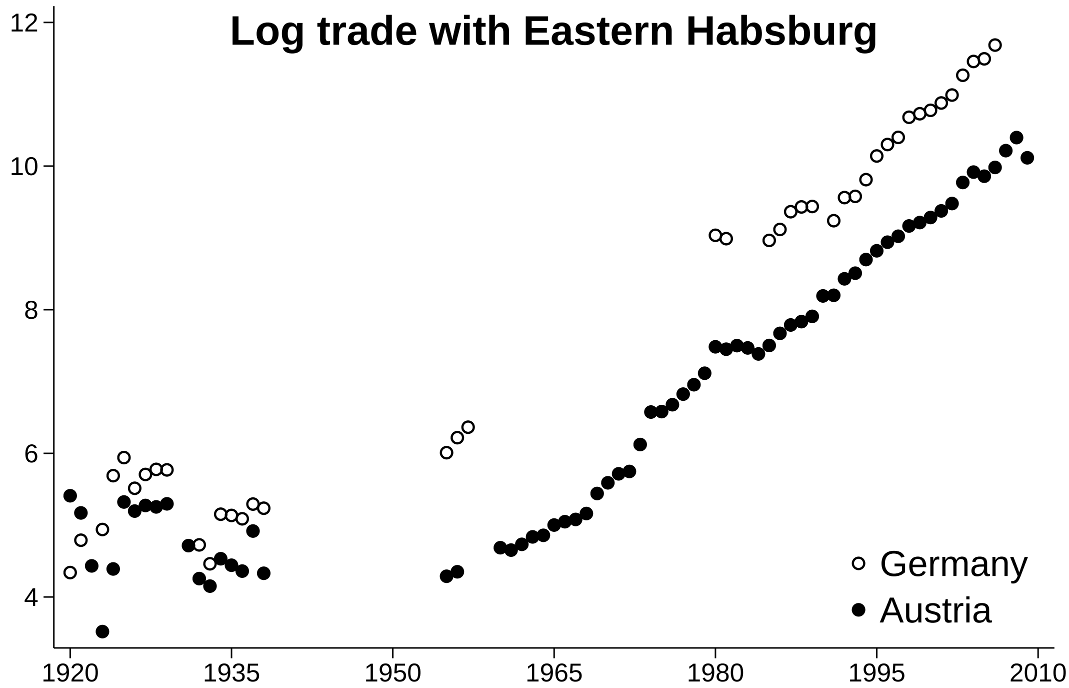

built up before 1990. It may have also been that the built up had been temporary due to the suspicion that the Iron Curtain may break with Austria first.

It is not straight forward to obtain trade data for these countries and time periods. To get some evidence we use the data by Barbieri (2002). These data cannot be directly compared to the Comtrade data we rely on in the rest of chapter 1. Some industries are missing, and the trade levels for flows we can compare seem much lower than those reported in Comtrade. They also do not match well with the numbers reported in the history section (Butschek, 1996; Lazarevic, 2010; and Pogany, 2010). Thus, while we do not want to over interpret the trade levels in this dataset, we think that they are informative regarding trends, just as Barbieri used data from unified sources to analyse trade developments over time. As the sources from the history section make clear, these trends take place at very low levels of trade.

Given concerns over the levels of these data, missing data and lack of availability we do not run a full gravity model here, but just report the relative trade levels of Germany and Austria with respect to Hungary and Czechoslovakia before 1990 in Figure A.6. This graph displays a few helpful facts on trade before 1990: Firstly, Austria and Germany display a similar trend of growth with the Eastern Habsburg Countries from 1950 to 1990. Thus, Austria seems not to particularly build up the surplus trade in the years before 1990. Secondly, Austria seems not to have had much of an advantage over Germany during the years of separation; while Germany's trade with the East grows from 1935 to 1955, Austria starts at levels slightly below the average from 1920-1940, and Austria never catches up to the German advantage during the entire century. (On average, Germany's trade level is about four times that of Austria in the period 1980-2010). Thirdly, these numbers also confirm the claim that the fall of the Iron Curtain had not been widely anticipated, and appears as a surprising event in these trade numbers. Fourthly, there may be some suggestion that the Austrian surplus slowly builds up from 1955 to 1990, and then declines again. Yet given the limitations of this exercise, and the data involved we urge the reader to treat these suggestions with caution.

Missing data

In the main parts of chapter 1 we treat missing trade flows in the CEPII or Comtrade data as a zero trade flow and include such observations in the analysis. We test the robustness of our results to treating zeros in different ways. First, we omit zeros from the sample or second, we replace zeros by 1. Again, our main conclusions in the main results table do not seem to be altered by these specifications. These results are available on request.

Figure A.7
Include internal trade.

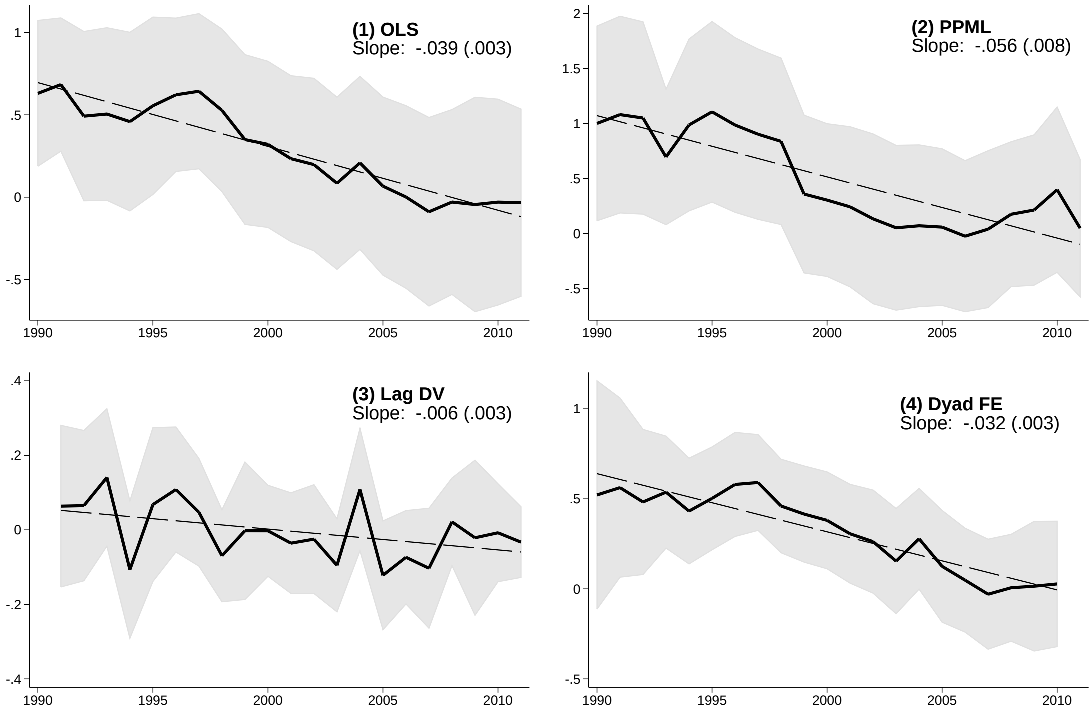

Internal trade

Yotov (2012) argues that gravity models should not just focus on international flows but explicitly take into account national integration, internal distance and internal trade costs. We follow the gravity literature on how to construct internal trade flows. We construct internal trade flows as the difference between GDP from the World Bank (World Development Indicators) and total exports from our Comtrade dataset. Related recent contributions (for example, Heid et al. 2015) use production data from the UN's Industrial statistics database at the industry level. This measure, however, suffers from significant missing observations in the early 1990s. Therefore, because our specification focuses on aggregate trade flows, we prefer using GDP data. Regarding the coding of covariates for internal flows, we adopt the measures suggested by the CEPII. These standard definitions also include suggested controls for internal trade flows. These are then consistent with our previous controls. For example, the population-weighted distance allows for a consistent use of internal and international distances. See Figure A.7 for the results of this exercise. As shown in Figure A.7 the internal trade specification does not change the slopes or picture much compared with the main estimate. We would have expected the results to be weaker, as we count the internal trade flow of Habsburg countries as part of the treatment effect, but don't expect that portion of the trade to be important. In other words, we add noise to the treatment. And indeed the OLS slope is slightly smaller in magnitude than in the main results specification. This decrease is neither very pronounced nor robust across all specifications.

Standard errors

As an alternative treatment of standard errors, we conduct a robustness test in which we cluster standard errors by bilateral country pairs. Coefficients remain identical, we verify that this does not change the significance of coefficients reported in Figure 3 in a meaningful way. Results indeed remain strongly significant. It should be quite apparent from the monotonic downward slope visible in that figure that the significance of this downward slope is strongly robust to other even more demanding specifications.

Appendix B

Appendix to Chapter 2

B.1 Matching RASFF rejections with HS4 product codes

One of the contributions made by chapter 2 is the method developed to assign product codes to the verbal descriptions provided for notifications on the RASFF portal.[1] Attributing product codes is a prerequisite for matching sanitary rejections with Chinese export data.

To assign a product code to each notification, we exploit information on *product category* (e.g. 'alcoholic beverages') and *subject* (e.g. 'undeclared sulphite in Wine from Chile') reported by the RASFF authorities. We assign observations to the HS classification in which our Chinese firm-level data are coded. We code to the HS 4-digit level - the most disaggregated level at which we can identify notifications. We use the 2002 revision of the HS classification.

A manual assignment of HS4 codes on an individual basis is not possible given the number of notifications in our database (14,860 observations for the period 2000-

[1] url: *http://ec.europa.eu/food/safety/rasff/index_en.htm*

2011 after the cleaning procedure described in Section 2.3). Therefore, to assign product codes we implement the following approach. We first split *subject* in order to extract the relevant information on the product (e.g. 'wine'). Next, we rearrange some *product* categories and align them more directly with HS2 sectors (e.g. 'fish and fish products' and 'farmed fish and products thereof - other than crustaceans and molluscs' are combined). We also conduct some re-assignments of observations across *product* categories to ensure consistency. Finally, we disregard observations from *product* category 'food contact materials' as we are only interested in agri-food products (HS chapters 01-24).

We identify the sector (HS2) wherever possible, and assign the HS4 product code using Stata's *regexm* function. *Regexm* searches for keywords associated with a specific HS4 code. For example, within *product* 'fish', 'frozen hake fillets' can be assigned HS4 code 0304 ('Fish fillets and other fish meat - whether or not minced, fresh, chilled or frozen') using keywords 'fillets' and 'frozen'. Using the same method "chilled hake" is assigned HS4 code 0302 ('Fish, fresh or chilled, excluding fish fillets and other fish meat of heading No 0304'). The full Stata do-files with the matching correspondence and code mapping RASFF notifications and HS codes are available on request from the authors.

This methodology has several advantages. Firstly, it is easily checked, verified, and replicated, and ensures consistent treatment of RASFF observations. Secondly, it can be extended to more data at a very low cost. For example, it can be applied to additional observations as more RASFF notifications become available over time.

Using this strategy, we successfully match 89 percent of rejections with an HS4 code (13,241 out of 14,860). Among border rejections applied to China we match 91 percent (1,537 out of 1,690). The incidence of rejections is fairly heterogeneous across products but is clustered in some sectors. Our rejections are split over 115

Table B.1
Chinese border rejections and percent of agri-food exports by HS2 (2000-2011)

HS Chapter	% Chinese agri-food exports	Nb. of rejections
01 Live animals	0.1	0
02 Meat and edible meat offal	0.6	32
03 Fish and crustaceans, molluscs	24.5	258
04 Dairy produce	0.9	59
05 Products of animal origin	9.9	40
06 Live trees and other plants	0.6	0
07 Edible vegetables	11.2	75
08 Edible fruits and nuts	4.6	24
09 Coffee, tea, mate and spices	2.9	76
10 Cereals	0.2	65
11 Products of the milling industry	0.2	0
12 Oil seeds and oleaginous fruits	7.4	698
13 Lac; gums, resins	1.6	2
14 Vegetable plaiting materials	0.4	0
15 Animal or vegetable fats and oils;	1.2	1
16 Prep. of meat, of fish or of crustaceans, molluscs	5.5	1
17 Sugar and sugar confectionery	1.0	26
18 Cocoa and cocoa preparations	0.6	1
19 Preparations of cereals, flour, starch or milk	2.0	87
20 Preparations of vegetables, fruit, nuts	17.7	29
21 Miscellaneous edible preparations	1.4	51
22 Beverages, spirits and vinegar	0.7	4
23 Residues and waste from the food industries	2.2	8
24 Tobacco	2.7	0

different HS4 codes out 201 potential ones in the 24 chapters of agri-food products (for China we identify 67 different HS4 products). If we look at all the rejections, the majority of notifications concern HS08 'Edible fruits and nuts', HS03 'Fish and Crustaceans, Molluscs', and HS12 'oil seeds and oleaginous fruits'. For China, HS12 and HS03 are the two main chapters affected by border rejections. We conduct an additional visual check of the mapping in Table B.1. We compare the percentage of Chinese exports and rejections by HS2 product category. While we do not expect a strong correlation (small export sectors could plausibly be affected by a disproportionate number of rejections), we are able to confirm that there are no large sectors without rejections and no tiny agri-food sectors with many rejections.

B.2 Chinese firm-level exports

Table B.2 investigates whether aggregation of the observations at the 4-digit level is a potential source of bias. If rejections occur at the HS6 product level but our analysis is performed at the HS4 sector level, we could observe automatic higher survival rates (and lower levels of exit) for larger firms. Large firms might export multiple HS6 products within an HS4 sector. Even if one firm's HS6 product is affected by rejections, other HS6 products may remain unaffected. Thus, at the HS4 level, we may observe large firms as less likely to exit the RASFF market.

To address this issue, we record the number of HS6 products exported by a firm within each HS4 sector. Table B.2 summarizes the results. Columns 1 to 5 report the fractions of firm-HS4 exports that have the underlying number of HS6 products. We observe that firms – even multi-HS4 firms – usually export only one HS6 product within each HS4 sector. 89.66 percent of firms present in only one HS4 sector export just one HS6 product within that HS4 sector (and 8.52 percent of these firms export two HS6 products within that HS4 sector). At the other end of the spectrum, for firms present in 10 or more HS4 sectors, only one HS6 product per HS4 sector is exported in 73.84 percent of the cases (and two products in 18.11 percent of the cases).

Table B.2
Percentage of HS6 products within HS4 sectors for Chinese firms (2000-2011)

Nb. of HS4	Nb. of HS6 within HS4 1	2	3	4	5+	% firms	% exports
1	89.66	8.52	1.19	.43	.19	12.69	15.48
2	86.52	10.71	1.71	.73	.33	10.2	10.96
3	84.12	12.16	2.27	.83	.62	8.42	8.85
4	82.38	13.41	2.39	1.01	.81	7.4	8.64
5	79.83	15.3	3.08	.91	.87	6.75	7.91
6	77.56	16.05	3.9	1.45	1.03	5.88	6.61
7	76.29	16.4	4.36	1.69	1.26	5.34	6.26
8	75.15	16.43	4.7	2	1.71	4.88	5.46
9	75.68	16.56	4.74	1.44	1.59	3.82	4.39
10+	73.84	18.11	4.86	1.58	1.61	34.62	25.43

Notes: Excluding wholesalers.

Appendix C

Appendix to Chapter 3

C.1 Data

MeinFernbus booking data

The route-day level bookings are constructed from an underlying dataset provided by MeinFernbus. It contains the universe of MFB bookings for all route combinations of 33 large German cities over the sample period for departure days from September 1^{st} to December 31^{st} 2014 – roughly 1.7 million observations. The dataset also includes individuals who departed in the sample period, but who booked their ticket outside the sample period.

The dataset provides detailed information on each booking such as the origin, destination, date and departure times of each service, as well as details on the individual booking process such as the time, date and whether a booking was via the web or an agency. The majority (>80%) of all bookings are made directly via the MeinFernbus website. For each booking via the internet an anonymized e-mail identifier is provided. Assuming for simplicity customer e-mails remain the same over time, this variable allows tracking individual booking behaviour over time. For agency bookings no data on individual e-mails is available. Furthermore, for each individual's e-mail the dataset records the first time a booking has been undertaken even if this was before the sample period. This allows classifying each bus customer into new and returning passengers. On the one hand, approximately 75 percent of bookings only appear once. On the other hand, about one percent of all individuals in the sample period travel regularly (more than seven times over the sample period).

In addition to the bookings, the dataset includes information on the supply of MeinFernbus services. The dataset identifies the total capacity of each bus, the line number and bus partner, as well as information on the prices charged. This allows identifying each individual journey (by bus id and route), and calculation of the total capacity of MeinFernbus buses for each departure day.

The set of routes includes all route combination of 33 large cities as depicted in Figure 3.2.[1] The cities and routes are spread across the entirety of Germany. Route selection was based on the most important cities in the bus network which approximately corresponds to the largest German cities. The choice of each city was justified based on the frequency of searches from a large online price comparison website. The data cover roughly 40 percent of the German inter-city bus market.[2] Exceptions are the exclusion of Bochum and Wuppertal as they are in the densely populated Ruhr-valley. To protect local public services, German law requires inter-city bus services to cover a minimum distance of 100km. Cities in the Ruhr-valley are frequently at a closer distance so no data on inter-city buses would be retrieved. I retain Ruhr-cities Dortmund and Essen. Furthermore, I include Freiburg because it is an important university town and Wuerzburg for its geographical centrality in Germany. Given the 33 cities in the sample there are 1056 possible routes spanning the simplex of these cities. 588 are served at least once. I focus on an even larger subset of routes: those routes that are served

[1] Note that I consider routes to be directional. For example, I treat Hamburg–Berlin and Berlin–Hamburg as two separate markets.

[2] The author thanks the team of Fernbusse.de for making data on search queries available.

almost every day; i.e. not without at least one customer for more than 10 days in the sample.

A bus station is included if it is within 15 kilometre of the city centre. If there exist multiple bus stops within one city, my dataset includes information on all offered combination of stops. However, I retain only the service between the main bus terminals. Second, I exclude origins and destinations that are airports. All airports are sufficiently outside cities that consumers are likely to prefer a bus service to the city center. Thirdly, the MFB booking data includes itineraries that involve stopovers, even though I do not observe data on these. This, however, is not a major concern. The German bus market primarily operates as a point-to-point service: the majority of passengers travel directly, meaning few connect to other buses. Buses typically have multiple stops on a line, so the travellers on a given bus may travel very different routes.

DB Emergency timetables and web-crawled itineraries

I construct a dataset of DB service cancellations and expected delays:

Emergency timetables measure the heterogeneity of different routes exposed to the rail strike. DB published emergency timetables for all inter-city (IC) and inter-city express (ICE) lines during the strikes. A route may be served by multiple rail lines and the emergency timetable only includes information on the changed frequency of each DB line (e.g. IC line 31 which operates from Frankfurt to Hamburg via Cologne usually operates every two hours but its service was cancelled entirely during the strike). However, actual travel itineraries are significantly more complex because they often involve stopovers.

To address the issue of stopovers, I gather an additional dataset using an electronic 'web crawler' linked to an online price comparison website for the week April 18-

24, 2016.[3] DB has changed timetables twice since 2014, but changes have been minor and after matching with rail lines the data are comparable to the DB service offered in 2014. The web-crawled data includes all travel itineraries for the routes of the dataset in a complete week. A travel itinerary is defined as the specific departure times, stopovers and train numbers a traveller needs to take on a rail journey.

Only the combination of emergency timetables and the web-crawled travel itineraries, allow me to construct the exposure of each route to the rail strike. Using correspondence tables of rail lines and train numbers, I match the emergency timetable data with the crawled dataset. I construct the variables *fraction services cancelled* and *additional travel time* as follows: I construct a variable measuring the trains per hour for the normal and 'treatment' (i.e. strike) period. For example, the route Hamburg–Berlin is served with 1.2 trains per hour during normal operations and 0.2 trains per hour during the strike. Multiplying these numbers by 24 gives the daily number of trains operating on the route; i.e. 28.8 trains during normal operations and 4.8 daily trains during the strikes for Hamburg–Berlin. Using these data, calculating the fraction of services that were cancelled is straightforward (i.e. 0.83 for Hamburg–Berlin). The expected additional travel time travellers have to incur to reach their destination is calculated as the time a traveller has to wait for the next train if his service is cancelled. For simplicity, I assume that the number of daily connections are evenly spaced throughout the day. For example, travellers on a route which is served by one train per hour in normal operations, and only one train every two hours during the strikes had to endure an additional travel time of one hour. I report the calculated fraction of service cancelled and additional travel time in Figure 3.3 in the data section of chapter 3.

[3] The web crawling methodology closely follows a small but growing airline literature. See Williams (2013) or Siegert and Ulbricht (2015).

One data limitation, however, remains: the DB emergency timetables do not include information on regional trains. Regional and local trains are likely to have been cancelled in a similar fashion to IC/ICE lines reflecting the local power of the GDL. Since I have no information on the disruption of regional trains, I drop all routes where more than 90 percent of all services offered involve the use of RE and RB trains. This is not a major concern, however, as the large majority of inter-city services is conducted by ICE and IC trains.

The dataset contains all trains, stopovers and travel times for the remaining routes in the sample. Using this information I construct a variable for the frequency in which each route is served per hour. For example, Hamburg–Berlin is served by 1.2 trains per hour on average, while Munich–Berlin is only served by 0.5 trains per hour.

Table C.1
Definition of variables used in Equations 3.1 to 3.6

Variable:	Definition:
Dependent variables:	
ln ticket sales_{ijt}	Log total MFB ticket sales on route ij on departure date t
ln ticket $\text{sales}_{ijt}^{\text{new}}$	Log total MFB ticket sales to new customers (NC) in the final three days to departure.
Channel variables (channel$_{ij}$):	
Fraction services cancelled	Dummy = 1 if the fraction of DB services cancelled on a route is above the median (i.e. above 63%).
Additional travel time	Dummy = 1 if the additional travel time on a route is above the median (i.e. longer than 78.5 minutes).
Relative travel time	Dummy = 1 if the relative travel time (bus travel time / rail travel time)on a route is below the median (i.e. below ratio 1.64).
Absolute travel time diff.	Dummy = 1 if the absolute travel time difference (bus travel time − rail travel time) on a route is below the median (i.e. shorter than 109.9 minutes).
Bus travel time	Dummy = 1 if the bus travel time on a route is below the median (i.e. shorter than 265 minutes).
Control variables (X_{ijt}):	
School holiday	Dummy = 1 if school holiday in German state (Bundesland). Either origin or destination must be in state.
Public holiday	Dummy = 1 if national or state specific holiday.
Bundesliga (Div. 1)	Dummy = 1 if division 1 football game at origin or destination.
Bundesliga (Div. 2)	Dummy = 1 if division 2 football game at origin or destination.
Munich Oktoberfest	Dummy = 1 if route to or from Munich during Oktoberfest (20/09/2014–03/10/2014).
Stuttgart Wasen	Dummy = 1 if route to or from Stuttgart during Wasen (26/09/2014–12/10/2014).

C.2 Potential transmission channels: additional regression tables

Table C.2
Transmission channel: relative travel time difference

	Dep. variable: ln ticket sales$^{\text{new}}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Channel × Strike wave 1	-0.0664	-0.0616	-0.0683	-0.176**
	(0.0654)	(0.0635)	(0.0628)	(0.0755)
Channel × Strike wave 2	0.0124	0.00719	-0.000987	-0.0637
	(0.0627)	(0.0608)	(0.0588)	(0.0694)
Channel × Strike wave 3	0.0335	-0.000881	0.00898	0.0752
	(0.0538)	(0.0535)	(0.0542)	(0.0718)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.743	0.750	0.754	0.814
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.3
Transmission channel: time delay

	Dep. variable: ln ticket sales$^{new}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Channel × Strike wave 1	-0.101 (0.0672)	-0.0650 (0.0663)	-0.0605 (0.0647)	-0.136 (0.0834)
Channel × Strike wave 2	-0.0583 (0.0612)	-0.0297 (0.0617)	-0.0258 (0.0596)	-0.153** (0.0730)
Channel × Strike wave 3	-0.103* (0.0600)	-0.0807 (0.0574)	-0.0744 (0.0534)	-0.0348 (0.0685)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.743	0.751	0.754	0.814
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.4
Transmission channel: fraction cancelled

	Dep. variable: ln ticket sales$^{\text{new}}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Channel × Strike wave 1	-0.0693 (0.0666)	-0.0199 (0.0655)	-0.0273 (0.0650)	-0.0552 (0.0820)
Channel × Strike wave 2	-0.0896 (0.0652)	-0.0464 (0.0650)	-0.0555 (0.0637)	-0.206*** (0.0712)
Channel × Strike wave 3	-0.0536 (0.0560)	-0.0173 (0.0531)	-0.0346 (0.0568)	0.0189 (0.0729)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.743	0.750	0.754	0.814
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.5
Transmission channel: triple interaction – time delay, bus travel time

	Dep. variable: ln ticket sales$^{new}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Time delay × Strike wave 1	0.172 (0.216)	0.151 (0.207)	0.194 (0.206)	0.159 (0.212)
Time delay × Strike wave 2	-0.364** (0.181)	-0.331* (0.177)	-0.282 (0.179)	-0.315 (0.204)
Time delay × Strike wave 3	-0.246 (0.169)	-0.204 (0.164)	-0.136 (0.158)	-0.215 (0.171)
Duration bus × Strike wave 1	0.274*** (0.0728)	0.283*** (0.0732)	0.274*** (0.0715)	0.292*** (0.0774)
Duration bus × Strike wave 2	0.344*** (0.0597)	0.333*** (0.0604)	0.322*** (0.0596)	0.219*** (0.0649)
Duration bus × Strike wave 3	0.442*** (0.0531)	0.419*** (0.0521)	0.403*** (0.0488)	0.343*** (0.0556)
Duration bus × Time delay × Strike wave 1	-0.0488 (0.132)	-0.0581 (0.128)	-0.0878 (0.127)	-0.112 (0.127)
Duration bus × Time delay × Strike wave 2	0.240** (0.116)	0.206* (0.113)	0.172 (0.111)	0.278** (0.128)
Duration bus × Time delay × Strike wave 3	0.175* (0.103)	0.114 (0.104)	0.0856 (0.0955)	0.186* (0.104)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.748	0.754	0.757	0.816
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.6
Transmission channel: triple interaction – fraction cancelled, bus travel time

	Dep. variable: ln ticket sales$_{ijt}^{new}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.- Day FE
Fraction cancelled × Strike wave 1	0.0323 (0.218)	-0.00689 (0.211)	0.0547 (0.206)	-0.0307 (0.221)
Fraction cancelled × Strike wave 2	-0.00471 (0.174)	-0.0168 (0.172)	0.0534 (0.170)	-0.00633 (0.190)
Fraction cancelled × Strike wave 3	-0.329* (0.183)	-0.323* (0.180)	-0.224 (0.171)	-0.298 (0.186)
Duration bus × Strike wave 1	0.254*** (0.0726)	0.261*** (0.0731)	0.254*** (0.0725)	0.271*** (0.0782)
Duration bus × Strike wave 2	0.404*** (0.0609)	0.386*** (0.0611)	0.378*** (0.0608)	0.266*** (0.0678)
Duration bus × Strike wave 3	0.432*** (0.0502)	0.403*** (0.0489)	0.391*** (0.0498)	0.337*** (0.0572)
Duration bus × Fraction cancelled × Strike wave 1	0.0280 (0.131)	0.0212 (0.128)	-0.0135 (0.125)	-0.0316 (0.133)
Duration bus × Fraction cancelled × Strike wave 2	0.0218 (0.116)	0.00724 (0.114)	-0.0317 (0.110)	0.102 (0.122)
Duration bus × Fraction cancelled × Strike wave 3	0.204* (0.110)	0.157 (0.113)	0.123 (0.104)	0.202* (0.112)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.748	0.754	0.757	0.816
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.7
Transmission channel: triple interaction – time delay, absolute travel time difference

	Dep. variable: ln ticket sales$^{new}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Time delay × Strike wave 1	0.324 (0.270)	0.392 (0.260)	0.425* (0.256)	0.363 (0.272)
Time delay × Strike wave 2	-0.336 (0.233)	-0.265 (0.228)	-0.231 (0.222)	-0.325 (0.254)
Time delay × Strike wave 3	-0.249 (0.242)	-0.167 (0.231)	-0.118 (0.202)	-0.139 (0.210)
Absolute duration × Strike wave 1	0.312*** (0.112)	0.377*** (0.111)	0.391*** (0.107)	0.357*** (0.120)
Absolute duration × Strike wave 2	0.262*** (0.100)	0.285*** (0.100)	0.302*** (0.0978)	0.215** (0.0967)
Absolute duration × Strike wave 3	0.312*** (0.104)	0.333*** (0.0997)	0.373*** (0.0865)	0.366*** (0.0899)
Absolute duration × Time delay × Strike wave 1	-0.102 (0.157)	-0.190 (0.153)	-0.227 (0.152)	-0.140 (0.171)
Absolute duration × Time delay × Strike wave 2	0.271* (0.144)	0.197 (0.140)	0.156 (0.136)	0.358** (0.165)
Absolute duration × Time delay × Strike wave 3	0.247* (0.139)	0.168 (0.135)	0.106 (0.120)	0.120 (0.135)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	8300	8300	8300	8000
R^2	0.773	0.783	0.787	0.844
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.8
Transmission channel: triple interaction – fraction cancelled, absolute travel time difference

	Dep. variable: ln ticket sales$^{new}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Fraction cancelled × Strike wave 1	0.0710 (0.269)	0.104 (0.262)	0.196 (0.254)	0.0554 (0.280)
Fraction cancelled × Strike wave 2	0.141 (0.237)	0.128 (0.231)	0.239 (0.225)	0.139 (0.245)
Fraction cancelled × Strike wave 3	-0.557** (0.231)	-0.555** (0.222)	-0.391** (0.193)	-0.444** (0.199)
Absolute duration × Strike wave 1	0.235** (0.110)	0.297*** (0.110)	0.328*** (0.109)	0.281** (0.116)
Absolute duration × Strike wave 2	0.412*** (0.107)	0.413*** (0.104)	0.452*** (0.101)	0.359*** (0.107)
Absolute duration × Strike wave 3	0.225** (0.0910)	0.226*** (0.0833)	0.302*** (0.0761)	0.297*** (0.0788)
Absolute duration × Fraction cancelled × Strike wave 1	0.0311 (0.156)	-0.0495 (0.153)	-0.116 (0.148)	0.0170 (0.162)
Absolute duration × Fraction cancelled × Strike wave 2	-0.0227 (0.147)	-0.0580 (0.143)	-0.136 (0.136)	0.0724 (0.153)
Absolute duration × Fraction cancelled × Strike wave 3	0.396*** (0.134)	0.355*** (0.133)	0.238** (0.111)	0.261** (0.119)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	8300	8300	8300	8000
R^2	0.773	0.783	0.787	0.845
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.9
Transmission channel: triple interaction – time delay, relative travel time difference

	Dep. variable: ln ticket sales$_{ijt}^{\text{new}}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Time delay × Strike wave 1	-0.159 (0.298)	-0.0474 (0.295)	-0.0459 (0.284)	0.174 (0.339)
Time delay × Strike wave 2	-0.227 (0.249)	-0.177 (0.251)	-0.171 (0.250)	-0.108 (0.257)
Time delay × Strike wave 3	0.0507 (0.293)	0.0882 (0.284)	0.112 (0.253)	0.0328 (0.299)
Relative duration × Strike wave 1	-0.137 (0.132)	-0.0949 (0.132)	-0.106 (0.127)	-0.0715 (0.145)
Relative duration × Strike wave 2	0.0113 (0.105)	0.00686 (0.108)	-0.00242 (0.109)	-0.0768 (0.105)
Relative duration × Strike wave 3	0.156 (0.128)	0.150 (0.126)	0.143 (0.110)	0.114 (0.125)
Relative duration × Time delay × Strike wave 1	0.187 (0.181)	0.0664 (0.177)	0.0465 (0.171)	-0.0311 (0.191)
Relative duration × Time delay × Strike wave 2	0.188 (0.158)	0.121 (0.157)	0.0957 (0.158)	0.161 (0.178)
Relative duration × Time delay × Strike wave 3	0.0506 (0.167)	-0.0158 (0.162)	-0.0608 (0.147)	0.00329 (0.178)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	8300	8300	8300	8000
R^2	0.770	0.780	0.784	0.842
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.10
Transmission channel: triple interaction – fraction cancelled, relative travel time difference

	Dep. variable: ln ticket sales$_{ijt}^{new}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Fraction cancelled × Strike wave 1	-0.0610 (0.275)	-0.0537 (0.267)	-0.0368 (0.262)	0.0537 (0.305)
Fraction cancelled × Strike wave 2	0.232 (0.259)	0.207 (0.252)	0.234 (0.249)	0.364 (0.270)
Fraction cancelled × Strike wave 3	-0.119 (0.263)	-0.178 (0.248)	-0.114 (0.241)	-0.226 (0.274)
Relative duration × Strike wave 1	-0.112 (0.118)	-0.0962 (0.116)	-0.100 (0.118)	-0.101 (0.124)
Relative duration × Strike wave 2	0.156 (0.117)	0.133 (0.114)	0.133 (0.114)	0.0637 (0.116)
Relative duration × Strike wave 3	0.0952 (0.102)	0.0628 (0.0967)	0.0739 (0.0946)	0.0409 (0.107)
Relative duration × Fraction cancelled × Strike wave 1	0.0896 (0.169)	0.0235 (0.162)	0.00540 (0.159)	-0.0230 (0.173)
Relative duration × Fraction cancelled × Strike wave 2	-0.0938 (0.159)	-0.126 (0.154)	-0.150 (0.153)	-0.0977 (0.172)
Relative duration × Fraction cancelled × Strike wave 3	0.0974 (0.154)	0.0781 (0.144)	0.0296 (0.139)	0.0963 (0.164)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	8300	8300	8300	8000
R^2	0.769	0.780	0.784	0.842
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.1. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

C.3 Robustness and additional results

Table C.11
Robustness: continuous dependent variable: ln(*bus travel time*)

	Dep. variable: ln ticket sales			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Treated × Strike wave 1	0.906***	0.975***	0.994***	0.991***
	(0.204)	(0.206)	(0.202)	(0.222)
Treated × Strike wave 2	1.821***	1.735***	1.752***	1.576***
	(0.160)	(0.161)	(0.156)	(0.191)
Treated × Strike wave 3	2.228***	2.095***	2.092***	1.992***
	(0.168)	(0.169)	(0.156)	(0.183)
Treated × Post	1.513***	1.427***	1.416***	1.454***
	(0.108)	(0.103)	(0.101)	(0.118)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	26832	26832	26832	26488
R^2	0.876	0.879	0.881	0.913
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.12
Robustness: treatment absolute travel time difference

	Dep. variable: ln ticket sales			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Treated × Strike wave 1	0.0293	0.0445	0.0542	0.0374
	(0.0488)	(0.0471)	(0.0451)	(0.0530)
Treated × Strike wave 2	0.130***	0.134***	0.145***	0.154***
	(0.0421)	(0.0405)	(0.0386)	(0.0416)
Treated × Strike wave 3	0.219***	0.202***	0.225***	0.238***
	(0.0425)	(0.0426)	(0.0388)	(0.0455)
Treated × Post	0.140***	0.147***	0.177***	0.179***
	(0.0276)	(0.0256)	(0.0247)	(0.0274)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	26832	26832	26832	26488
R^2	0.871	0.874	0.878	0.910
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.13
Robustness: treatment routes from or to East German cities

	Dep. var.: ln ticket sales$^{new}_{ijt}$		
	(1) Basic DD	(2) DD +trend	(3) DD + controls
Channel × Strike wave 1	0.113	0.127	0.135
	(0.128)	(0.125)	(0.130)
Channel × Strike wave 2	0.0655	0.0515	0.0771
	(0.117)	(0.0941)	(0.125)
Channel × Strike wave 3	0.163**	0.183*	0.215**
	(0.0794)	(0.0968)	(0.106)
Add. Controls			✓
Route - trend		✓	
Origin - trend			✓
Destination - trend			✓
Day FEs	✓	✓	✓
Route FEs	✓	✓	✓
Observations	15600	15600	15600
R^2	0.721	0.736	0.728
Clustered SEs	✓		✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level.)

Table C.14
Robustness: excluding Berlin

	Dep. variable: ln ticket sales			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Treated × Strike wave 1	0.172***	0.187***	0.195***	0.235***
	(0.0446)	(0.0457)	(0.0445)	(0.0487)
Treated × Strike wave 2	0.351***	0.325***	0.334***	0.355***
	(0.0345)	(0.0342)	(0.0328)	(0.0427)
Treated × Strike wave 3	0.413***	0.387***	0.396***	0.418***
	(0.0373)	(0.0370)	(0.0342)	(0.0396)
Treated × Post	0.283***	0.264***	0.277***	0.286***
	(0.0217)	(0.0211)	(0.0205)	(0.0248)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	21672	21672	21672	20984
R^2	0.824	0.828	0.832	0.879
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.15

Robustness: dependent variable ln(*total ticket sales*); excluding post-strike period

	Dep. variable: ln ticket sales$^{new}_{ijt}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.- Day FE
Channel × Strike wave 1	0.131***	0.148***	0.131***	0.128***
	(0.0462)	(0.0453)	(0.0436)	(0.0464)
Channel × Strike wave 2	0.291***	0.277***	0.258***	0.225***
	(0.0356)	(0.0349)	(0.0344)	(0.0394)
Channel × Strike wave 3	0.387***	0.355***	0.330***	0.326***
	(0.0377)	(0.0376)	(0.0328)	(0.0396)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	15600	15600	15600	15400
R^2	0.881	0.885	0.888	0.917
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.16
Robustness: including two days before and after each strike, and intermediate period

	Dep. variable: ln ticket sales			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Treated × Strike wave 1	0.134***	0.158***	0.159***	0.164***
	(0.0378)	(0.0377)	(0.0373)	(0.0410)
Treated × Strike wave 2	0.306***	0.301***	0.300***	0.276***
	(0.0287)	(0.0294)	(0.0288)	(0.0355)
Treated × Strike wave 3	0.386***	0.361***	0.355***	0.354***
	(0.0294)	(0.0297)	(0.0273)	(0.0330)
Treated × Post	0.255***	0.242***	0.229***	0.260***
	(0.0169)	(0.0169)	(0.0158)	(0.0185)
Add. Controls		✓	✓	✓
Origin - trend			✓	✓
Destination - trend			✓	✓
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	33384	33384	33384	32956
R^2	0.877	0.880	0.882	0.914
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.17
Robustness: excluding within-German flights

	Dep. variable: ln ticket sales$_{ijt}^{new}$			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
Channel × Strike wave 1	0.235***	0.252***	0.234***	0.273***
	(0.0632)	(0.0635)	(0.0625)	(0.0697)
Channel × Strike wave 2	0.425***	0.406***	0.385***	0.333***
	(0.0526)	(0.0526)	(0.0526)	(0.0664)
Channel × Strike wave 3	0.458***	0.427***	0.396***	0.371***
	(0.0476)	(0.0471)	(0.0458)	(0.0544)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	14800	14800	14800	14600
R^2	0.740	0.746	0.749	0.811
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

Table C.18
Robustness: excluding return ticket bookings

	Dep. variable: ln ticket sales			
	(1) Basic DD	(2) DD + controls	(3) DD + trend	(4) Orig.-, Dest.-Day FE
treated × Strike wave 1	0.153***	0.159***	0.162***	0.152***
	(0.0375)	(0.0384)	(0.0375)	(0.0410)
treated × Strike wave 2	0.355***	0.341***	0.345***	0.314***
	(0.0296)	(0.0299)	(0.0296)	(0.0392)
treated × Strike wave 3	0.367***	0.350***	0.351***	0.337***
	(0.0313)	(0.0315)	(0.0295)	(0.0340)
treated × Post	0.255***	0.243***	0.244***	0.255***
	(0.0170)	(0.0169)	(0.0166)	(0.0193)
Add. Controls		✓	✓	
Origin - trend			✓	
Destination - trend			✓	
Day FEs	✓	✓	✓	
Route FEs	✓	✓	✓	✓
Origin-Day FEs				✓
Destination-Day FEs				✓
Observations	26832	26832	26832	26488
R^2	0.873	0.875	0.877	0.909
Clustered SEs	✓	✓	✓	✓

Notes: Estimated coefficients from Equation 3.2. Standard errors in parentheses, clustered at the route level (166 clusters). ***/**/* indicate significance at the 1%/5%/10% level. Column 4 refers to the inclusion of γ_{it} and γ_{jt} – origin- and destination-day specific fixed effects.)

List of Figures

List of Tables

Bibliography

References to Preface

[1] Anderson, J. E. (2000). Why do nations trade (so little)?. *Pacific Economic Review 5*(2), 115-134.

[2] Anderson, J. E., and E. Van Wincoop (2004). Trade costs. *Journal of Economic Literature 42*(3), 691-751.

[3] Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains?. *American Economic Review 102*(1), 94-130.

[4] Armenter, R., and M. Koren (2014). A balls-and-bins model of trade. *American Economic Review 104*(7), 2127-2151.

[5] Autor D., D. Dorn, and G. H. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6), 2121-2168.

[6] Baldwin, R. E. (1970). Nontariff distortions of international trade. *Brookings Institution, Washington, D.C.*, Ch. 3.

[7] Becker, S. O., K. Boeckh, C. Hainz, and L. Woessmann (2016). The empire is dead, long live the empire! Long-run persistence of trust and corruption in the bureaucracy. *Economic Journal 126*(590), 40-74.

[8] Bernard, A. B., S. J. Redding, and P. K. Schott (2011). Multiproduct firms

and trade liberalization. *Quarterly Journal of Economics 126*(3), 1271-1318.

[9] Das, S., M. J. Roberts, and J. R. Tybout (2007). Market entry costs, producer heterogeneity, and export dynamics. *Econometrica 75*(3), 837-873.

[10] Dixit, A., and Norman, V. (1980). Theory of international trade: A dual, general equilibrium approach. *Cambridge University Press.*

[11] Eckel, C., and Neary, J. P. (2010). Multi-product firms and flexible manufacturing in the global economy. *Review of Economic Studies 77*(1), 188-217.

[12] Felbermayr, G., and J. Gröschl (2014). Within US trade and the long shadow of the American secession. *Economic Inquiry 52*(1), 382-404.

[13] Felbermayr, G., J. Prat, and H. J. Schmerer (2011). Globalization and labor market outcomes: Wage bargaining, search frictions, and firm heterogeneity. *Journal of Economic Theory 146*(1), 39-73.

[14] Feyrer, J. (2009a). Distance, trade, and income – the 1967 to 1975 closing of the Suez canal as a natural experiment. NBER Working Papers 15557, National Bureau of Economic Research.

[15] Feyrer, J. (2009b). Trade and income – exploiting time series in geography. NBER Working Papers 14910, National Bureau of Economic Research.

[16] Grossman, G. M. (1998) Comment on Determinants of Bilateral Trade: Does Gravity Work in a Neoclassical World?. University of Chicago Press, Chicago.

[17] Grossman, G. M., and E. Rossi-Hansberg (2008). Trading tasks: A simple theory of offshoring. *American Economic Review 98*(5), 1978-1997.

[18] Guiso, L., P. Sapienza, and L. Zingales (2009). Cultural biases in economic exchange?. *Quarterly Journal of Economics 124*(3), 1095-1131.

[19] Head, K., and T. Mayer (2013). What separates us? Sources of resistance to globalization. *Canadian Journal of Economics 46*(4), 1196-1231.

[20] Head, K. and T. Mayer (2013). Gravity equations: Workhorse, toolkit, and cookbook, Ch. 4 in *Handbook of International Economics*, edited by Gopinath et al.

[21] Head, K., T. Mayer, and J. Ries (2010). The erosion of colonial trade linkages after independence. *Journal of International Economics 81*(1), 1-14.

[22] Helpman, E., M. Melitz, and Y. Rubinstein (2008). Estimating trade flows: Trading partners and trading volumes. *Quarterly Journal of Economics 123*(2), 441-487.

[23] Hornok, C., and M. Koren (2015). Per-shipment costs and the lumpiness of international trade. *Review of Economics and Statistics 97*(2), 525-530.

[24] Hornok, C., and Koren, M. (2015). Trade and development in a globalized world: The roadmap for a research agenda, Ch. 6 in *Economics Without Borders. Economic Research for European Policy Challenges*, edited by Blundell et al.

[25] Hummels, D.,J. Ishii , and K. M. Yi (2001). The nature and growth of vertical specialization in world trade. *Journal of International Economics 54*(1), 75-96.

[26] Hummels, D. L., and G. Schaur (2013). Time as a trade barrier. *American Economic Review 103*(7), 2935-2959.

[27] Irarrazabal, A., A. Moxnes, and L. D. Opromolla (2015). The tip of the iceberg: a quantitative framework for estimating trade costs. *Review of Economics and Statistics 97*(4), 777-792.

[28] Krugman, P. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of International Economics 9*(4), 469-479.

[29] Krugman, P. (1980). Scale economies, product differentiation, and the pattern of trade. *American Economic Review 70*(5), 950-959.

[30] Larcom, S., F. Rauch, and T. Willems (2015). The benefits of forced experimentation: striking evidence from the London Underground network. CEPR Discussion Papers 10854.

[31] Lileeva, A., and D. Trefler (2010). Improved Access to Foreign Markets Raises Plant-level Productivity... For Some Plants. *Quarterly Journal of Economics, 125*(3), 1051-1099.

[32] Melitz, M. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica 71*(6), 1695-1725.

[33] Morgan, P., and R. Manning (1985). Optimal search. *Econometrica 53*(4), 923-44.

[34] Nitsch, V., and N. Wolf (2013). Tear down this wall: On the persistence of borders in trade. *Canadian Journal of Economics 46*(1), 154-179.

[35] Porter, M. (1991). America's green strategy, *Scientific American 264*(4), 33-35.

[36] Trefler, D. (1995). The Case of the Missing Trade and Other Mysteries. *American Economic Review 85*(0), 1029-1046.

[37] Verhoogen, E. A. (2008). Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector. *Quarterly Journal of Economics, 123*(2), 489-530.

[38] Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641-654.

[39] Yeaple, S. R. (2013). The multinational firm. *Annual Review of Economics 5*(1), 193-217.

References to Chapter 1

[40] Algan, Y., and P. Cahuc (2010). Inherited trust and growth, *American Economic Review 100*(5), 2060-2092.

[41] Anderson, J. E., and E. Van Wincoop (2003). Gravity with gravitas: A solution to the border puzzle, *American Economic Review, 93*(1), 170-192.

[42] Anderson, J. E., and Y. V. Yotov. (2012). Gold standard gravity. NBER Working Papers 17835, National Bureau of Economic Research.

[43] Baltagi, B. H. (1995). Econometric analysis of panel data. Wiley.

[44] Barbieri, K. (2002). The liberal illusion: does trade promote peace?. University of Michigan Press.

[45] Becker, S. O., K. Boeckh, C. Hainz, and L. Woessmann (2016). The empire is dead, long live the empire! Long-run persistence of trust and corruption in the bureaucracy. *Economic Journal 126*(590), 40-74.

[46] Boyer, J. W. (1989). Some reflections on the problem of Austria, Germany, and Mitteleuropa. *Central European History 22*(3-4), 301-315.

[47] Butschek, F. (1996). Statistische Reihen zur Österreichischen Wirtschaftsgeschichte, WIFO Wien.

[48] Clark, C. (2013). The Sleepwalkers, How Europe went to war in 1914. Penguin Books.

[49] De Menil, G. and M. Maurel (1994). Breaking up a customs union: the case of the Austro-Hungarian empire. *Weltwirtschaftliches Archiv 130*(3), 553-575.

[50] Dell, M., B. F. Jones, and B. A. Olken (2014). What do we learn from the weather? The new climate–economy literature. *Journal of Economic Literature 52*(3), 740-798.

[51] Disdier, A.-C., and T. Mayer (2007) Je t'aime, moi non plus: Bilateral Opinions and International Trade. *European Journal of Political Economy 23*(4), 1140-1159.

[52] Dixon, W. J., and B. E. Moon (1993), Political similarity and American foreign trade patterns, *Political Research Quarterly 46*(1), 5-25.

[53] Djankov, S. and C. Freund (2002). Trade Flows in the Former Soviet Union, 1987 to 1996, *Journal of Comparative Economics 30*(1), 76-90.

[54] Dornbusch, R. (1992). Monetary problems of post-communism: Lessons from the end of the Austro-Hungarian empire. *Review of World Economics 128*(3), 391-424.

[55] Eddie, S. (1989). Economic policy and economic development in Austria-Hungary, 1867-1913. *The Cambridge Economic History of Europe 8*, 814-886.

[56] Egger, P. (2000). A note on the proper econometric specification of the gravity equation, *Economics Letters 66*(1), 25-31.

[57] Egger, P. H., and A. Lassmann (2015). The causal impact of common native language on international trade: Evidence from a spatial regression discontinuity design. *Economic Journal 125*(584), 699-745.

[58] Erdösi, F. (1999). Regional characteristics of the development of the transportation and telecommunication during the systematic change, in *Regional processes and spatial structures in Hungary in the 1990s*, edited by Hajdu, Center for regional studies.

[59] Feenstra, R. C. (2004). Advanced international trade. Princeton University Press, Princeton, NJ.

[60] Feenstra, R. C., R. E. Lipsey, H. Deng, A. C. Ma, and H. Mo (2005) World trade flows: 1962-2000. NBER Working Papers 11040, National Bureau of

Economic Research.

[61] Felbermayr, G., and J. Gröschl (2014). Within US trade and the long shadow of the American secession. *Economic Inquiry 52*(1), 382-404.

[62] Felbermayr, G. J., and W. Kohler (2006). Exploring the intensive and extensive margins of world trade, *Review of World Economics 142*(4), 642-674.

[63] Felbermayr, G. J., and F. Toubal (2010). Cultural proximity and trade, *European Economic Review 54*(2) 279-293.

[64] Guillaume, G. and S. Zignago (2010). BACI: international trade database at the product-level (the 1994-2007 version).

[65] Good, D. F. (1984). The economic rise of the Habsburg empire 1750-1918. University of California Press. Berkeley, CA.

[66] Guiso, L., P. Sapienza, and L. Zingales (2009). Cultural biases in economic exchange?. *Quarterly Journal of Economics 124*(3), 1095-1131.

[67] Hartmuth, M. (2011). Images of imperial legacy: Modern discourses on the social and cultural impact of Ottoman and Habsburg rule in Southeast Europe. LIT Verlag Münster. Vol. 10.

[68] Head, K. and T. Mayer (2013). Gravity equations: Workhorse, toolkit, and cookbook, Ch. 4 in *Handbook of International Economics*, edited by Gopinath et al.

[69] Head, K., T. Mayer, and J. Ries (2010). The erosion of colonial trade linkages after independence. *Journal of International Economics 81*(1), 1-14.

[70] Heid, B., M. Larch, and Y. V. Yotov (2015). A simple method to estimate the effects of non-discriminatory trade policy within structural gravity models. Working Paper

[71] Jeffreys, D. (2007). The Austro-Hungarian Empire – New Holland History

Maps, New Holland Publishers Ltd.

[72] Lameli, A., V. Nitsch, J. Südekum, and N. Wolf (2015) Same same but different: Dialects and trade. *German Economic Review 16*(3), 290-306

[73] Lazarevic, Z. (2010). Yugoslavia: Economic aspects of the position between East and West, in *Gaps in the Iron Curtain: Economic Relations between Neutral and Socialist Countries in Cold War Europe*. edited by Enderle-Burcel et al. Jagiellonian University Press.

[74] Martin, W., and C. S. Pham (2015). Estimating the gravity model when zero trade flows are frequent and economically determined. World Bank Policy Research Working Paper 7308.

[75] Mayer, T., and S. Zignago (2011). Notes on CEPII's distances measures: The GeoDist database.

[76] Melitz, J., and F. Toubal (2014). Native language, spoken language, translation and trade. *Journal of International Economics 93*(2), 351-363.

[77] Michaels, G., and X. Zhi (2010). Freedom fries, *American Economic Journal: Applied Economics 2*(3), 256-281.

[78] Nickell, S. (1981). Biases in dynamic models with fixed effects, *Econometrica 49*(6), 1417-1426.

[79] Nitsch, V., and N. Wolf (2013). Tear down this wall: on the persistence of borders in trade. *Canadian Journal of Economics 46*(1), 154-179.

[80] Pogany, A. (2010). Co-operation through the Iron Curtain: economic relations between Austria and Hungary after the Second World War, in *Gaps in the Iron Curtain: Economic Relations between Neutral and Socialist Countries in Cold War Europe*. edited by Enderle-Burcel et al. Jagiellonian University Press.

[81] Rauch, F. (2016), The geometry of the distance coefficient in gravity equations

in international trade. *Review of International Economics 24*(5), 1167–1177

[82] Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics 48*(1), 7-35.

[83] Redding, S. J. and D. M. Sturm (2008). The costs of remoteness: Evidence from German division and reunification, *American Economic Review 98*(5), 1766-1797.

[84] Resch, A. (2010). Forty years of foreign trade between Austria and the CMEA area, in *Gaps in the Iron Curtain: Economic Relations between Neutral and Socialist Countries in Cold War Europe.* edited by Enderle-Burcel et al. Jagiellonian University Press.

[85] Schlipphacke, H. (2014). The temporalities of Habsburg nostalgia. *Journal of Austrian Studies 47*(2), 1-16.

[86] Schulze, M.-S. (2010). Patterns of growth and stagnation in the late nineteenth century Habsburg economy, *European Review of Economic History 4*(3), 311-340.

[87] Schulze, M.-S. and N. Wolf (2009). On the origins of border effects: insights from the Habsburg Empire. *Journal of Economic Geography 9*(1), 117-136.

[88] Schulze, M.-S. and N. Wolf (2012). Economic nationalism and economic integration: the Austro-Hungarian Empire in the late nineteenth century. *Economic History Review 65*(2), 652-673.

[89] Silva, J. S., and S. Tenreyro (2006). The log of gravity, *Review of Economics and Statistics 88*(4), 641-658.

[90] Spolaore, E. and R. Wacziarg (2009). The diffusion of development. *Quarterly Journal of Economics 124*(2), 469-529.

[91] Spolaore, E. and R. Wacziarg (2015). Ancestry, language and culture. NBER

Working Papers 21242, National Bureau of Economic Research.

[92] Thom, R., and B. Walsh (2002). The effect of a currency union on trade: Lessons from the Irish experience, *European Economic Review 46*(6), 1111-1123.

[93] Umana Dajud, C. (2013). Political proximity and international trade, *Economics and Politics 25*(3), 283-312.

[94] Wank, S. (1997). Some reflections on the Habsburg empire and its legacy in the nationalities question. *Austrian History Yearbook 28*, 131-146.

[95] Yotov, Y. V. (2912). A simple solution to the distance puzzle in international trade. *Economics Letters 117*(3), 794-798.

References to Chapter 2

[96] Acemoglu, D. and J. A. Robinson (2006). De facto political power and institutional persistence. *American Economic Association Papers and Proceedings 96*(2), 325–330.

[97] Baylis, K., A. Martens, and L. Nogueira (2009). What drives import refusals. *American Journal of Agricultural Economics 91*(5), 1477–1483.

[98] Baylis, K., L. Nogueira, and K. Pace (2011). Food import refusals: Evidence from the European Union. *American Journal of Agricultural Economics 93*(2), 566–572.

[99] Becker, G. S. (1957). The Economics of discrimination. University Chicago Press.

[100] Caballero, R. J. and V. Corbo (1989). The effect of real exchange rate uncertainty on exports: Empirical evidence. *World Bank Economic Review 3*(2), 263–278.

[101] Chen, M. and A. Mattoo (2008). Regionalism in standards: good or bad for trade? *Canadian Journal of Economics 41*(3), 838–863.

[102] Chen, N. and D. Novy (2012). On the measurement of trade costs: direct vs. indirect approaches to quantifying standards and technical regulations. *World Trade Review 11*(3), 401–414.

[103] Disdier, A.-C., L. Fontagné, and M. Mimouni (2008). The impact of regulations on agricultural trade: Evidence from the SPS and TBT agreements. *American Journal of Agricultural Economics 90*(2), 336–350.

[104] Essaji, A. (2008). Technical regulations and specialization in international trade. *Journal of International Economics 76*(2), 166–176.

[105] Falvey, R. E. (1989). Trade, quality reputations and commercial policy. *International Economic Review*(3), 607–622.

[106] Feng, L., Z. Li, and D. L. Swenson (2014). Trade policy uncertainty and exports: Evidence from China's WTO accession. CESifo Working Paper 4780.

[107] Fontagné, L., G. Orefice, R. Piermartini, and N. Rocha (2015). Product standards and margins of trade: Firm level evidence. *Journal of International Economics 97*(1), 29–44.

[108] Freund, C. L., and N. Rocha (2011). What constrains Africa's exports? *World Bank Economic Review 25*(3), 361–386.

[109] Grobar, L. M. (1993). The effect of real exchange rate uncertainty on LDC manufactured exports. *Journal of Development Economics 41*(2), 367–376.

[110] Grundke, R., and C. Moser (2014). Hidden protectionism? Evidence from non-tariff barriers to trade in the United States. ETH Zurich Working Papers 369.

[111] Guimbard, H., S. J. M. Mimouni, and X. Pichot (2012). MAcMap-HS6 2007,

an exhaustive and consistent measure of applied protection in 2007. *International Economics 130*(2), 99–121.

[112] Handley, K. (2014). Exporting under trade policy uncertainty: Theory and evidence. *Journal of International Economics 94*(1), 50–66.

[113] Héricourt, J., and S. Poncet (2015). Exchange rate volatility, financial constraints, and trade: Empirical evidence from Chinese firms. *World Bank Economic Review 29*(3), 550–578.

[114] Hooper, P. and S. W. Kohlhagen (1978). The effect of exchange rate uncertainty on the prices and volume of international trade. *Journal of International Economics 8*(4), 483–511.

[115] Jaud, M., O. Cadot, and A. Suwa-Eisenmann (2013). Do food scares explain supplier concentration? an analysis of EU agri-food imports. *European Review of Agricultural Economics 40*(5), 873–890.

[116] Javorcik, B. S. (2004). The composition of foreign direct investment and protection of intellectual property rights: Evidence from transition economies. *European Economic Review 48*(1), 39-62.

[117] Jouanjean, M.-A., J.-C. Maur, and B. Shepherd (2015). Reputation matters: Spillover effects for developing countries in the enforcement of us food safety measures. *Food Policy 55*, 81–91.

[118] Kee, H. L., A. Nicita, and M. Olarreaga (2009). Estimating trade restrictiveness indices. *Economic Journal 119*(534), 172–199.

[119] Knowles, J., N. Persico, and P. Todd (2001). Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy 109*(1), 203–229.

[120] Mayer, T. and G. Ottaviano (2008). The happy few: The internationalisation of European firms. *Intereconomics 43*(3), 135-148.

[121] McQuade, T. J., S. W. Salant, and J. A. Winfree (2012). Markets with untraceable goods of unknown quality: Beyond the small-country case. Resources for the Future Discussion Paper 09-31-rev.

[122] Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica 71*(6), 1695–1725.

[123] Nabokin, T. (2014). Global Investment Decisions and Patent Protection: Evidence from German Multinationals. Munich Discussion Paper 2014-40

[124] Nordas, H. K., and R. Piermartini (2004). Infrastructure and trade. WTO Staff Working Papers 2004-04.

[125] Osnago, A., R. Piermartini, and N. Rocha (2015). Trade policy uncertainty as barrier to trade. *WTO Staff Working Papers 2015-05.*

[126] Shapiro, C. (1983). Premiums for high quality products as returns to reputations. *Quarterly Journal of Economics 98*(4), 659–680.

[127] Wooldridge, J. (2010). Econometric Analysis of Cross Section and Panel Data. MIT Press, Cambridge MA.

References to Chapter 3

[128] Anderson, J. E., and E. Van Wincoop (2003). Gravity with gravitas: A solution to the border puzzle, *American Economic Review, 93*(1), 170-192.

[129] Angrist, J. D., and J. S. Pischke (2014). Mastering'metrics: the path from cause to effect. Princeton University Press, Princeton NJ.

[130] Ashenfelter, O., and D. Card (1985). Using the longitudinal structure of earnings to estimate the effect of training programs. *Review of Economics and Statistics 67*(4), 648-660.

[131] Autor, D. (2003). Outsourcing at will: The contribution of unjust dismissal

doctrine to the growth of employment outsourcing. *Journal of Labor Economics 21*(1), 1-42.

[132] Bataille, M., and A. Steinmetz (2013). Intermodal competition on some routes in transportation networks: The case of inter urban buses and railways. Düsseldorf Institute for Competition Economics (DICE) Working Paper 84.

[133] Bauernschuster, S., T. Hener, and H. Rainer (2015). When labor disputes bring cities to a standstill: The impact of public transit strikes on traffic, accidents, air Pollution, and health. Annual Conference German Economic Association 112957.

[134] Baumol, W. J., and R. E. Quandt (1964). Rules of thumb and optimally imperfect decisions. *American Economic Review 54*(2), 23-46.

[135] Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates?. *Quarterly Journal of Economics 119*(1), 249-275.

[136] Böckers, V., J. Haucap, U. Heimeshoff, and S. Thorwarth (2015). Auswirkungen der Fernbusliberalisierung auf den Schienenpersonenfernverkehr. In *List Forum für Wirtschafts-und Finanzpolitik 41*(1), 75-90. Springer Berlin Heidelberg.

[137] Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review 99*(4), 1145-1177.

[138] Coates, D., and T. Harrison (2005). Baseball strikes and the demand for attendance, *Journal of Sports Economics 6*(3), 282-302.

[139] Dürr, N. S., S. Heim, and K. Hüschelrath (2015). Deregulation, competition, and consolidation: The case of the German interurban bus industry. ZEW Discussion Papers, 15-061.

[140] Evangelinos, C., M. Mittag, and A. Obermeyer (2015) Die ökonomischen Risiken einer zu naiven Marktliberalisierung-der Fall des deutschen Fernbusmarktes, *Zeitschrift für Verkehrswissenschaft 86*(1).

[141] Felbermayr, G. J., and W. Kohler (2006). Exploring the intensive and extensive margins of world trade, *Review of World Economics 142*(4), 642-674.

[142] Foster, L., J. C. Haltiwanger, and C. Syverson (2012). The slow growth of new plants: Learning about demand?. NBER Working Papers 17853, National Bureau of Economic Research.

[143] Fricke, H. (2015). Identification based on Difference-in-Differences Approaches with Multiple Treatments. University of St Gallen Working Paper 1510.

[144] Friebel, G., and M. Niffka (2009). The functioning of inter-modal competition in the transportation market: Evidence from the entry of low-cost airlines in Germany, *Review of Network Economics8*(2).

[145] Gagnepain, P., M. Ivaldi, and C. Vibes (2011). The industrial organization of competition in local bus services, in *A Handbook of Transport Economics*, 744-762.

[146] Guihéry, L., M. R. T. Environnement, and F. Pontoise (2016). New interurban coach services in Germany and France: which European perspective? The key role of bus station... but where to locate them?, ch. 14 in *Mobility in a Globalised World 2015.* edited by Sucky et al. University of Bamberg Press.

[147] Hirschhausen, C. von, M. Walter, F. Haunerland, and R. Moll (2008). Das Potenzial des Fernlinienbusverkehrs in Deutschland. Transport Economics Working Paper 15.

[148] Kennan, J. (1986). The economics of strikes, *Handbook of Labor Economics 2*, 1091-1137.

[149] Larcom, S., F. Rauch, and T. Willems (2015). The benefits of forced experimentation: striking evidence from the London Underground network. CEPR Discussion Papers 10854.

[150] Morgan, P., and R. Manning (1985). Optimal Search. *Econometrica 53*(4), 923-44.

[151] Nunn, N., and N. Qian (2011). The Potato's Contribution to Population and Urbanization: Evidence From A Historical Experiment. *Quarterly Journal of Economics 126*(2), 593-650.

[152] Porter, M. (1991). America's Green Strategy, *Scientific American 264*(4), 33-35.

[153] Repetto, L. (2016). Political budget cycles with informed voters: evidence from Italy. Uppsala University Working Paper 2016- 6.

[154] Siegert, C. and R. Ulbricht (2015). Dynamic oligopoly pricing: Evidence from the airline industry. Working Paper SSRN 2413180.

[155] van Exel, N. J. A., and P. Rietveld (2001). Public transport strikes and traveller behaviour. *Transport Policy 8*(4), 237-246.

[156] Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641-654.

[157] Williams, K. R. (2013). Dynamic Airline Pricing and Seat Availability. *Working Paper.*