Social Networks, Optimal Contract Design, and Present Bias: Four Essays in Applied Economic Theory

Inaugural-Dissertation

zur Erlangung des Grades

Doctor oeconomiae publicae (Dr. oec. publ.)

an der Ludwig-Maximilians-Universität München

2016

vorgelegt von

Marco Alexander Schwarz

Referent:	Prof. Dr. Klaus M. Schmidt
Korreferent:	Prof. Dr. Lutz Arnold
Promotionsabschlussberatung:	10.05.2017

Datum der mündlichen Prüfung: 10.04.2017

Namen der Berichterstatter: Klaus M. Schmidt, Lutz Arnold, Simeon Schudy

Acknowledgements

This dissertation would not have been possible without the support of numerous people.

First of all, I would like to thank my advisor Klaus Schmidt for being supportive but also challenging and critical. This combination and his detailed comments improved this dissertation a lot. I would also like to thank my second advisor Lutz Arnold for his patience and his encouraging comments and helpful suggestions. In addition, I would like to thank Simeon Schudy for being my third committee member and his insights along the way. Moreover, I am grateful to Florian Englmaier, Matthias Fahn, Fabian Herweg, and Takeshi Murooka not only for being great coauthors but also for many comments and discussions about my other research projects and keeping my morale up. Working with all of my coauthors was a pleasure and I hope that there are many more joint projects to come.

My time as a doctoral student would not have been the same without the many colleagues, friends, professors, and teaching assistants I got to know in Munich, in Berkeley, and via the Bavarian Graduate Program in Economics. Thank you for the many fruitful and also the less fruitful discussions, your support, your comments and suggestions, your inspiration, and simply the fun times I could enjoy with you.

I am also grateful to the student research assistants and the administrative staff for making my life easier and to the guests at our department and several seminar and conference participants for their helpful comments.

Financial support from the Elite Network of Bavaria through the Bavarian Graduate Program in Economics and from the Deutsche Forschungsgemeinschaft through SFB/TR 15 are gratefully acknowledged.

Finally, I would like to thank my family, my girlfriend, and my friends for their support, motivation, trust, and understanding.

Contents

List of Figures

Preface

This dissertation consists of four chapters which aim to explain different economic and societal issues and to derive policy implications with methods from microeconomic theory. All chapters are self-contained and can be read independently. The first chapter explains how cost overruns in procurement can be the result of a constrained-optimal award procedure. The second chapter shows that poorly-timed choice-enhancing policies can backfire when some consumers are naive present-biased, but well-timed policies can robustly improve consumer welfare. The third chapter analyzes how firms may exploit naive present-biased employees by providing them a menu of contracts. The fourth chapter derives how the increasing advancements in communication technology first decrease and then increase the amount of disagreement in a society.

Chapter 1 relies on standard assumptions, whereas the other three chapters contain deviations from traditional microeconomic theory. Those three employ approaches from behavioral economics and the last chapter also incorporates methods from the economics of networks.

Behavioral economics extends traditional microeconomic theory by relaxing some assumptions about human behavior. While traditional microeconomic theory assumes that decision-makers behave as homines economici, behavioral economics, often inspired by psychology, analyzes and incorporates systematic deviations from rational and selfish behavior. More concretely, the models in Chapters 2 and 3 allow decision-makers to have a so-called present bias. Instead of discounting with a factor of $\delta \in (0,1)$ between any two periods as in the standard model, decision-makers discount the utility of future periods with an additional $\beta \in (0,1)$ when comparing it to utility in the presence. This means that the decision-makers with a present bias overweight the presence in their decisions and therefore may behave time-inconsistently. For example, an individual might plan to go to the gym in seven days. For the next six days the person has no doubt that she will go to the gym on that day, but on the seventh day she cannot exert enough self-control to actually go but rather postpones going to the gym to a later time, even if she has not received any new information.[1] Strotz (1955) was the first to formalize inconsistent time preferences, Phelps and Pollak (1968) developed the $\beta-\delta$-model, which is widely used because of its portability and is also deployed in this dissertation. O'Donoghue and Rabin (1999a) introduced the concept of naivete and sophistication. While sophisticated individuals with a present bias are aware of their self-control problem,

[1] DellaVigna and Malmendier (2004, 2006) analyze individuals' behavior of (not) going to the gym.

naive ones believe that they will behave time-consistently in the future. DellaVigna (2009) surveys the empirical evidence from the field. Chapters 2 and 3 of this dissertation make use of the concept of naive present bias to explain phenomenona which would be difficult to explain without behavioral economics and derive policy implications. Chapter 4 analyzes the behavior of individuals with correlation neglect. Unlike individuals in the standard model, who update (quasi-)rationally, these individuals do not take into account a potential correlation of signals when they update their beliefs. Eyster and Weizsäcker (2011), Enke and Zimmermann (2015), and Enke (2015) provide experimental evidence. For this bias, the literature has not converged to a simple work horse model yet. There are different approaches (Levy and Razin 2015b, Ortoleva and Snowberg 2015) and Chapter 4 adds another one.

The economics of networks extends the standard framework by relaxing some common assumptions about the interaction of individuals in large groups. In traditional microeconomics, the payoffs of different parties are usually only dependent on the actions of the directly involved parties. Alternatively, all information is aggregated and transactions are realized via a central market which all individuals have access to. In the economics of networks, the connections between individuals are modeled explicitly, employing concepts from graph theory. The economics of networks studies how connections between individuals affect outcomes, for example, how information spreads and whether and under what conditions a transaction takes place. Chapter 4 analyzes how individuals form their beliefs when they do not have access to all information in a society but only to information they or their direct contacts observe directly. Only over time, they also get access to the information of their indirect contacts. Potentially – if they do not have direct or indirect connections to all members of the society – they might even never get access to some information. Whether and how fast all information is aggregated depends on the pattern of connections. Moreover, the connections in Chapter 4 are not exogenously given but are endogenously formed. The formation of connections comes with costs and benefits and thus with trade-offs for decision-makers. Depending on the costs and benefits, different network structures may arise and consequently, information may be aggregated differently. Chapter 4 is one of the few works combining endogenous network formation and belief formation – the others are Acemoglu, Bimpikis and Ozdaglar (2014), Tan (2015), and Song (2016).

Chapter 1 is joint work with Fabian Herweg. We analyze procurement auctions in which the procurer can specify an ex-ante design but cannot commit not to renegotiate the design and the price. We show that the common cost-overruns in public procurement are not necessarily due to inefficient project management or strategizing politicians but can be the result of a constrained-optimal award procedure. This means that we do not have to automatically be concerned about the taxpayers' money when we observe cost-overruns, but it is worth to take a closer look at what led to the cost-overruns. Determining empirically to what degree the different explanations contribute to the cost-overruns in public procurement is left for future research.

In our model, a procurer runs an auction off a contract to buy a good or service which can vary in its complexity. More complex projects provide a higher utility to the procurer, but they are also more expensive. Our model assumes that the cost differences between different firms are smaller for simpler projects. Under this assumption, the procurer finds it optimal to fix a simple design ex ante and to renegotiate to a more complex one ex post, which will come at an additional cost. The firms anticipate the renegotiation and compete most of the surplus they might gain from the renegotiation away. The competition is most intense when the firms are most similar, which is the case for a simple design. Moreover, we show that the results are unchanged if the procurer can use a maximum bid for the auction.

We also investigate so-called scoring auctions. Here, the bidders submit a price and a design and get a score for the combination of the two. The firm with the highest score wins the auction. Surprisingly, the procurer cannot benefit from using such a multi-dimensional auction, but the optimal scoring rule depends only on the price. The reason is that allowing the firms to differentiate their bids along different dimensions would reduce competition but not benefit the procurer, because the renegotiation always leads to an efficient outcome.

Our results change when the renegotiation is not efficient, which we model as an exogenous risk of a break down of the renegotiations. Then, a procurer will not always choose the simplest ex-ante design any more, because she is not guaranteed that the efficient design will be implemented later on. However, there remains a tendency to specify an ex-ante design that is simpler than the expected efficient design. Thus, renegotiations more often lead to more complex designs and higher prices than simpler designs and lower prices.

Chapter 2, which is joint work with Takeshi Murooka, deals with automatic enrollments and renewals that are common for many products and services such as cell-phone contracts, credit cards, and cable contracts. We analyze the optimal timing of choice-enhancing policies (e.g., policies that motivate consumers to make an active choice) when firms can strategically react to them. One example of such a policy is requiring firms to inform consumers how to cancel a contract. One might think that any policy that reduces the cancellation or switching costs for consumers increases consumer welfare. However, we show that this is not always the case and that such a policy can even decrease consumer welfare if it is poorly-timed.

In our model, a firm provides an automatic enrollment or renewal to consumers. A fraction of consumers are rational and the other consumers are naive present-biased. We show that a conventional choice-enhancing policy, which decreases consumers' switching costs when they are initially enrolled, can be detrimental to consumer and social welfare. The reason for this is that the firm would have to lower the price to keep all its customers in response to the policy. Then, however, it might be optimal to rather increase the price and only serve present-biased customers. We suggest an alternative policy that decreases consumers' switching costs whenever the firm increases its price and show that this policy increases consumer and social welfare robustly. Thus, the timing of policies matters.

We also analyze various extensions to investigate how robust our finding is in less simplified settings. First, we show that our finding does not change if consumers' sign-up decisions are endogenous and if the time horizon of the enrollment is extended and firm can charge consumers several times instead of only once. Moreover, we find that as the number of times a firm can charge consumers increases, exploiting naive consumers becomes more attractive for the firm. Furthermore, we show that our suggested policy that decreases consumers' switching cost whenever the firm increases the price increases consumer and social welfare as robustly as a policy that always reduces consumers' switching cost, even though our suggested policy is arguably less invasive. Second, we prove that our results remain qualitatively unchanged if there is competition among firms or if there is a fraction of sophisticated present-biased consumers in addition to naive present-biased and rational consumers. Third, even if the policy maker can use deadlines as an additional policy instrument, our intuition about the timing of choice-enhancing policies survives. Only in a set-up where a firm sells a base product and an add-on, there is no competition on the base product, and consumers have heterogeneous valuations for the products, our results with respect to the specific policies potentially change. However, our result that the timing of policies matters remains robust. Therefore, it is important for policy design to take into account the equilibrium response of firms, because otherwise well-intended policies may backfire.

Joint with Florian Englmaier and Matthias Fahn, I investigate in Chapter 3 how firms may exploit naive present-biased employees by providing them a menu of contracts. One contract in this menu, which we call the *virtual* contract, requires the employee to work hard in the first period and promises a lot of benefits thereafter. Because the employee thinks that he will work hard in the next period and enjoy the benefits after that, he is willing to work for less than his outside option in the current period. Therefore he chooses the other contract in the menu, the *real* contract, which gives the employee less than his outside option for the current period, but allows him to choose the virtual contract in the next period. However, when the next period comes, he rather opts for the real contract again, which does not require him to work so hard then, but promises the option to work hard in the following period with a lot of benefits thereafter. This continues forever unless the employee learns the degree of his present bias. This finding matches several situations with immediate costs but delayed benefits in employment settings, for example, the promises of a future career if one works hard today, or the lucrative international mobility that comes at the cost of organizing the stay abroad. An implication is that the employer can benefit from being able to offer a long-term contract rather than a series of short-term contracts, even though the employee's effort choice only affects the current period.

While it is not surprising given the literature on the exploitation of consumers (see e.g., DellaVigna and Malmendier, 2004, Eliaz and Spiegler, 2006, and Heidhues and Kőszegi, 2010) that a firm can exploit an employee who does not stick to his planned action, our model also makes several new predictions which are unique to the employment setting. First, if the employee is protected by limited liability, a naive present-bias agent may work harder in equilibrium than his rational or sophisticated counterpart. Usually, present-biased individuals appear lazy. The reason for our result is that the only way to exploit an employee whose wage and bonus is reduced to the minimum is to let him provide an inefficiently high effort. Second, if the time-horizon of the employment is finite, employees who will stay longer can be exploited more. This means that younger workers can be exploited more than older workers, which is in line with actual labor market data. The reason is that young workers can be incentivized more with promises for future periods. Third, our model suggests that lower outside options for the employer, e.g., because of higher firing costs, may hurt the employee. They enable the employer to more credibly commit to higher promises for the employee's future, which allows the employer to exploit the employee more today (and also in the future, because the employee always chooses the real contract). Furthermore, we find that an employee who is present-biased in the effort domain can benefit from a present-bias in the monetary domain. Then, future financial benefits seem less attractive and thus limit the extent of exploitation. Finally, we show that our results are qualitatively robust to implementing moral hazard, unobservable agent types, and bargaining.

In Chapter 4, I combine endogenous network formation with learning to investigate the question how social media impact belief formation in a society. In the model, agents have different types that can be ordered on a line and live in different regions. The cost of forming a link to another agent depends on how similar the type is and whether the other agent lives in the same region. Agents form links to other agents so that they can access these other agents' signals. I show that when the importance and availability of social media increases, i.e., the cost of forming links to agents in other regions decreases, it first decreases and then increases the amount of long-run disagreement in a society. The intuition is as follows. If the cost of forming links to agents in other regions is very high, agents only form links with agents in their own region and there is no information exchange across regions. Then agents in different regions do not converge to the same beliefs. If the cost of forming links to agents in other regions is intermediate, agents form links to agents in their own region and to agents of the same type in other regions. Eventually, all information is aggregated and all agents hold the same beliefs in the long run. If the cost of forming links to agents in other regions is very low, agents only form links with agents of the same type. Then there is no information exchange across types and different types disagree with each other. This effect is amplified if agents suffer from correlation neglect. The lack of information aggregation may not be important for one agent to let her form different links in the first place. However, on the aggregate level, it might have a big impact. For example, an agent's vote is unlikely to be pivotal in an election. Therefore, it does not make sense to form costly links to acquire information about the different options. If information is not aggregated in a society, however, the society is more likely to make the wrong choice. A policy maker who wants the society to make the right choice may try to encourage local information exchange, for example, by subsidizing local events or infrastructure, in order to prevent the segregation of the society.

The model also makes some further predictions. I find that agents of the same type hold increasingly homogeneous beliefs and therefore also make increasingly homogeneous decisions. For agents with very high or very low types, it is more costly than for agents with intermediate types to form links to others in their own region. Thus, they are more likely to interact with similar agents online and to consequently hold extreme beliefs. Moreover, I propose a simple way of modeling correlation neglect, i.e., when agents overlook the potential correlation of signals they receive. I derive conditions of when agents with correlation neglect benefit from additional signals, and show that these agents may be made worse off by the growing influence of social media, thereby widening the utility gap between rational agents and agents with correlation neglect.

All chapters have their own appendices that can be found after Chapter 4. At the end of this dissertation there is a bibliography containing the references of all chapters.

Chapter 1

Optimal Cost Overruns: Procurement Auctions with Renegotiation*

1.1 Introduction

Renegotiation of procurement contracts awarded by public authorities are ubiquitous. The initial contract is awarded via competitive tendering; i.e., via an auction. The terms of the initial contract, however, are often subject to renegotiation with the result that the ultimate price is (by far) higher than the price which the parties initially agreed upon. Prominent recent examples of public procurement projects that are by far more expensive than initially planned are the Elbphilharmonie, a concert hall in Hamburg, the Big Dig, a highway artery in Boston, and the North-South metro line in Amsterdam.[1] What is often considered as the most severe case of a cost overrun in modern construction history is the Sydney Opera House.[2] The project was completed 10 years late at a price of 14.6 times the initial price.[3] During the construction of the Sydney Opera House, plenty of design changes had taken place. For instance, the change to the ribbed ellipsoidal roof increased the cost for the roof by 65%. Also the use of the two halls changed during construction. The major hall was originally planned as a multipurpose opera/concert hall but became solely a concert hall (the Concert Hall). The minor hall, called the Joan Sutherland Theatre today, had been adjusted

*This chapter is based on joint work with Fabian Herweg.

[1] Regarding the Elbphilharmonie the accepted offer from the underwriting group in 2006 was 241 million euro. The current costs are estimated at around 800 million euro and the project is not yet completed (Kostka and Anzinger 2015). For the Amsterdam metro line the initial budget was set at 1.46 billion euro in 2002 but the costs had risen to 3.1 billion euro in 2009. Recent estimates suggest that it will be completed in 2017 (Chang, Salmon and Saral 2016). For the Boston highway artery the ultimate price exceeded the initial price by 1.6 billion US dollar (Bajari, Houghton and Tadelis 2014).

[2] We define as cost overrun the difference between the final price and the initial price at which the procurement order has been awarded.

[3] When controlling for inflation, the cost overrun reduces to a factor of 7.5 (Newton, Skitmore and Love 2014).

so that it can incorporate opera performances. All these changes that arguably made the Sydney opera House more complex are responsible for a large part of its cost overrun.[4]

While the costs of complex and large-scale projects are hard to estimate, the ultimate price should on average coincide with the initial estimates if the estimates are unbiased. Casual observations and empirical studies, however, suggest that the prices typically increase through renegotiation.[5] According to public opinion, these cost overruns are sign of inefficient project management by bureaucrats or of strategizing politicians and thus a waste of taxpayers' money. In contrast to this widespread public opinion, we argue that these seemingly inefficient cost overruns can be the result of a constrained-optimal award procedure that minimizes the expected final price for the procurer.

The baseline model is fairly simple. A procurer needs an indivisible good or service, which can take one of several designs; i.e., the good can be more or less complex (a bridge with two or three traffic lanes). The good can be delivered by several suppliers that may differ in their privately known production costs. Moreover, the ex post efficient design depends on the contractor's production cost – i.e., on the cost type of the supplier who has been awarded with the contract. Initially the procurer runs an auction in order to allocate the contract. Importantly, the contract specified by the auction is a specific performance contract that can be enforced by courts. First, we assume that the procurer can collect bids only on prices and thus has to select one particular design of the good. More precisely, the procurement contract for the given design is awarded to one supplier via a standard second price sealed bid auction. The design specified in the initial contract may turn out to be inefficient, given the cost type of the contractor. In this case, we assume that the parties engage in Coasian bargaining and implement the efficient good ex post. Renegotiation is expected by the suppliers and thus incorporated in their bidding behavior. The rent the contractor (the supplier who won the auction) receives depends on his cost advantage compared to the second-lowest bidder with regard to the initial design. If cost differences are more pronounced for more complex designs, then it is optimal for the procurer to fix a rather simple design ex ante because this enhances competition in the initial auction. In other words, when commitment not to renegotiate is not feasible, it is optimal for the procurer to choose a simple design ex ante and to renegotiate to a more complex and costlier design ex post. The outcome is efficient because the supplier with the lowest cost wins the auction. He can not only produce the initial design at lower cost but also benefits more from contract renegotiation than suppliers with higher costs. Moreover, the outcome is constrained-optimal for the procurer because it minimizes the ultimate price she has to pay.

[4]See Newton et al. (2014) and Drew (1999) for more detailed discussions of the construction and cost increases of the Sydney Opera House.

[5]Substantial price increases resulting from contract renegotiation are reported by Decarolis (2014) for Italian procurement contracts and by Bajari et al. (2014) for Californian procurement contracts. German procurement contracts and the their cost increases are listed by Kostka and Anzinger (2015). They also report that some projects perform exceptionally well. For instance, the Chemikum, a building of the University of Erlangen-Nuremberg, was completed a year earlier than planned and at a cost of only 80 million euro instead of the planned 140 million euro.

In an extension we allow the procurer to set a public maximum bid. We show that our main findings are robust regarding the introduction of a maximum bid. It is still optimal to choose a simple design ex ante so that contract renegotiation always leads to a more complex design which requires an upward price adjustment.

In our general model, we allow for multi-dimensional auctions – i.e., scoring auctions. The procurer now asks for bids containing a price and a design. The supplier who places the bid leading to the highest score – determined by a commonly known scoring rule – wins the auction. The procurer's initial choice is the scoring function, which we restrict to be linear in price. If the scoring function reflects the procurer's true preferences, each supplier offers the optimal design given his cost. In this case, contract renegotiation can be avoided. The optimal scoring function, however, does not reflect the procurer's true preferences. It depends solely on the price offer and not on the offered design. Thus, all suppliers offer the most simple and cheapest design. The most efficient supplier wins the initial auction and the parties agree to implement a more complex design at a higher price ex post via renegotiation. In other words, a standard auction for the simplest design (only price bids are collected) outperforms scoring auctions, where suppliers place multi-dimensional bids containing a price and a design. The reason is that a multi-dimensional auction allows for differentiation of the suppliers' bids, which relaxes competition between suppliers ex ante and thus leads to higher ultimate prices. This finding is in contrast to the existing literature on scoring auction that assumes the procurer can commit not to renegotiate the contract (Dasgupta and Spulber 1989-1990, Che 1993, Chen-Ritzo, Harrison, Kwasnica and Thomas 2005).

As a robustness check, we analyze what happens when there is an exogenous risk that the renegotiation breaks down and the parties are stuck with the initial contract. If this risk is rather low, it is still optimal to choose a simple design ex ante. As the risk becomes larger, taking into account the situation when renegotiation fails becomes more important, so the optimal design becomes more complex and closer to the ex post optimal design. However, as we demonstrate in an example, upward price adjustments seem to be by far more likely than downward price adjustments if renegotiation takes place; i.e., if there had been a risk of bargaining breakdown but the parties succeed in finding an agreement.

The paper is structured as follows. After having discussed the related literature, which is done in the following paragraphs, we introduce the model in Section 1.2. The model is analyzed in Section 1.3. This section concludes with our first main finding; Renegotiation of the contract determined by the constrained-optimal auction always leads to an upward price adjustment (Proposition 1.1). We extend the baseline model and allow for multi-dimensional auctions in Section 1.4, where we show that price-only auctions are optimal (Proposition 1.3). We critically discuss our assumptions in Section 1.5 and, in particular, show that our main findings are robust toward introducing a risk of bargaining breakdown. The final Section 1.6 summarizes our findings and concludes. All proofs that are not presented in the main text are deferred to Appendix B.

Related Literature

Investigation of procurement contracts is an important and classic topic of contract theory.[6] A seminal contribution analyzing procurement and renegotiation is Tirole (1986). He analyzes the contractual relationship between a single procurer and a single supplier with a focus on how initial contracts can enhance non-contractible relationship specific investments.

Dasgupta and Spulber (1989-1990), Che (1993), and Chen-Ritzo et al. (2005) analyze procurement auctions for the case that the procurer can commit not to renegotiate the contract. All three articles show that the optimal scoring auction outperforms price-only auctions. We demonstrate that if this commitment is absent, the optimal price-only auction outperforms scoring auctions.

There is only a small extent literature that analyzes auctions without perfect commitment; i.e., that allow either bidders to renege on their bids or to engage in contract renegotiation.[7] Waehrer (1995), Harstad and Rothkopf (1995), and Roelofs (2002) allow bidders to withdraw the winning bid ex post. In these models suppliers are initially uncertain about their costs and thus may underestimate it. The possibility to default on the initial commitments enhances competition in the auction, which in turn is beneficial to the procurer.[8] Waehrer (1995) also analyzes a scenario where the procurer and the winner renegotiate a new contract. Here, however, renegotiation takes place after the default of the winner and thus the initial contract has no impact on the outcome of renegotiation.

A similar form of renegotiation is analyzed by Wang (2000) and Shachat and Tan (2015). In these models the procurer either accepts the lowest bid or rejects all bids. In case of rejection, the procurer negotiates with the supplier who placed the lowest bid; i.e., if renegotiation takes place the initial contract concluded by the auction is not binding. In such a setup renegotiation always leads to lower prices, which is exactly the opposite from what we study.

The initial contract has an impact on the outcome of renegotiation in Chang et al. (2016). Here, suppliers' production costs have an ex-ante unknown common component. Some of the suppliers are wealth constraint, while other have deep pockets and this is private information of each supplier. Allowing for contract renegotiation is advantageous to wealth-constraint suppliers who can credibly threaten to default. The prices increase with renegotiation in order to avoid bankruptcy of the contractor who is faced by unexpectedly high costs. In our model, the parties agree to a different design of the project ex post, which is more costly to produce and thus the final price exceeds the initial price.

[6]For an excellent discussion of the standard contract theoretical analysis of procurement see Laffont and Tirole (1993).

[7]There is also a small literature that analyzes screening mechanisms if the principal (the procurer) cannot commit not to renegotiate; e.g., Beaudry and Poitevin (1995). This literature typically assumes that there is only one buyer and one seller and focuses on the constraints limited commitment power imposes on the implementable allocations.

[8]The effects of limited liability on more general mechanisms than auctions are investigated by Burguet, Ganuza and Hauk (2012).

A few papers directly deal with the issue of cost overruns. Birulin and Izmalkov (2013) analyze what shares of a price are optimally paid before and after a potential extra cost to the supplier realizes when suppliers are protected by limited liability. This paper is orthogonal to this dissertation chapter, because it does not provide an explanation for cost overruns but rather assumes its existence.[9] Closer to our work is Ganuza (2007). Here, suppliers are differentiated à la Salop (1979). The buyer does not know her preferences – her location – but can invest in obtaining a noisy signal. A procurement order for the expected optimal design is awarded via a price-only auction. Ex post, the buyer's preferences are common knowledge and the winner of the auction can make a take-it-or-leave-it renegotiation offer. The main result is that the buyer under-invests in learning her preferences because this enhances competition in the initial auction. This is related to our results that the initial design is chosen to enhance competition in the initial auction. There are, however, crucial differences. For instance, in our model the buyer's preferences are known ex ante and suppliers are not horizontally differentiated.

Finally, costly renegotiation of incomplete procurement contracts is analyzed by Bajari and Tadelis (2001) and Herweg and Schmidt (2016). The former paper analyzes when fixed-price contracts outperform cost-plus contracts, while the latter one derives conditions so that bilateral negotiations outperform procurement auctions.[10]

1.2 The Model

1.2.1 Players and Payoffs

A procurer P, say a government agency, wants to buy one unit of an indivisible good, e.g., a bridge. The good can be produced with one of various designs $x \in \{x_1, \dots, x_J\} \equiv X$, with $J \geq 2$. The procurer's valuation of the good depends on the design x and is denoted by $v(x)$. There are $n \geq 2$ suppliers, indexed by $i = 1, \dots, n$, that can produce the good required by the procurer. A supplier's production cost depends on the design x and his cost type $\theta \in [\underline{\theta}, \overline{\theta}] \equiv \Theta$, and is denoted by $c(x, \theta)$. Ex ante, the cost type θ is private information of each supplier. The n cost types are drawn independently according to an identical cumulative distribution function $F(\theta)$. Let the corresponding probability density be $f(\theta)$, with $f(\theta) > 0$ for all $\theta \in \Theta$. When design x is delivered at price p, the procurer's and the contractor's, i.e., the selected supplier's, ex post utilities are

$$\begin{aligned} u &= v(x) - p && (1.1) \\ \text{and} \quad \pi &= p - c(x, \theta), && (1.2) \end{aligned}$$

respectively. All parties are assumed to be risk neutral and the outside option utilities are all set equal to zero.

[9] A similar model is analyzed by Birulin (2014).

[10] An empirical analysis of incomplete procurement contracts is provided by Crocker and Reynolds (1993). They argue that contracts are left incomplete intentionally to economize on the cost of the ex ante design.

We denote the ex post efficient design by

$$x^*(\theta) \in \arg\max_{x \in X}\{v(x) - c(x,\theta)\} \tag{1.3}$$

and the corresponding social surplus by

$$S(\theta) = \max_{x \in X}\{v(x) - c(x,\theta)\}. \tag{1.4}$$

We make the following assumptions on the designs, the procurer's valuation, and the suppliers' cost functions. First, we posit that the designs can be ordered so that designs that are ordered higher are associated with higher gross benefits for the procurer but are more costly to produce for the suppliers.[11]

Assumption 1.1 (designs are ordered)**.** *For all $j \in \{1, \ldots, J-1\}$ and all $\theta \in \Theta$:*

(i) $v(x_j) \leq v(x_{j+1})$*;*

(ii) $c(x_j, \theta) \leq c(x_{j+1}, \theta)$*.*

Secondly, regarding the impact of the type θ on the production cost, we assume that lower types are more efficient.

Assumption 1.2 (Types are ordered)**.** *For all $\theta_1, \theta_2 \in \Theta$ with $\theta_1 < \theta_2$ and all $x \in X$:*

$$c(x, \theta_1) < c(x, \theta_2).$$

Assumption 1.2 implies that $S(\theta)$ is decreasing.[12]

Finally, we assume that the cost advantage is lowest for the cheapest design x_1. For more complex designs, the differences in costs between suppliers of different types are larger.

Assumption 1.3 (Increasing differences)**.** *For all $\theta_1, \theta_2 \in \Theta$ with $\theta_1 < \theta_2$ and all $x \in X \backslash x_1$ it holds that:*

$$c(x, \theta_2) - c(x, \theta_1) > c(x_1, \theta_2) - c(x_1, \theta_1).$$

This assumption is weaker than assuming that the cost advantage is increasing in the complexity of the design; i.e., assuming that the costlier the design is the more important it is to select the most efficient supplier.[13] The latter assumption corresponds to the usual

[11] Designs that lead to higher costs and lower benefits than some other design, will never be implemented. All the designs that are not dominated by another design can be ordered so that higher elements are associated with higher benefits and higher costs. Thus, Assumption 1.1 is merely a relabeling of the designs.

[12] We impose Assumption 1.2 for simplicity but it is not necessary. In fact, we can allow for $c(x, \theta_1) > c(x, \theta_2)$ for some x, as long as $S(\theta_1)$ is sufficiently larger than $S(\theta_2)$.

[13] Formally: For all $\theta_1 < \theta_2$ and $x_l, x_H \in X$ with $L < H$:

$$c(x_H, \theta_2) - c(x_H, \theta_1) > c(x_L, \theta_2) - c(x_L, \theta_1).$$

single-crossing property, which is a standard assumption in the auction literature (Krishna 2010).

1.2.2 Award Procedure and Renegotiation

The procurement contract is awarded to one supplier, called the contractor, via an auction. For now, we assume that the procurer has to use a simple auction; i.e., a second price sealed bid auction without maximum bid for a pre-specified design $\overline{x} \in X$.[14] In other words, the procurer can collect only price bids for a given design. A scoring auction which maps bids containing a price and a design into a single score cannot be used. Moreover, we assume that the used procurement mechanism is a second-price sealed bid auction without maximum bid. Not specifying a maximum bid is optimal if the procurer has to buy the good for sure – i.e., the bridge is urgently needed and thus the loss in surplus is huge if no construction contract is signed.

When specifying the initial auction, the procurer has only one choice variable, the design $\overline{x} \in X$. The procurement order for the good with design $\overline{x} \in X$ is auctioned off between the n suppliers. Each supplier i places a secret price bid p_i. The supplier with the lowest bid is selected as the contractor and the specified price equals the second lowest bid. If the lowest bid is made by several suppliers, one of these suppliers is selected at random as the contractor.

With the procurer being restricted to simple auctions and suppliers' types being stochastic, the initial design $\overline{x}$ may not be optimal given the contractor's type ex post. In this case there is scope for renegotiation. We posit that the contractor's type is observed by the procurer after the award of the contract and thus the parties engage in (efficient) Coasian bargaining ex post. In practice, the contractor starts working on the project before the parties agree to renegotiate. The procurer monitors the contractor and thus obtains an informative signal – next to the contractor's bid – about its efficiency. For simplicity we focus on the extreme case where the contractor's type is perfectly observed.[15] Doubtlessly, this is a strong assumption, which (partly) drives our results but also simplifies the exposition significantly.[16] The surplus from renegotiation is split between the procurer and the contractor according to the generalized Nash bargaining solution (GNBS), i.e., the renegotiation contract is

$$(\hat{x}, \hat{p}) \in \arg \max_{x \in X, p \in \mathbb{R}} [p - c(x, \theta) - d_C]^{\alpha} \times [v(x) - p - d_P]^{1-\alpha}, \tag{1.5}$$

[14] The restriction to a second price auction is without loss of generality because the revenue equivalence theorem holds in our model. This implies in particular that the main propositions also hold in case of a first price auction. The restriction to simple auctions – i.e., neither a maximum bid nor a scoring auction, is relaxed at the end of Section 1.3 and in Section 1.4, respectively.

[15] Alternatively the procurer might receive a signal about the contractor's type due to information acquisition which would have been too costly before the selection of a certain supplier.

[16] If the true type of the contractor is observed ex post, the procurer could use a mechanism that makes payments contingent on the true type. For a survey of this literature see Skrzypacz (2013). This chapter is more in the spirit of the literature on incomplete contracts and therefore does not allow for such mechanisms.

where $\alpha \in (0,1)$ denotes the contractor's relative bargaining power ex post.[17] The disagreement payoffs of the two parties are determined by the initial contract $(\overline{x}, \overline{p})$:[18]

$$d_C = \overline{p} - c(\overline{x}, \theta) \tag{1.6}$$

$$d_P = v(\overline{x}) - \overline{p}. \tag{1.7}$$

The sequence of events is described in Figure 1.1. We employ perfect Bayesian equilibrium in symmetric strategies as equilibrium concept.

- Nature draws $\theta_1, ..., \theta_n$.
- Procurer selects $\overline{x}$ and auctions off contract.

- Suppliers submit bids.
- Lowest bidder wins contract $(\overline{x}, \overline{p})$, where $\overline{p}$ is the second lowest bid.

- Contractor's cost type is observed by procurer.
- Parties may renegotiate to a new contract $(\hat{x}, \hat{p})$.

Figure 1.1: Timeline.

1.3 The Analysis

1.3.1 Contract Renegotiation and Bidding Behavior

We start the analysis with the renegotiation game. Suppose the procurer awarded a supplier with cost type θ with the contract $(\overline{x}, \overline{p})$. If the design $\overline{x}$ is not the efficient design given the contractor's cost type, $\overline{x} \neq x^*(\theta)$, then there is scope for renegotiation. The social surplus can be increased by moving $\overline{x}$ to $x^*(\theta)$ and this additional surplus is split between the two parties according to their relative bargaining power. The outcome of renegotiation is characterized in the following lemma.

Lemma 1.1 (Renegotiation outcome)**.** *Let $(\overline{x}, \overline{p})$ be the initial contract and suppose the contractor's cost type is θ. Ex post, the procurer and the contractor agree to trade design $\hat{x} = x^*(\theta)$ at price*

$$\hat{p} = \overline{p} + \alpha[v(x^*(\theta)) - v(\overline{x})] + (1-\alpha)[c(x^*(\theta), \theta) - c(\overline{x}, \theta)].$$

A supplier taking part in the auction is aware that the contract may be renegotiated ex post. In particular, he knows that if he wins the auction, he may obtain additional profits generated by contract renegotiation. These additional profits from renegotiation are

[17] For a detailed description of the Nash bargaining solution see Muthoo (1999). A non-cooperative foundation for the Nash bargaining solution is provided by Binmore, Rubinstein and Wolinsky (1986).

[18] Exactly the same findings are obtained with an alternative bargaining game, where the GNBS is replaced by a take-it-or-leave-it offer game. With probability α the contractor can make a take-it-or-leave-it offer in the renegotiation game, while with probability $(1-\alpha)$ the procurer can.

incorporated in a supplier's bidding behavior. Supplier θ's ex post payoff from being awarded with the procurement contract $(\overline{x}, \overline{p})$ is

$$\begin{aligned} \pi(\overline{x}, \overline{p}, \theta) &= \hat{p}(\overline{p}, \theta) - c(x^*(\theta), \theta) \\ &= \overline{p} + \alpha[v(x^*(\theta)) - c(x^*(\theta), \theta)] - \alpha v(\overline{x}) - (1-\alpha)c(\overline{x}, \theta). \end{aligned} \tag{1.8}$$

The price bid affects directly the probability of winning the auction but only indirectly the price the supplier receives when being awarded with the contract. Thus, placing the lowest feasible bid that allows the supplier to break-even, even when he is awarded with the contract at a price equal to his bid, is optimal. The equilibrium bidding behavior is formally described in the next lemma.

Lemma 1.2 (Outcome of the auction)**.** *Suppose that Assumption 1.2 holds and that the procurement order for design $\overline{x}$ is awarded via a second price sealed bid auction. In the unique equilibrium in undominated strategies, each supplier uses the bidding function*

$$\overline{p}(\theta) = \alpha v(\overline{x}) + (1-\alpha)c(\overline{x}, \theta) - \alpha S(\theta)$$

and the supplier with the lowest type wins the auction.

It is important to note that – according to Lemma 1.2 – the auction selects the most efficient supplier. In other words, productive efficiency is still guaranteed by a second price auction even if contract renegotiation is feasible. This relies on the assumption that a more efficient type has not only lower production costs for producing design $\overline{x}$ but also generates a higher surplus by adjusting the design via contract renegotiation.

1.3.2 Constraint Optimal Auction

The procurer does not only care about ex post efficiency but also about the ultimate price she has to pay for the good. The initial price – i.e., the price specified in the procurement contract, is determined by the auction and depends on the cost of the second lowest type, type $\ddot{\theta}$. It is given by

$$\overline{p}(\ddot{\theta}) = \alpha v(\overline{x}) + (1-\alpha)c(\overline{x}, \ddot{\theta}) - \alpha S(\ddot{\theta}). \tag{1.9}$$

The ultimate price, paid by the procurer and received by the contractor, depends not only on the second lowest type but also on the lowest type $\dot{\theta}$. The procurer's ex post utility, for given realizations of $\dot{\theta}$ and $\ddot{\theta}$, is

$$\begin{aligned} u(\overline{x}, \dot{\theta}, \ddot{\theta}) &= v(x^*(\dot{\theta})) - \overline{p}(\ddot{\theta}) - \alpha[v(x^*(\dot{\theta})) - v(\overline{x})] \\ &\qquad - (1-\alpha)[c(x^*(\dot{\theta}), \dot{\theta}) - c(\overline{x}, \dot{\theta})] \\ &= (1-\alpha)S(\dot{\theta}) + \alpha S(\ddot{\theta}) - (1-\alpha)[c(\overline{x}, \ddot{\theta}) - c(\overline{x}, \dot{\theta})]. \end{aligned} \tag{1.10}$$

The first part of the procurer's ex post utility can be written as $S(\ddot{\theta}) + (1-\alpha)[S(\dot{\theta}) - S(\ddot{\theta})]$; i.e., the procurer obtains the whole surplus generated by the second most efficient type due to the competitive award procedure. On top of that, the procurer obtains the share $1-\alpha$ of the rents that are generated by the excess efficiency of type $\dot{\theta}$ compared to type $\ddot{\theta}$. This, however, is only half the story. Different supplier types benefit differently from contract renegotiation ex post. The most efficient type benefits more from contract renegotiation than the second most efficient type because he can produce design $\overline{x}$ at lower cost. Therefore, the contractor obtains a rent which equals his advantage from contract renegotiation as compared to type $\ddot{\theta}$, plus the share α of the additional surplus that he generates, $\alpha[S(\dot{\theta}) - S(\ddot{\theta})]$.

Now, we can state the first main finding of the paper.

Proposition 1.1. *Suppose that Assumptions 1.1-1.3 hold. The procurer optimally chooses design $\overline{x} = x_1$ ex ante. If and only if $x^*(\dot{\theta}) \neq x_1$, renegotiation takes place and the ultimate price exceeds the initial price determined by the auction; i.e., $x^*(\dot{\theta}) \neq x_1 \iff \hat{p} - \overline{p} > 0$. Moreover, the price increase $\hat{p} - \overline{p}$ is increasing in the contractor's ex post bargaining power.*

Proof. The procurer's expected utility ex ante is

$$\mathbb{E}[u(\overline{x}, \dot{\theta}, \ddot{\theta})] = \mathbb{E}\left[(1-\alpha)S(\dot{\theta}) + \alpha S(\ddot{\theta}) - (1-\alpha)[c(\overline{x}, \ddot{\theta}) - c(\overline{x}, \dot{\theta})]\right]. \tag{1.11}$$

The expected utility is maximized by the design $\overline{x} \in X$ that minimizes

$$\mathbb{E}[c(\overline{x}, \ddot{\theta}) - c(\overline{x}, \dot{\theta})]. \tag{1.12}$$

By Assumption 1.3 the above expression is minimized for $\overline{x} = x_1$.

From Lemma 1.1 it is readily obtained that $\hat{p} - \overline{p} > 0$ if renegotiation takes place and Assumption 1.1 holds. Moreover, from Lemma 1.1 it follows immediately that $\hat{p} - \overline{p}$ is increasing in α. □

The contractor's bid already reflects that renegotiation may take place. In other words, part of the contractor's profits made by contract renegotiation are competed away in the initial auction. The profits from renegotiation that are not competed away can be decomposed into two parts. The first part is the additional surplus the contractor generates compared to the second most efficient supplier, $S(\dot{\theta}) - S(\ddot{\theta})$. The second part, $c(\overline{x}, \ddot{\theta}) - c(\overline{x}, \dot{\theta})$, is due to the fact that the contractor's disagreement payoff is higher than the one of the second most efficient supplier; i.e., the contractor can produce $\overline{x}$ at lower costs than all other suppliers. From the procurer's perspective, the first part is a random variable, which does not depend on her choice variable, the initial project design $\overline{x}$. The second part, on the other hand, depends on the initial design. The more complex the initial design is, the larger is the difference in disagreement payoffs between suppliers of different types. Hence, in order to minimize this difference, the procurer optimally specifies the most simple design ex ante.[19]

[19] If $\alpha = 1$, i.e., the contractor has all the bargaining power, the procurer is indifferent between all $\overline{x} \in X$. In this case the revelation of the contractor's type is irrelevant and we are in the standard framework: The payoffs are pinned down by the utility of the least efficient type in combination with the incentive constraints; i.e., all types are revealed truthfully (Myerson 1981).

When renegotiation is inefficient or there is a risk it might breakdown, however, there is a trade-off and the optimal ex ante design is not necessarily the most simple one any more, as we demonstrate in Section 1.5.2.

1.3.3 Maximum Price Bid

Without renegotiation, a second-price auction with a maximum bid (a reserve price in a selling context) is an optimal mechanism. Proposition 1.1 crucially relies on the assumption that the difference in production costs between the most efficient and the second most efficient type is minimized for the simplest design. If the initial price is sometimes determined by the maximum bid instead of the production costs of the second-lowest type, design $\overline{x} = x_1$ might not be optimal anymore. As we will show, this reasoning is not true and the simplest design is also optimal when a maximum bid is specified.

Now, we suppose the procurer specifies a maximum bid R; i.e., only price bids $p \leq R$ are allowed in the second-price auction. The maximum bid is publicly announced before suppliers place their bids. The supplier who places the lowest bid is awarded with the procurement contract. If at least two suppliers placed an admitted bid, the price equals the second lowest bid. If only one supplier placed an admitted bid, then the price equals the maximum bid R. As before, the contractor's type is observed by the procurer ex post and the parties may renegotiate the contract. If none of the suppliers placed an admitted bid, no initial contract is awarded. In this case, the procurer does not buy the good and there is also no (re)negotiation.

The optimal bidding strategy for a supplier who takes part in this auction with maximum bid is the same as before; i.e., a supplier of type θ places the bid

$$\overline{p}(\theta) = \alpha v(\overline{x}) + (1 - \alpha)c(\overline{x}, \theta) - \alpha S(\theta).$$

This bid is admitted only if $\overline{p}(\theta) \leq R$. Thus, suppliers with types so that $\overline{p}(\theta) > R$ do not submit a bid. Let r denote the highest type who places an admitted bid; i.e., $\overline{p}(r) \equiv R$. Notice that there is a one-to-one relationship between R and r. Thus, we can formulate the procurer's problem as a problem of choosing a design $\overline{x}$ and threshold type r.

If $\overline{p}(\dot{\theta}) \leq \overline{p}(\ddot{\theta}) \leq R$, the maximum bid is not binding and the procurer's ex post utility is the same as before. If, on the other hand, $\overline{p}(\dot{\theta}) \leq R < \overline{p}(\ddot{\theta})$, then the initial price is equal to the maximum bid and thus given by

$$R \equiv \overline{p}(r) = \alpha v(\overline{x}) + (1 - \alpha)c(\overline{x}, r) - \alpha S(r). \tag{1.13}$$

The procurer's ex post utility in this case is given by

$$\begin{aligned} u(\overline{x}, \dot{\theta}, r) &= v(x^*(\dot{\theta})) - R - \alpha[v(x^*(\dot{\theta})) - v(\overline{x})] \\ &\qquad - (1-\alpha)[c(x^*(\dot{\theta}), \dot{\theta}) - c(\overline{x}, \dot{\theta})] \\ &= \alpha S(r) + (1-\alpha)S(\dot{\theta}) - (1-\alpha)[c(\overline{x}, r) - c(\overline{x}, \dot{\theta})]. \end{aligned} \tag{1.14}$$

For any threshold type $r \in (\underline{\theta}, \overline{\theta})$ the above expression is maximized at the simplest design $\overline{x} = x_1$ due to Assumption 1.3. Thus, initially specifying a simple design that will probably be renegotiated to a more complex and costlier design is optimal even when a maximum bid is specified by the procurer.

Proposition 1.2. *For any maximum bid $R \in (\overline{p}(\underline{\theta}), \overline{p}(\overline{\theta}))$, the procurer optimally chooses design $\overline{x} = x_1$ ex ante. If $x^*(\dot{\theta}) \neq x_1$, renegotiation takes place and the ultimate price exceeds the initial price.*

1.4 Scoring Auctions

So far we assumed that the procurer has to specify the good she wants to procure completely ex ante, i.e., before the auction takes place. In the auction, the procurer collected bids only on prices and the supplier who offered the lowest price has been awarded with the contract. Different types of suppliers do not only have different production costs but also differ in the optimal design – i.e., the design that maximizes the joint surplus. Therefore, it may be profitable for the procurer to ask suppliers for bids on price *and* design.[20]

1.4.1 The Model with Multi-Dimensional Auctions

In the following we consider a second score auction. Each supplier places a bid containing a price $p \in \mathbb{R}$ and a design $x \in X$. Each bid (x, p) is mapped into a single score. The supplier who placed the bid giving rise to the highest score wins the auction and is required to match the highest rejected score – i.e., the second highest score. The outcome $(\overline{x}, \overline{p})$ determines a binding specific-performance contract between the procurer and the winner (the contractor). Nevertheless, this contract can be renegotiated after the auction as before.

The procurer does not choose a design when particularizing the auction. She specifies a scoring function, $G: X \times \mathbb{R} \to \mathbb{R}$, that maps bids into a single score. We focus on quasi-linear scoring functions of the form

$$G(x, p) = g(x) - p.$$

If the procurer can commit not to renegotiate the contract, the optimal quasi-linear scoring function implements the second-best allocation (Che 1993). Here, the procurer cannot com-

[20] Scoring auctions where bids are multi-dimensional (e.g., price and quality) are analyzed by Che (1993) and Asker and Cantillon (2008). An excellent short review of this literature is provided by Asker and Cantillon (2010).

mit not to engage in contract renegotiation. However, as we will show below, if the scoring function represents the procurer's true preferences, i.e., $g(x) \equiv v(x)$, contract renegotiation can be avoided.

1.4.2 The Analysis of Multi-Dimensional Auctions

As before, we solve the game by backward induction. The outcome of the renegotiation game is independent of the award procedure. In other words, Lemma 1.1 still holds and the implemented design will always be ex post efficient. Thus, the ex post utility of a supplier of type θ who has been awarded procurement contract $(\overline{x}, \overline{p})$ is

$$\begin{aligned}\pi(\overline{x}, \overline{p}, \theta) &= \hat{p}(\overline{x}, \overline{p}, \theta) - c(x^*(\theta), \theta) \\ &= \overline{p} + \alpha S(\theta) - \alpha v(\overline{x}) - (1-\alpha)c(\overline{x}, \theta). \end{aligned} \tag{1.15}$$

Optimal bidding behavior in the second-score auction is described by the following result.

Lemma 1.3. *The (reduced) second score auction game has a dominant strategy equilibrium. The equilibrium bid of each supplier of type θ is*

$$\begin{aligned}x^b(\theta) &\in arg\max_{x \in X} \{g(x) - \alpha v(x) - (1-\alpha)c(x, \theta)\}, \\ p^b(\theta) &= \alpha v(x^b(\theta)) + (1-\alpha)c(x^b(\theta), \theta) - \alpha S(\theta).\end{aligned}$$

By Lemma 1.3, each supplier bids the optimal design, $x^b(\theta) = x^*(\theta)$, if the scoring function represents the procurer's true preferences – i.e., if $v(x) \equiv g(x)$. If, on the other hand, the scoring function does not reflect the true preferences of the procurer, then it is likely that suppliers propose designs that are not efficient.

According to Lemma 1.3, the score offered by a supplier of type θ amounts to

$$\begin{aligned}G(\theta) &\equiv g(x^b(\theta)) - p^b(\theta) \\ &= g(x^b(\theta)) - \alpha v(x^b(\theta)) - (1-\alpha)c(x^b(\theta), \theta) + \alpha S(\theta).\end{aligned} \tag{1.16}$$

As before, the most efficient type places the bid that leads to the highest score and thus wins the auction.

Lemma 1.4. *Suppose that Assumption 1.2 holds. Then for all $\theta_1 < \theta_2$ it holds that:*

$$G(\theta_1) > G(\theta_2).$$

The winner of the auction, type $\dot{\theta}$, has to match the second highest score but is otherwise free in its choice of $(\overline{x}, \overline{p})$. Thus, the winner of the auction chooses the initial contract $(\overline{x}, \overline{p})$ in order to maximize

$$\overline{p} + \alpha S(\dot{\theta}) - \alpha v(\overline{x}) - (1-\alpha)c(\overline{x}, \theta)$$

subject to $g(\overline{x}) - \overline{p} = G(\ddot{\theta})$. Hence, the initial contract specifies $\overline{x}(\dot{\theta}, \ddot{\theta}) = x^b(\dot{\theta})$ and $\overline{p}(\dot{\theta}, \ddot{\theta}) = g(x^b(\dot{\theta})) - G(\ddot{\theta})$. The ultimate price paid by the procurer is

$$\begin{aligned}\hat{p}(\dot{\theta}, \ddot{\theta}) = g(x^b(\dot{\theta})) - g(x^b(\ddot{\theta})) + \alpha v(x^b(\ddot{\theta})) + (1-\alpha)c(x^b(\ddot{\theta}), \ddot{\theta}) \\ - \alpha S(\ddot{\theta}) + \alpha[v(x^*(\dot{\theta})) - v(x^b(\dot{\theta}))] \\ + (1-\alpha)[c(x^*(\dot{\theta}), \dot{\theta}) - c(x^b(\dot{\theta}), \dot{\theta})]. \quad (1.17)\end{aligned}$$

1.4.3 The Optimality of Price-Only Auctions

The procurer's ex post utility – for given realizations of $\dot{\theta}$ and $\ddot{\theta}$ – is given by

$$\begin{aligned}u(\dot{\theta}, \ddot{\theta}) &= v(x^*(\dot{\theta})) - \hat{p}(\dot{\theta}, \ddot{\theta}) \\ &= (1-\alpha)S(\dot{\theta}) + \alpha S(\ddot{\theta}) \\ &\quad + \{g(x^b(\ddot{\theta})) - \alpha v(x^b(\ddot{\theta})) - (1-\alpha)c(x^b(\ddot{\theta}), \ddot{\theta})\} \\ &\quad - \{g(x^b(\dot{\theta})) - \alpha v(x^b(\dot{\theta})) - (1-\alpha)c(x^b(\dot{\theta}), \dot{\theta})\}. \qquad (1.18)\end{aligned}$$

The procurer chooses the scoring function $g(\cdot)$ that maximizes her expected payoff $\mathbb{E}[u(\dot{\theta}, \ddot{\theta})]$. As it turns out, a constant scoring function – i.e., a scoring rule that depends only on the price – is optimal.

Proposition 1.3. *Suppose Assumptions 1.1–1.3 hold. Then, the optimal quasi-linear scoring rule is independent of the design, i.e.,* $g(x) = \overline{g} \in \mathbb{R}$. *Each supplier* θ *bids the design* $x^b(\theta) = x_1$. *Renegotiation takes place if* $x^*(\dot{\theta}) \neq x_1$. *In this case, the ultimate price exceeds the initial price,* $\hat{p} - \overline{p} > 0$.

Proof. Define

$$A(x, \theta) \equiv g(x) - \alpha v(x) - (1-\alpha)c(x, \theta)$$

and notice that $A(x, \theta)$ is maximized at $x = x^b(\theta)$. Moreover, for all $\theta_1 < \theta_2$ it holds that

$$A(x^b(\theta_1), \theta_1) \geq A(x^b(\theta_2), \theta_2), \qquad (1.19)$$

irrespective of the shape of the scoring function $g(\cdot)$. The procurer's maximization problem can be restated as:

$$\min \quad \mathbb{E}_{\dot{\theta}, \ddot{\theta}}[A(x^b(\dot{\theta}), \dot{\theta}) - A(x^b(\ddot{\theta}), \ddot{\theta})]. \qquad (1.20)$$

Note that

$$A(x^b(\dot{\theta}), \dot{\theta}) - A(x^b(\ddot{\theta}), \ddot{\theta}) \geq (1-\alpha)[c(x^b(\ddot{\theta}), \ddot{\theta}) - c(x^b(\ddot{\theta}), \dot{\theta})] \qquad (1.21)$$

The left-hand side coincides with the right-hand side if $x^b(\dot{\theta}) = x^b(\ddot{\theta})$. Moreover, by Assumption 1.3, the right-hand side – the lower bound – is minimized if $x^b(\ddot{\theta}) = x_1$. Thus, if for all possible $\dot{\theta} \leq \ddot{\theta}$ we have $x^b(\dot{\theta}) = x^b(\ddot{\theta}) = x_1$, then the corresponding $g(\cdot)$ function is a solution to the above minimization problem. For $g(x) = \overline{g}$, each supplier bids the design x_1. This follows from Lemma 1.3 and Assumption 1.1. Hence, $G(x,p) = -p$ is an optimal scoring function.

The price mark-up in case renegotiation takes place is

$$\hat{p}(\dot{\theta},\ddot{\theta}) - \overline{p} = \alpha[v(x^*(\dot{\theta})) - v(x^b(\dot{\theta}))] + (1-\alpha)[c(x^*(\dot{\theta}),\dot{\theta}) - c(x^b(\dot{\theta}),\dot{\theta})], \tag{1.22}$$

where $x^b(\dot{\theta}) = x_1$. Hence, if $x^*(\dot{\theta}) \neq x_1$ – renegotiation takes place, then $\hat{p} - \overline{p} > 0$, by Assumption 1.1. □

According to Proposition 1.3, if the buyer is unable to commit not to renegotiate, she cannot benefit from using a scoring auction. A scoring auction by its multi-dimensionality allows suppliers to differentiate their bids, which reduces price competition. In other words, a more efficient supplier can offer a design that leads to a higher score than a less efficient supplier. By doing so the more efficient supplier may be able to win the auction even if his price bid is relatively high. This makes the usage of a scoring auction expensive and thus less attractive to the procurer. Hence, a simple auction where the procurer collects price bids for a given design is optimal. The given design is rather simple, so that the differences between suppliers regarding their costs for delivering this design are relatively low. This enhances the competition at the auction stage and leads to a very low initial price. Even though the ex post price can be significantly higher than the initial price, the effect on the initial price dominates.

Moreover, Proposition 1.3 implies that avoiding renegotiation by specifying an appropriate scoring function is not in the procurer's best interest.

Corollary 1.1. *For $\alpha \in (0,1)$ the procurer strictly prefers the optimal simple second-price auction, i.e., fixing $\overline{x} = x_1$ to the second score auction with the scoring function reflecting her true preferences, i.e., to $g(x) \equiv v(x)$.*

If the scoring function represents the procurer's true preferences, each supplier θ bids the efficient design $x^*(\theta)$. In this case, there is no scope for renegotiation. Avoiding renegotiation, however, is not in the procurer's interest. This is due to the fact that we assume efficient – Coasian – bargaining ex post and that the gains from renegotiation are incorporated in the initial price bids.

1.5 Discussion and Extensions

1.5.1 Discussion of Our Main Assumptions

Strikingly as the findings are, they rely on a couple of strong assumptions. First, we assume that there is a design that is relatively cheap to produce for all types and the differences in production costs between types are lowest for this design. If the differences in production costs across types are lowest for a more expensive design, specifying the cheapest design ex ante is no longer optimal. A potential micro foundation of our assumption might be the following: Suppose that designs consist of different components. The more complex a design, the more components it consists of. A supplier's type is now a distribution from which the various cost types for the different components are drawn. If the draws are independent for each component, then suppliers' expected differences in overall costs are increasing in the number of components; i.e., the expected cost differences are larger for more complex projects.

Second, and related to the point above, suppliers' preferences regarding designs are aligned; i.e., a change in the design leads either to a cost increase or decrease for all suppliers. To see the importance of this assumption imagine that – for the sake of the argument – all designs give rise to the same gross benefit. A supplier's type determines which design the supplier can deliver at lowest cost. Now, specifying one design ex ante reduces competition and leads to a high ultimate price. A scoring auction is now likely to be optimal because it enhances competition. Our assumption that costs are perfectly aligned is sufficient but not necessary for our results. Allowing for "weakly" aligned costs makes the analysis more tedious without adding much insight.

Third, the suppliers anticipate the contract renegotiation and incorporate the potential ex post profits in their ex ante bid. This, however, implies that if the initial contract is enforced – because the parties could not find an agreement at the renegotiation stage – the contractor may make a loss and is able to cover this loss. If suppliers are protected by limited liability, the initial price is bounded from below by the production cost. Due to this price bound, the supplier's outside option at the renegotiation stage is now more valuable. Hence, renegotiation becomes more expensive for the procurer and thus she may have an incentive to specify a design which is close to the ex post efficient design.

Fourth, we assume that renegotiation takes place under symmetric information.[21] Suppose the contractor's cost type is private information ex post and that either the procurer or the contractor can make a take-it-or-leave-it offer at the renegotiation stage. Now, the optimal renegotiation offer of the procurer might be such that it is rejected by certain supplier types. If this is the case, then the initial contract is executed. Thus, auctioning off the simplest design which creates only little value to the procurer is likely to be suboptimal if there is a high chance that the initial contract will be executed due to a breakdown of renegotiation. This story is true, however, only if the procurer cannot deduce the contractor's type from the initial bid.

[21] Assuming that suppliers' types are public information already ex ante does not change our main results.

The game is now a signaling game, where the initial bid may contain information about the supplier's type. If there is a fully separating equilibrium, contract renegotiation basically takes place under full information. Thus, behavior at the renegotiation stage is as in our model with complete information. Suppliers' bidding behavior ex ante might nevertheless be different. An analysis of procurement with contract renegotiation under asymmetric information, which involves solving an intricate signaling game, is beyond the scope of this dissertation.

Renegotiation can be imperfect for reasons different from asymmetric information. This is discussed in detail in the following subsection.

1.5.2 Risk of Breakdown of Renegotiation

So far we assumed that ex post the parties engage in Coasian bargaining. In other words, the ex post efficient design is always implemented; i.e., there are neither adjustment costs of contract renegotiation nor a risk that renegotiation fails. In the following, we will show that our main findings are robust toward introducing frictions of contract renegotiation.

We model the imperfection of contract renegotiation in the following simple way. With fixed exogenous probability $b \in [0, 1)$ the parties cannot reach an agreement ex post – i.e., renegotiation fails. In this case the initial contract $(\overline{x}, \overline{p})$ is executed. With the converse probability $1 - b$ the parties reach an agreement and the outcome is determined by the GNBS. The parameter b measures how intricate or how costly contract renegotiation is. For $b = 0$ the model collapses to the one previously analyzed.[22]

The analysis of the model with a risk of renegotiation breakdown proceeds by the same steps as the analysis of Section 1.3.

Proposition 1.4. *Suppose that Assumptions 1.1-1.3 hold and that renegotiation fails with probability $b \in [0, 1)$. The procurer optimally chooses the design $\overline{x}$ that solves*

$$\max_{\overline{x}} \quad \mathbb{E}_{\dot{\theta},\ddot{\theta}}\Big[b[v(\overline{x})) - c(\overline{x},\ddot{\theta})] - (1-\alpha)(1-b)[c(\overline{x},\ddot{\theta}) - c(\overline{x},\dot{\theta})]\Big].$$

The procurer now faces a trade-off. On the one hand, as before, she wants to minimize the cost advantage that the most efficient supplier has in comparison to the second most efficient supplier in the production of design $\overline{x}$. This is achieved by setting $\overline{x} = x_1$. On the other hand, the procurer has an incentive to choose as initial design the design that is optimal when the second most efficient supplier obtains the contract. This is intuitive because if renegotiation fails the procurer obtains the surplus that is generated by the second most efficient type. This is likely to be achieved by a design $\overline{x}$ which is more complex than design x_1. If b is sufficiently low, the former concern dominates the latter and $\overline{x} = x_1$ is optimal. As

[22]Ganuza (2007) uses the same approach to model transaction costs of contract renegotiation. In his interpretation there are transaction costs associated with renegotiation and these costs are stochastic. With probability $1 - b$ the transaction costs are zero, while with probability b the transaction costs are prohibitively high so that renegotiation does not take place.

the risk of renegotiation failure increases, the optimal ex ante design becomes (weakly) more complex.[23]

If $\overline{x} \neq x_1$ is optimal, the ex post renegotiation can lead to an upward as well as downward price adjustment. The latter occurs if the most efficient type happens to be rather inefficient. How likely upward and downward price adjustments are, is intricate to characterize without further assumptions on the type distribution and the feasible designs. Therefore, we will present the results of a simple numerical example in the following. Let $\rho^U := Prob(x^*(\dot{\theta}) > \overline{x})$; i.e., it is the (conditional) probability with which renegotiation leads to an upward price adjustment (a cost overrun occurs).

Example 1.1. *Let the cost function be $c(x, \theta) = (1/2)\theta x^2$ and the value function be $v(x) = x$. The set of feasible designs is assumed to be continuous and given by $X = [1/2, 1]$, with $x_1 = 1/2$. The types of the $n \geq 2$ sellers are drawn independently from the uniform distribution with support $[1, 2]$. The ex post efficient design is $x^*(\theta) = 1/\theta \in [1/2, 1]$.*

The procurer optimally specifies design

$$\overline{x} = \max\left\{\frac{1}{2}, \frac{b(1+n)}{1-\alpha+b(2+\alpha+n)}\right\}$$

in the initial auction. The initial design is equal to the simplest design, $\overline{x} = x_1 = 1/2$ if $b \leq \hat{b}$, where

$$\hat{b} := \frac{1-\alpha}{n-\alpha}.$$

Notice that for $\alpha < 1$ the critical probability of breakdown $\hat{b}$ is strictly greater than zero. In other words, for a range of b-values specifying the simplest design initially is optimal.

A cost overrun – an upward price adjustment – occurs ex post if $x^*(\theta) > \overline{x}$, which is equivalent to $\dot{\theta} < \overline{x}^{-1}$. The probability of a cost overrun is bounded from below,

$$\rho^U \geq 1 - \frac{1}{\mathrm{e}^2} \approx 0.865,$$

where e *is Euler's number.*[24]

The example illustrates that even when there is a risk of renegotiation failure – and thus a rational for the procurer to choose a more complex design than x_1 – ex post adjustments typically lead to the implementation of a more complex and costlier design. Thus, if renegotiation takes place, an upward price adjustment is extremely likely. In the example, if renegotiation takes place the final price exceeds the initial price in more than 86% of the cases. Downward price adjustments are unlikely and occur in less than 14% of the cases. At

[23] We can interpret b as a bargaining inefficiency multiplier. As renegotiation becomes more costly, the optimal ex ante design becomes (weakly) more complex.

[24] Detailed calculations to Example 1 are presented in Appendix A.

first glance, these percentages may seem somewhat extreme. It is useful to remember that the procurer maximizes

$$\mathbb{E}_{\ddot{\theta}}\left[v(\overline{x}) - c(\overline{x}, \ddot{\theta})\right], \tag{1.23}$$

if renegotiation is very costly or her bargaining power is very low. The $\overline{x}$ that maximizes (1.23) is the most complex initial design that can be optimal. The more complex the design is, the less likely is a cost overrun ex post. The term (1.23) is maximized for the design at which the surplus generated by the second most efficient type is maximized in expectations. With the most efficient type being likely to be more efficient than the expected second most efficient type, even in this extreme situation upward adjustments are more likely than downward adjustments.[25] This explains why the lower bound on the conditional likelihood of a cost overrun is so high.

1.6 Conclusion

We analyzed competitive procurement mechanisms in an environment where the procurer is unable to commit not to renegotiate the contract ex post. Moreover, the cost function of the supplier who has been awarded with the initial contract is publicly observed ex post. Hence, if the initial design turns out to be ex post inefficient – for the given cost function – the parties adjust the initial design to the ex post efficient one; i.e., the parties engage in Coasian bargaining. The main finding is that the constrained-optimal award procedure is a simple auction. The procurer awards the contract for the simplest design via a standard auction. Ex post, a more complex and more costly design is implemented and thus the ultimate price typically is higher than the initial price determined by the auction. Interestingly, the optimal simple auction where the procurer collects only price bids for the simplest design outperforms multi-dimensional auctions, where a bid contains a price and a design. The reason is that a multi-dimensional auction reduces competition between the suppliers ex ante.

The finding that it is optimal to auction off the simplest design when using a price-only auction holds true also when the procurer can set a maximum bid. When renegotiation is inefficient or there is a risk it might breakdown, the optimal ex ante design is not necessarily the most simple one any more. However, as we demonstrate in an example, the optimal ex ante design still has a strong tendency to be relatively simple (simpler than the optimal ex post design) so that renegotiation typically leads to an upward price adjustment.

The findings of the paper rely on a couple of strong assumptions that often will not all be satisfied in practice. Hence, we do not argue based on these results that most of the projects with severe cost overruns that we observe in practice are always the result of efficient award procedures. However, our main assumption that commitment not to renegotiate is not feasible seems to be realistic. For instance, complex construction projects often cannot be executed exactly the way as initially specified, so contract renegotiation has to take place.

[25] This intuition holds true only for “nice” type distributions like the uniform distribution.

This chapter shows that severe cost overruns are not necessarily a sign of inefficient award procedures or project completion.

Chapter 2

The Timing of Choice-Enhancing Policies*

2.1 Introduction

Automatic enrollments and renewals are prevalent in many service industries. For example, cell-phone companies offer fixed-term contracts with automatic renewals.[1] Internet-connection providers automatically enroll their customers into anti-virus options with some grace period.[2] Retail banks often promote credit cards with very low interest rates for an initial teaser period, after which the interest rate rises.[3] With mounting evidence that some consumers exhibit systematic behavioral biases, there is concern that automatic enrollments and renewals may be used to exploit unsophisticated consumers. In response, choice-enhancing policies have been called for to protect such consumers.[4] Recent studies such as Carroll, Choi, Laibson, Madrian and Metrick (2009), Keller, Harlam, Loewenstein and Volpp (2011), and Chetty, Friedman, Leth-Petersen, Nielsen and Olsen (2014) have shown that motivating consumers to make an active choice can improve consumer welfare. However, two issues associated with such policies have been under-investigated. First, is there any adverse effect when firms can respond to such a policy? Second, if consumers can opt out of a service at different points in time, when should a policymaker motivate consumers to make an active choice?

---

*This chapter is based on joint work with Takeshi Murooka.

[1] In many countries, cell-phone companies offer a two-year contract with a mobile discount and automatic renewals of the contract. Although the loss these companies incur by offering the mobile discount is typically recouped by the monthly fees, some continue to charge the same monthly fees even after the initial two years.

[2] As a specific example, Kabel Deutschland, one of the largest Internet-connection providers in Germany, provided an optional three months of free antivirus software, a firewall, and parental control software with an automatic enrollment. After the first three months, the cost rose to €3.98 per month. See `http://www.kabel-internet-telefon.de/news/7214-kabel-deutschland-mit-neuem-sicherheitscenter-kabelsicherheit-de` (accessed March 1, 2016).

[3] Ausubel (1991) and DellaVigna and Malmendier (2004) document the evidence.

[4] To protect consumers, many countries have recently started using behavioral economics to improve policies. For example, the UK government created the Behavioral Insights Team (known as the Nudge Unit) and the US government built the Social and behavioral Science Team.

This chapter analyzes the welfare consequences of policies when a firm can change its pricing strategy in response to the policies. Section 2.2 introduces an illustrative model. A firm automatically enrolls consumers into a service. Some consumers are (partially or fully) naive present-biased à la O'Donoghue and Rabin (2001), whereas all others are time-consistent and rational. Each consumer incurs a positive switching cost when she opts out of the service, and she can do so at multiple points in time. By employing a choice-enhancing policy (e.g., by requiring firms to send an email with a simple cancellation format or to prominently inform how consumers can cancel the service), a policymaker can reduce the switching cost in a certain period.

Section 2.3 analyzes the illustrative model and presents our main results. If there is no policy intervention, the firm may exploit naive present-biased consumers by charging a high price for the service after a grace period.[5] We first analyze the effect of a policy that decreases the switching cost when consumers are enrolled. For example, in many countries firms are required to prominently inform consumers about how to cancel their service when they enroll consumers.[6] If the firm's pricing strategy is fixed, then such a conventional policy always (weakly) increases each consumer's utility. By contrast, we show that if the firm can change its pricing strategy in response, then this conventional policy can strictly decrease consumer welfare. Intuitively, because naive consumers may procrastinate their switching decision, time-consistent consumers are more responsive to the policy (i.e., more likely to opt out of the firm's service in response to the policy) than naive present-biased consumers. Then, the policy may increase the proportion of naive consumers among consumers who stay enrolled in the service. In response, the firm increases its prices to exploit naive consumers, thereby reducing naive consumers' long-run utility. This is a perverse result, as such policies typically aim to protect these unsophisticated consumers. In this case social welfare also decreases, because time-consistent consumers switch and thus incur a (socially wasteful) switching cost.

As an alternative policy, we then investigate a policy that decreases the switching cost whenever the firm increases the price for its service (or when the free-trial period ends). As a practical example of such an alternative policy, a firm could be required to prominently inform consumers about how to cancel its service upon an (expected or unexpected) price increase.[7] We show that—in contrast to the above conventional policy—this alternative policy always increases consumer (and social) welfare. Intuitively, because all consumers plan to switch in the same period under the alternative policy, the policy does not change the proportion of naive consumers who stay enrolled in the service. Hence, the firm's trade-off between exploiting naive consumers by setting a high price and serving to all consumers by setting a

[5] Although this chapter focuses on the case in which automatic enrollments or renewals can be used to exploit naive consumers, we note that automatic enrollment or renewal itself can be valuable for consumers in a general framework. We discuss this issue in Sections 2.5 and 4.6.

[6] This makes it easier for consumers to cancel at the time of the enrollment rather than to cancel later, because in the latter case they have to remember or find out how to cancel it.

[7] Note that in most countries, firms are required to announce an unexpected price increase or an unexpected termination of a free trial to their customers. Here, the firm is required to announce and provide a simple cancellation procedure even upon a known price increase or a previously-announced termination of the free-trial period.

moderate price is unaffected. As a result, the alternative policy does not have the perverse effect of inducing the firm to increase its price. We also show that if there is a per-period cost of decreasing the switching costs, the proposed alternative policy is better than decreasing the switching costs in all periods.

Section 2.4 analyzes some extensions of the model. As a primary extension, we endogenize consumers' decisions of signing up to the enrollment and analyze the case in which a firm can charge fees for its service multiple times. In this model, the firm sells a base product necessary to use the add-on service. Consumers who decide to purchase the base product are automatically enrolled into the firm's service. We show that if the number of periods the firm can charge its add-on fee is high, the firm is more likely to exploit naive consumers, resulting in higher total payments by them. Similar to the illustrative model, a conventional policy that decreases the switching when consumers are enrolled can decrease consumer and social welfare, whereas the alternative policy that decreases the switching cost whenever the firm increases its price always (weakly) increases consumer and social welfare. We also examine the robustness of our results to incorporating competition among firms, a fraction of sophisticated present-biased consumers, and heterogeneous product valuations among consumers.

Section 2.5 discusses potential alternative policies: reminders, automatic terminations of a service, regulating prices, and deadlines. We also discuss how each policy can interact with other potential behavioral biases such as forgetting or inattention to a switching opportunity. Section 4.6 concludes. Proofs are provided in Appendix B.

Related Literature This chapter contributes to the literature on behavioral public policy.[8] As the most closely related studies, Carroll et al. (2009), Keller et al. (2011), and Chetty et al. (2014) investigate the policy effects on active choice. These studies focus on cases in which a policymaker either decreases consumers' switching costs to zero or forces consumers to make an explicit choice. In contrast to these studies, we investigate the case in which a policymaker can reduce consumers' switching costs, but the reduced switching cost is still positive and consumers themselves decide whether to switch. We discuss the real-world applications and interpretations of such a policy in Section 2.2.2.

This chapter is also related to two theoretical literatures: pricing for unsophisticated present-biased consumers and the equilibrium effects of policies. First, the literature on behavioral industrial organization has studied how firms can exploit consumers' time inconsistency and naivete.[9] Building upon this stream of the literature, we focus on the policy implications of enhancing active choice and analyze how the timing of policies can affect consumer and social welfare.

[8] O'Donoghue and Rabin (2003, 2006) investigate the welfare effects of tax/subsidy policies under present bias and naivete. Baicker, Mullainathan and Schwartzstein (2015) analyze the design of health insurance under behavioral biases. For surveys of behavioral public policy, see, for examples, Mullainathan, Nöth and Schoar (2012) and Chetty (2015).

[9] See, for example, DellaVigna and Malmendier (2004), Kőszegi (2005), Gottlieb (2008), Heidhues and Kőszegi (2010), and Heidhues and Kőszegi (2015).

Second, recent theoretical and empirical studies have analyzed the equilibrium effects of policies when consumers are inattentive.[10] To the best of our knowledge, however, the timing of employing choice-enhancing policies and the resulting welfare effects have not been investigated in the previous literature. Complementing this literature, we thus highlight the adverse welfare effect of a conventional policy, analyze how the timing of policies affects welfare, and suggest an alternative policy that mitigates the adverse welfare effect and hence can more robustly improve welfare.

2.2 Illustrative Model

This section introduces our illustrative model. Section 2.2.1 sets up the model. Section 2.2.2 discusses key assumptions on procrastination and on consumers' switching costs.

2.2.1 Setup

A risk-neutral firm provides a service to a continuum of risk-neutral consumers (normalized to measure one). There are three periods: $t = 1, 2, 3$. In $t = 1$, the firm automatically enrolls consumers into the service which they value at $a > 0$ in each $t = 2, 3$. The firm offers a free trial: the price for the service $p^a \geq 0$ is charged at the end of the game (i.e., $t = 3$), and p^a is not charged if consumers opt out of the firm's service either in $t = 1$ (i.e., when consumers are enrolled into the service) or in $t = 2$ (i.e., when consumers use the free-trial service). This automatic-enrollment setting also encompasses the case in which a firm automatically renews a consumer's existing contract or provides a consumer's default option with a grace period; we analyze the case in which consumers endogenously make their initial enrollment decisions in Section 2.4.1. A competitive fringe also provides a service with the same value. For simplicity, we assume that the production cost of the service—and hence the price of the competitive fringe—is zero. Each consumer incurs a switching cost k_t by changing from the firm to the competitive fringe in period $t = 1, 2$.[11] At the beginning of the game, the policymaker decides whether to enact a choice-enhancing policy for each period. Without any such policy, $k_t = \overline{k} > 0$ for all t.[12] If a policymaker enacts the policy in period t, then the switching cost of that period is reduced to $k_t = \underline{k} \in (0, \overline{k})$. Denote by $\Delta_k := \underline{k}/\overline{k} \in (0, 1)$.

[10] For the theoretical literature, see Armstrong, Vickers and Zhou (2009), Armstrong and Chen (2009), Piccione and Spiegler (2012), Grubb (2015), de Clippel, Eliaz and Rozen (2014), Ericson (2014), and Spiegler (2015). For the empirical literature, see Duarte and Hastings (2012), Handel (2013), Grubb and Osborne (2015), and Damgaard and Gravert (2016). Relatedly, based on Gabaix and Laibson's (2006) shrouded-attribute model, Kosfeld and Schüwer (2014) analyze the effect of increasing the proportion of sophisticated consumers in the market and show that such an intervention can lower welfare. Intuitively, this intervention can increase the proportion of consumers who (socially inefficiently) substitute away from an add-on consumption. By contrast, we investigate how the timing of enacting a policy affects welfare when some consumers are naive present-biased and derive an alternative policy that robustly improves welfare.

[11] As shown below, consumers have no incentive to switch back from the competitive fringe to the firm.

[12] If k_t is endogenously chosen by the firm, then the firm would set it to the maximal amount. Without loss of generality, we can think of $\overline{k}$ as that amount.

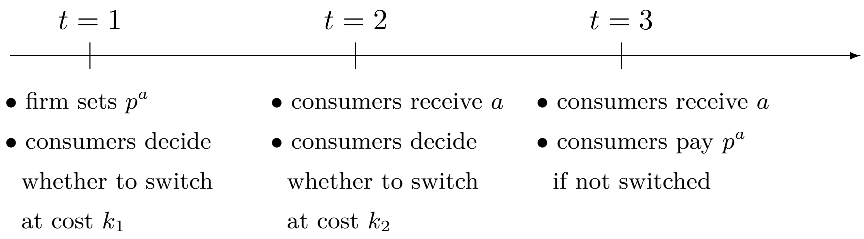

Figure 2.1: Timeline of the model.

Following O'Donoghue and Rabin (1999a, 2001), we assume that a proportion α of consumers are present-biased and (partially or fully) naive, whereas the remaining proportion of consumers are time-consistent and rational. To explain this, suppose that u_t is a consumer's period-t utility. In each period $t = 1, 2$, time-consistent consumers decide whether to opt out of the firm's service based on $u_t + \sum_{s=t+1}^{3} \delta^{s-t} u_s$, and they correctly expect their future behavior. By contrast, present-biased consumers decide whether to opt out based on $u_t + \beta \sum_{s=t+1}^{3} \delta^{s-t} u_s$, where $\beta \in (0, 1)$ represents the degree of their present bias. These present-biased consumers are (partially) naive about their future self-control problem: in $t = 1$, they think that their future present bias will be equal to $\hat{\beta} \in (\beta, 1]$ and that they will behave as if $\beta = \hat{\beta}$ in $t = 2$. When $\hat{\beta} = 1$, these consumers are unaware of their self-control problem; when $\hat{\beta} \in (\beta, 1)$, they are aware of it, but not to the full extent.[13] In what follows, we set $\delta = 1$ without loss of generality.

We investigate perception-perfect equilibria: each player maximizes her perceived utility in each subgame (O'Donoghue and Rabin 2001). We evaluate consumer welfare based on each consumer's long-run utility (i.e., $\sum_{t=1}^{3} u_t$). Figure 2.1 illustrates the timeline of the firm's pricing and consumers' decisions.

2.2.2 Discussion of Key Assumptions

This subsection discusses two key assumptions made in this chapter. First, some consumers may procrastinate their switching decisions.[14] Second, a policymaker can decrease consumers' switching costs by employing a choice-enhancing policy (but the policymaker may be unable to decrease them to zero).

Procrastination Recent empirical and experimental studies have shown that people often procrastinate their decisions.[15] In our model, consumers incur the switching cost now but make the payment later. Because of this discrepancy in the timing, naive present-biased consumers may procrastinate their switching decisions. This assumption is plausible in our

[13] We discuss how our results are robust to incorporating perfectly sophisticated consumers (i.e., $\hat{\beta} = \beta < 1$) into the model in Section 2.4.3.

[14] Following O'Donoghue and Rabin (2001) and DellaVigna (2009), we classify that a consumer "procrastinates" if ex-ante she anticipates switching in some period but does not actually switch in that period.

[15] See, for examples, Ariely and Wertenbroch (2002), DellaVigna and Malmendier (2006), and Skiba and Tobacman (2008).

real-world examples (Internet-connection providers, cell-phone companies, and retail banks), because customers incur an immediate effort cost to switch a service, whereas the change in the bill typically comes later (e.g., at the beginning of the following month). Moreover, the assumption on the timing can be relaxed when consumers possibly incur the payments multiple times, as we extend the model in Section 2.4.1.[16]

Choice-Enhancing Policies The literature on choice-enhancing and active-choice policies has focused on either of the following cases: (i) the policy enables consumers to make a switching decision without incurring any switching cost or (ii) the policy forces all consumers to make a switching decision. In contrast to the literature, this chapter analyzes the case in which a choice-enhancing policy decreases consumers' switching costs (but the switching cost can still be positive) and consumers themselves decide whether to switch. For example, suppose that a firm can automatically enroll its customers into an additional service and that customers need to take an extra action (e.g., register their personal information) if they want to use the additional service of another firm. In this case, a policymaker can decrease the switching cost (e.g., by requiring a simple cancellation format), but cannot decrease it to zero because of the cost of the new registration.

In practice, such automatic enrollments are illegal unless firms offer a grace period for the service. Along with this interpretation, we assume that a firm that automatically enrolls consumers charges no price for the first usage of the service (i.e., $p^a = 0$ in $t = 2$).[17]

Beyond automatic enrollments, our main logic and results are applicable to situations in which consumers find it easier to register with one firm compared with others. As an example, for customers of a retail bank, signing up for a credit card associated with the bank is often easier than doing so at other firms because the bank can use the customer information that it already holds.

2.3 Analysis

Consumer behavior We first characterize each consumer's behavior given prices and switching costs. Note that consumers do not take any action in $t = 3$. Note also that in the context of our model, consumers do not have an incentive to switch back from the competitive fringe to the firm.

[16] Intuitively, even when consumers incur the payment of the service now (rather than later), they may still think that they will opt out later to avoid future payments. Hence, if consumers possibly incur payments multiple times, our results hold qualitatively even when consumers face the switching cost and payment at the same time. By contrast, if consumers incur the payments for the service at most once (as in our illustrative model), our timing assumption is crucial.

[17] In Section 2.4.1, we show that this is without loss in a model with endogenous consumer enrollments because the initial payment upon enrollment is indistinguishable from a payment for the service in $t = 2$. In addition, a firm may have an incentive to offer a low introductory price (Shapiro 1983) or a free trial (Murooka 2016) if the service is an experience good.

If time-consistent consumers are still enrolled in the firm's service at the beginning of period 2, they switch to the competitive fringe if and only if $-k_2 + a > a - p^a$ or equivalently $p^a > k_2$. That is, time-consistent consumers switch in $t = 2$ if and only if the price exceeds the switching cost as in the classical switching-cost literature (Farrell and Klemperer 2007). Given this, period 1's switching behavior can be divided into the following two cases. First, if $p^a \leq k_2$, time-consistent consumers switch in period 1 if and only if $a - k_1 + a > a + (a - p^a)$ or equivalently $p^a > k_1$. Second, if $p^a > k_2$, time-consistent consumers switch in period 1 if and only if $a - k_1 + a > a + (a - k_2)$ or equivalently $k_2 > k_1$. Intuitively, when the switching cost in the second period is sufficiently high, time-consistent consumers switch in period 1 if the price exceeds the switching cost in the first period; otherwise, they switch in period 1 if the switching cost in that period is lower than that in period 2.

If present-biased consumers are still enrolled in the firm's service at the beginning of period 2, they switch to the competitive fringe if and only if $-k_2 + \beta a > \beta(a - p^a)$ or equivalently $p^a > \frac{k_2}{\beta}$. In contrast to time-consistent consumers, present-biased consumers may not switch even if the price exceeds the switching cost. Note also that $\hat{\beta}$ does not affect consumer behavior in period 2. Because these consumers underestimate their future self-control problem, in period 1 they think they will switch in the next period if and only if $\hat{\beta}p^a > k_2$. Given this belief, period 1's switching behavior can be divided into the following two cases. First, if $\hat{\beta}p^a \leq k_2$, naive present-biased consumers think they will keep using the firm's service in period 2. Hence, they switch in period 1 if and only if $a - k_1 + \beta a > a + \beta(a - p^a)$ or equivalently $p^a > \frac{k_1}{\beta}$. Second, if $\hat{\beta}p^a > k_2$, these consumers think they will switch in period 2. Hence, they switch in period 1 if and only if $a - k_1 + \beta a > a + \beta(a - k_2)$ or equivalently $\beta k_2 > k_1$. Intuitively, naive present-biased consumers are less likely to switch in period 2 because they procrastinate switching in period 1 as they underestimate their future impatience and because they are more impatient than time-consistent consumers. It is worth emphasizing that if $p^a = \frac{k_2}{\beta}$, then $\hat{\beta}p^a > k_2$ holds for *any* $\hat{\beta} > \beta$.

Firm behavior and Policy Effects We now analyze the optimal pricing of the firm and the effects of choice-enhancing policies. We first investigate the situation in which the policymaker does not employ any policy, i.e., $k_1 = k_2 = \overline{k}$. The firm faces a trade-off between exploiting naive consumers at a high price ($p^a = \frac{1}{\beta}\overline{k}$) and selling its service to all consumers at a moderate price ($p^a = \overline{k}$). The result is summarized as follows:

Lemma 2.1. *Suppose $k_1 = k_2 = \overline{k}$.*

If $\alpha > \beta$, the firm sets $p^a = \frac{1}{\beta}\overline{k}$. Time-consistent consumers do not pay p^a, whereas naive consumers pay p^a. The profits of the firm are $\pi = \frac{\alpha}{\beta}\overline{k}$.

If $\alpha \leq \beta$, the firm sets $p^a = \overline{k}$. All consumers pay p^a. The profits of the firm are $\pi = \overline{k}$.

The intuition is simple: the firm is more likely to exploit naive consumers if there are more naive consumers (larger α) or if naive consumers suffer from a more severe present bias (smaller β). Because naive consumers pay a high price, which they initially did not anticipate paying, consumer welfare is lower when the firm sells only to naive consumers than when it sells to both naive and time-consistent consumers. Social welfare is also lower because time-consistent consumers pay the switching cost.

The result in Lemma 2.1 does not depend on the extent to which naive consumers are aware of their present bias. To see the intuition, suppose that the firm sets $p^a = \frac{k_2}{\beta}$. Note first that the consumer behavior in $t = 2$ does not depend on $\hat{\beta}$. In $t = 1$, partially naive consumers think that they will switch in $t = 2$ if and only if $p^a > \frac{k_2}{\hat{\beta}}$. Since $\frac{k_2}{\beta} > \frac{k_2}{\hat{\beta}}$ for any $\hat{\beta} > \beta$, these naive consumers do not switch in $t = 1$ if and only if $k_1 \geq \beta k_2$. Consequently, consumer behavior in both $t = 1$ and $t = 2$ does not depend on $\hat{\beta}$ when the firm sets $p^a = \frac{k_2}{\beta}$. Hence, akin to Heidhues and Kőszegi (2010), the firm can make partially naive consumers procrastinate and can exploit them at the same amount irrespective of $\hat{\beta}$.[18] This intuition is also applied to the following results.

We next investigate the situation in which the switching cost is decreased in the first period, i.e., $k_1 = \underline{k}$, $k_2 = \overline{k}$. This is the case if the policymaker employs a policy that reduces the switching cost when consumers are enrolled. The firm still faces the same type of trade-off as above. The equilibrium cut-off condition becomes different, however. On the one hand, time-consistent consumers switch in period 1 if $p^a > \underline{k}$. On the other hand, naive consumers in period 1 prefer to switch in period 2 rather than immediately if $-\underline{k} \leq -\beta\overline{k}$ or equivalently $\Delta_k \geq \beta$. In this case, the firm can set $p^a = \frac{1}{\beta}\overline{k}$ and naive consumers end up paying the price. The result is summarized as follows:

Lemma 2.2. *Suppose* $k_1 = \underline{k}, k_2 = \overline{k}$.

(i) Suppose $\Delta_k \geq \beta$. *If* $\alpha > \beta\Delta_k$, *the firm sets* $p^a = \frac{1}{\beta}\overline{k}$. *Time-consistent consumers switch in period* 1 *and do not pay* p^a, *whereas naive consumers pay* p^a. *The profits of the firm are* $\pi = \frac{\alpha}{\beta}\overline{k}$.
If $\alpha \leq \beta\Delta_k$, *the firm sets* $p^a = \underline{k}$. *All consumers pay* p^a. *The profits of the firm are* $\pi = \underline{k}$.

(ii) Suppose $\Delta_k < \beta$. *If* $\alpha > \beta$, *the firm sets* $p^a = \frac{1}{\beta}\underline{k}$. *Time-consistent consumers switch in period* 1 *and do not pay* p^a, *whereas naive consumers pay* p^a. *The profits of the firm are* $\pi = \frac{\alpha}{\beta}\underline{k}$.
If $\alpha \leq \beta$, *the firm sets* $p^a = \underline{k}$. *All consumers pay* p^a. *The profits of the firm are* $\pi = \underline{k}$.

Lemma 2.2 (i) means that the firm may still be able to charge a high price and exploit naive consumers even if the switching cost in $t = 1$ is decreased. Intuitively, naive consumers procrastinate switching if the decrease in the switching cost in period 1 is not large ($\Delta_k \geq \beta$): in period 1, they do not switch because they prefer to switch in period 2 and (wrongly) think that they will do so. In period 2, however, naive consumers actually do not switch if $p^a \leq \frac{1}{\beta}\overline{k}$. By contrast, Lemma 2.2 (ii) shows that when $\Delta_k < \beta$, the policy in $t = 1$ can decrease the

[18] Specifically, Heidhues and Kőszegi (2010) show that in a general contracting setting, an ex-ante incentive compatibility constraint (in our model, the condition that partially naive consumers procrastinate switching in $t = 1$) does not bind for any $\hat{\beta} > \beta$ if an ex-post incentive compatibility constraint (in our model, the condition that partially naive consumers do not switch in $t = 2$) binds.

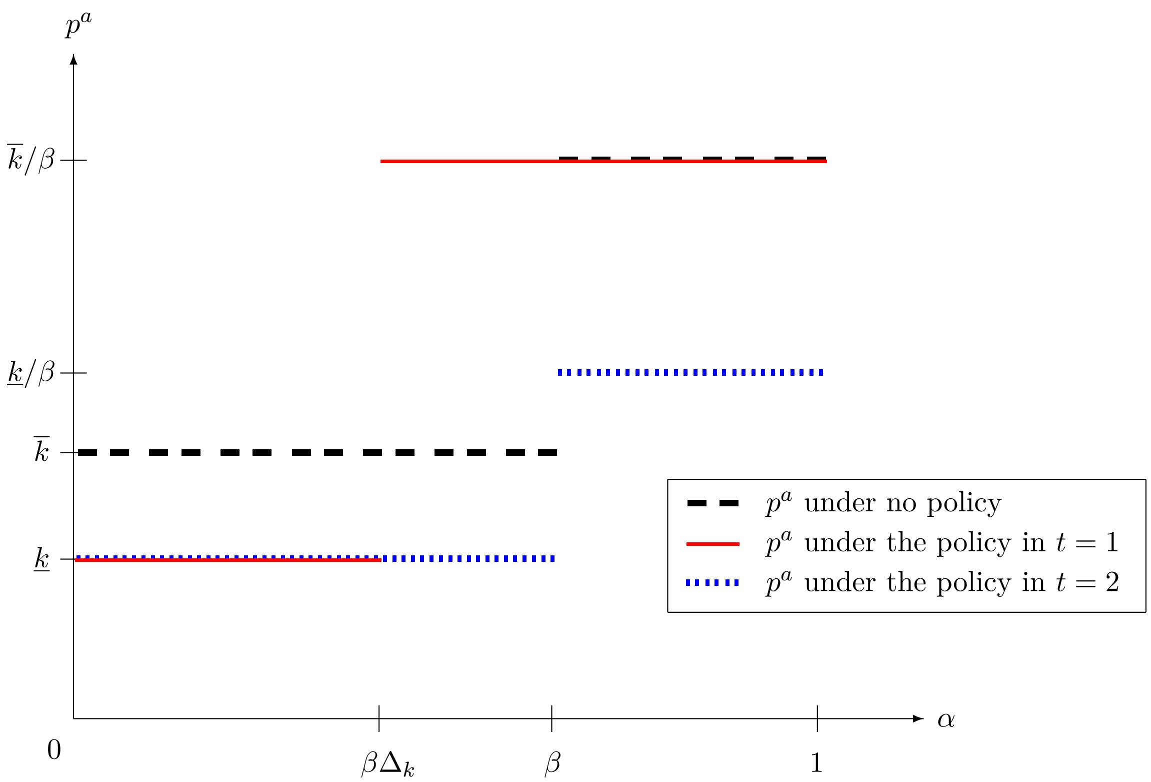

Figure 2.2: Equilibrium price under different policies when $\Delta_k \geq \beta$.

price and increase consumer welfare. Intuitively, if naive consumers do not procrastinate switching, the firm needs to decrease its price in response to the policy.

Comparing Lemma 2.1 with Lemma 2.2 leads to our first main result:

Proposition 2.1. *Suppose the choice-enhancing policy is employed in $t = 1$. If $1 \geq \frac{\alpha}{\beta} > \Delta_k \geq \beta$, the policy increases the equilibrium price and decreases naive consumers' welfare and social welfare. If in addition $\frac{\alpha}{\beta} + (1 - \alpha)\Delta_k > 1$, it also decreases consumer welfare.*

Proposition 2.1 highlights that the policy that decreases the switching cost when consumers are enrolled can lower naive consumers' welfare and social welfare. Intuitively, the firm faces a trade-off between exploiting naive consumers at a high price and selling to all consumers at a moderate price. Because the policy reduces the maximum price time-consistent consumers are willing to pay, it makes exploiting naive consumers relatively more attractive for the firm. When the firm changes its pricing strategy in response to the policy, time-consistent consumers pay the switching cost, which decreases social welfare. Precisely, this perverse result occurs when naive consumers procrastinate switching ($\Delta_k \geq \beta$), the firm sells to both types of consumers without the policy ($\beta \geq \alpha$), and it sells only to naive consumers and exploits them with the policy ($\alpha > \beta\Delta_k$). Under these parameters, the policy also lowers naive consumers' long-run utility and can lower consumer welfare when the proportion of naive consumers is sufficiently large.

Figure 2.2 shows how the policy decreasing the switching cost in $t = 1$ changes the firm's equilibrium pricing when the reduction in the switching cost is insufficient (i.e., when

$\Delta_k \geq \beta$). There are three cases depending on the proportion of naive consumers. When most consumers are time-consistent (i.e., $\alpha \leq \beta\Delta_k$), the firm always chooses a price at which no consumers pay a switching cost, and hence the policy decreases the equilibrium price. When most consumers are naive (i.e., $\alpha > \beta$), the firm sets a high price both before and after the policy and time-consistent consumers pay a switching cost. When the composition of consumers is in-between, however, the firm increases its price in response to the policy—adversely affecting welfare—as stated in Proposition 2.1.

As an alternative policy, we investigate the situation in which the switching cost is decreased in the second period, i.e., $k_1 = \overline{k}$, $k_2 = \underline{k}$. This is the case if the policymaker decreases the switching cost whenever the firm starts charging a (higher) price. Similar to the above analysis, the firm faces a trade-off between exploiting naive consumers at a high price ($p^a = \frac{1}{\beta}\underline{k}$) and selling its service to all consumers at a moderate price ($p^a = \underline{k}$). The result is summarized as follows:

Lemma 2.3. *Suppose* $k_1 = \overline{k}, k_2 = \underline{k}$.

If $\alpha > \beta$, *the firm sets* $p^a = \frac{1}{\beta}\underline{k}$. *Time-consistent consumers switch in period* 2 *and do not pay* p^a, *whereas naive consumers pay* p^a. *The profits of the firm are* $\pi = \frac{\alpha}{\beta}\underline{k}$.

If $\alpha \leq \beta$, *the firm sets* $p^a = \underline{k}$. *All consumers pay* p^a. *The profits of the firm are* $\pi = \underline{k}$.

The parameters under which the firm chooses to exploit the naive consumers (i.e., $\alpha > \beta$) are the same as the ones under no policy. By comparing Lemma 2.1 with Lemma 2.3, we have the following result:

Proposition 2.2. *Suppose the choice-enhancing policy is employed in* $t = 2$. *This policy always strictly increases consumer welfare and weakly increases social welfare. It strictly increases social welfare if* $\alpha > \beta$.

Proposition 2.2 implies that the policy that decreases the switching cost when a firm increases the price (in this case, when a firm starts charging a positive fee) does not have the perverse effect as described in Proposition 2.1. As depicted in Figure 2.2, such a policy always decreases the equilibrium price. Thus, it always strictly increases consumer welfare relative to the no-policy case. It also increases social welfare when time-consistent consumers pay a switching cost in equilibrium, because the policy directly reduces the switching cost.

Furthermore, the comparison between Lemma 2.2 and Lemma 2.3 leads to the following result:

Proposition 2.3. *Under any parameters, both consumer and social welfare are weakly higher when enacting the choice-enhancing policy in* $t = 2$ *than when enacting it in* $t = 1$. *Consumer welfare is strictly higher if* $\frac{\alpha}{\beta} > \Delta_k \geq \beta$, *and social welfare is strictly higher if in addition* $1 \geq \frac{\alpha}{\beta}$.

Proposition 2.3 highlights that the timing of enacting the policy matters for both consumer and social welfare. If a policymaker enacts a choice-enhancing policy when consumers are enrolled, then a firm may change its pricing strategy in response to the policy, and hence the perverse welfare effect can occur. By contrast, as depicted in Figure 2.2, the alternative policy does not have such an adverse effect, and hence is welfare enhancing.

It is worth emphasizing that the choice-enhancing policy in $t = 2$ increases consumer and social welfare robustly to the proportion of naive consumers α, whereas that in $t = 1$ does not. In this sense, our proposed policy is in line with "asymmetric paternalism," which benefits consumers who make errors, while it imposes no (or relatively little) harm on consumers who are fully rational (Camerer, Issacharoff, Loewenstein, O'Donoghue and Rabin 2003).

Optimal Choice-Enhancing Policy We have thus far compared policies reducing the switching cost in $t = 1$ and $t = 2$. Another natural candidate policy is to reduce the switching cost in both periods. Corollary 2.1 summarizes the comparison between the policy employed in $t = 2$ and that employed in both periods:

Corollary 2.1. *(i) If there is no cost of decreasing the switching costs, then social welfare under a policy that reduces the switching cost only in $t = 2$ is equal to social welfare under a policy that reduces the switching cost in both $t = 1$ and $t = 2$.*

(ii) If there is a per-period cost of decreasing the switching costs, then social welfare under a policy that reduces the switching cost only in $t = 2$ is higher than social welfare under a policy that reduces the switching cost in both $t = 1$ and $t = 2$.

In summary, if there is no cost of implementing such a policy (as we assumed), a policy that reduces the switching cost only in $t = 2$ has the same effect as a policy that reduces the switching cost in all periods. If there is a positive per-period implementation cost, however, reducing the switching cost only in $t = 2$ is uniquely optimal unless the implementation cost is very high. If the implementation cost is very high, then implementing no policy is optimal.

2.4 Extensions

This section investigates extensions and modifications of our illustrative model. Section 2.4.1 endogenizes consumers' enrollment decisions and incorporates the possibility of multiple payments by assuming that the firm offers a base product and can enroll consumers in an add-on subscription when consumers buy the base product. Section 2.4.2 incorporates base-product competition among firms. Section 2.4.3 discusses models incorporating sophisticated present-biased consumers and heterogeneous (base-product or add-on) values among consumers.

2.4.1 Endogenous Enrollment and Multiple Payments

We have thus far assumed that consumers are enrolled into the firm's service and that the firm only decides how to price the service. We now modify the model such that consumers endogenously decide whether they take up a base product that automatically enrolls them into an add-on service, while the firm decides how to price both the base product and the add-on service. We also extend the model such that consumers may use the add-on service in multiple periods—which is plausible in the context of add-on subscriptions—and this allows us to derive additional comparative statics.

Suppose $T+1$ periods: $t = 0, 1, 2, \cdots, T$ where $T \geq 3$. The firm produces two types of products: a base product and an add-on. Consumers value the base product at $v > 0$ and can consume it only once in $t = 1$. The firm automatically enrolls consumers who buy the base product into its add-on service. Consumers value the add-on at $a > 0$ in each $t = 2, \cdots, T$, where they can use the add-on only combined with the base product. If consumers do not buy the base product, they receive an outside option with utility $\bar{u} \in [0, v)$ in period T. The production cost of the base product is $c^v \in (0, v - \bar{u})$. Both the firm and a competitive fringe can produce the same add-on at zero cost.

The timing of the game is modified as follows. In period 0, both the firm and the policymaker decide and commit whether to enact the policy (i.e., decreasing consumers' switching cost) for each period. If either or both of them enact the policy for period t, the switching cost in period t is $k_t = \underline{k}$; otherwise, it is $k_t = \overline{k}$. Then, the firm sets and commits to its prices: a price for the base product $p^v \geq 0$, which is charged in period 1 and prices for the add-on $p^a_t \geq 0$ which are charged in $t = 3, \cdots, T$.[19] After observing the prices and switching costs, consumers decide whether to buy the base product at the end of period 0. In period 1, consumers who bought the base product receive v and pay p^v. They also decide whether to opt out of the firm's add-on (and buy the add-on from the competitive fringe) at switching cost k_1. Then, in each period $t = 2, \cdots, T$, consumers who use the add-on receive a. In addition, in period $t = 2, \cdots, T-1$, if consumers have not opted out of the firm's add-on, they decide either to opt out at the switching cost k_t incurred in period t or to pay p^a_{t+1} in period $t+1$. The game ends at the end of period T.

As in Proposition 2.1, the policy that decreases the switching cost when consumers are enrolled into the add-on service can lower social welfare. In contrast to Proposition 2.1, however, the firm needs to appropriately discount its base-product price to attract consumers. As a result, the policy also decreases consumer welfare whenever it decreases social welfare:

Proposition 2.4. *Suppose that a policymaker enacts a choice-enhancing policy in $t = 1$. If $1 \geq \frac{\alpha + (T-3)(1-\beta)\alpha}{\beta} > \Delta_k \geq \beta$, the policy increases the equilibrium add-on prices and decreases both consumer and social welfare.*

[19] As discussed in Section 2.2, $t = 2$ is a free-trial period and hence $p^a_2 = 0$.

We next investigate an alternative policy. Interestingly, merely imposing a low switching cost in period 2 is insufficient to unambiguously improve welfare, because the firm could react to the policy by setting a low p_3^a, leading naive consumers not to switch in that period, and exploiting consumers afterwards.[20] Hence, we propose a policy in which the policymaker forces the firm to lower the switching cost whenever it increases the add-on price. As an example, suppose that the firm sets $p_3^a > 0$. Since the add-on price increases from $p_2^a = 0$ to $p_3^a > 0$, the policy requires lowering consumers' switching costs in $t = 2$. In practice, a policymaker could force firms to send an email with a simple cancellation format to consumers upon a price increase, even when the price increase was known and previously announced to consumers.

Under such a policy, the firm may have an incentive to voluntarily decrease its switching cost to $\underline{k}$ in all periods after it is forced to do so. Intuitively, voluntarily lowering switching costs makes naive consumers more likely to believe that they will switch in the future and hence it makes them more likely to procrastinate their switching decision.[21] The welfare effects of the alternative policy are summarized as follows:

Proposition 2.5. *Suppose the policymaker enacts a policy that requires the firm to lower its switching cost whenever it increases the add-on price. The policy always weakly increases consumer and social welfare. It strictly increases consumer and social welfare if* $\alpha + (T - 3)(1 - \beta)\alpha > \beta$.

Note that the policy in Proposition 2.5, which requires the firm to lower the switching cost whenever the firm raises add-on prices, is more likely to increase consumer and social welfare as T rises.

Furthermore, the following result demonstrates that it is not necessary to force the firm to reduce the switching cost in every period to improve welfare; a milder intervention as in Proposition 2.5 has the same consequence.

Proposition 2.6. *The equilibrium outcomes under a policy that forces the firm to reduce the switching cost whenever the firm increases the add-on price are the same as those under a policy that forces the firm to reduce the switching cost in every period.*

Hence, akin to the results in Corollary 2.1, our suggested policy may be preferable when there is a positive cost for forcing firms to reduce a switching cost.

[20] Formally, consider the case in which naive consumers face a switching decision in period 2, $k_2 = \underline{k}$, and $k_t = \overline{k}$ for all $t \geq 3$. In this case, naive consumers (wrongly) think that they will switch in period 3 if $p_4^a = \frac{1-\beta}{\beta}\overline{k}$. Given that, they do not switch in period 2 if $-\underline{k} \leq -\beta(p_3^a + \overline{k})$. Hence, if $\Delta_k \geq \beta$, the firm can make naive consumers procrastinate their switching decisions by lowering its add-on price.

[21] To see this, suppose that naive consumers face a switching decision in period t with $k_t = \underline{k}$ because of the increase in the add-on price and the policy. In period t, the condition for naive consumers to procrastinate switching to the next period is $-\underline{k} \leq -\beta(p_{t+1}^a + k_{t+1})$. Note that naive consumers always switch in period t if $\Delta_k < \beta$ and $k_{t+1} = \overline{k}$. Hence, if $\Delta_k < \beta$, the firm decreases k_{t+1} from $\overline{k}$ to $\underline{k}$ voluntarily in order to lead naive consumers to procrastinate their switching decisions. Interestingly, even if $\Delta_k \geq \beta$, the firm has an incentive to decrease k_{t+1} under that policy; see the proof of Lemma B.3 for details.

2.4.2 Competition on the Base Product

In Section 2.4.1, we assumed that only one firm can provide the base product. In this subsection, we analyze the case in which $N \geq 2$ firms sell a homogeneous base product. We investigate a symmetric pure-strategy equilibrium in which all firms offer the same contract in $t = 0$ and equally split each type of consumers in the case of tie-breaking. For simplicity, we focus on the case in which $T = 3$.

Under competition on the base product, market outcomes depend on whether setting negative prices is feasible or not. We first discuss the case in which firms can set any base-product prices:

Proposition 2.7. *Suppose there are $N \geq 2$ firms selling the base product and $p^v \in \mathbb{R}$.*

Then, all firms earn zero profits in any equilibrium. Under any parameters, both consumer and social welfare are weakly higher when enacting the choice-enhancing policy in $t = 2$ than when enacting it in $t = 1$. Consumer welfare is strictly higher if $\frac{\alpha}{\beta} > \Delta_k \geq \beta$, and social welfare is strictly higher if in addition $1 \geq \frac{\alpha}{\beta}$.

Intuitively, if firms can compete down their base-product prices, they will do so as in standard Bertrand-type price competition. Although all profits from exploitation are passed on to consumers and all firms earn zero profits, the timing of the policies still matters. In addition, a cross-subsidization from naive consumers to time-consistent consumers may occur under competition, because the presence of naive consumers decreases the equilibrium base-product price (Gabaix and Laibson 2006).

In practice, however, firms may be unable to profitably set overly low prices. Heidhues, Kőszegi, and Murooka (2016b, 2016c) investigate how the possibility of arbitrage can endogenously generate a price floor of the base product.[22] To investigate such a case in a simple manner, suppose that a base-product price is restricted to $p^v \geq 0$. In this case, firms may earn positive profits even under competition:

Proposition 2.8. *Suppose there are $N \geq 2$ firms selling the base product and $p^v \geq 0$.*

When the choice-enhancing policy is enacted in $t = 2$ or when it is enacted in $t = 1$ and $\Delta_k < \beta$, a positive-profit equilibrium in which $(p^v = 0, p_3^a = \frac{1}{\beta}\underline{k})$ exists if $\frac{1}{N}(\frac{\alpha}{\beta}\underline{k} - c^v) > \max\{\underline{k} - c^v, 0\}$. When the choice-enhancing policy is enacted in $t = 1$ and $\Delta_k \geq \beta$, a positive-profit equilibrium in which $(p^v = 0, p_3^a = \frac{1}{\beta}\overline{k})$ exists if $\frac{1}{N}(\frac{\alpha}{\beta}\overline{k} - c^v) > \max\{\underline{k} - c^v, 0\}$.

Under any parameters, both consumer and social welfare are weakly higher when enacting the choice-enhancing policy in $t = 2$ than when enacting it in $t = 1$. Consumer and social welfare are strictly higher if $\Delta_k \geq \beta$ and $\frac{1}{N}(\frac{\alpha}{\beta}\overline{k} - c^v) > \max\{\underline{k} - c^v, 0\} \geq \frac{1}{N}(\frac{\alpha}{\beta}\underline{k} - c^v)$.

[22] See also Armstrong and Vickers (2012) and Grubb (2015) for an analysis with price floors for the base product.

To see the intuition, suppose that $k_1 = k_2 = k > \frac{\beta}{\alpha}c^v$. Then, each firm earns profits $\frac{1}{N}(\frac{\alpha}{\beta}k - c^v) > 0$ by setting $(p^v = 0, p_3^a = \frac{k}{\beta})$ and charging the add-on price only to naive consumers. While no firm can decrease p^v, each firm can still attract consumers by lowering p^a. From the deviation, however, each firm would earn profits of at most $k - c^v$. Hence, such deviations may not be profitable for firms.[23] The effects on policies are qualitatively the same as those in Section 2.4.1.

2.4.3 Further Extensions

Incorporating Sophisticated Consumers We have thus far assumed that all present-biased consumers are (either partially or fully) unaware of their self-control problem. When sophisticated consumers ($\beta = \hat{\beta} < 1$) are also in the market, the condition in which the firm chooses to exploit naive consumers changes.[24] Intuitively, sophisticated consumers might make no purchase if they are afraid of being exploited because of a high add-on price. Our policy implications are, however, still valid in that the policy in $t = 1$ can be worse than no policy whereas the policy in $t = 2$ is better than no policy. The analysis is provided in the Supplementary Material.

Heterogeneous Demand We now discuss the cases in which consumers' valuation of the base product or of the add-on is heterogeneous. In this case, the equilibrium base-product price may differ. Similar to Grubb (2015), under downward-sloping demand, a choice-enhancing or an active-choice policy may increase the equilibrium base-product price. The intuition is as follows. As in a simple monopoly problem, a firm faces a trade-off between charging a high price for the base product (but only serving few consumers) and serving many consumers by setting a low price for the base product (but only making a small profit per consumer). In addition to the profits from the base product, the firm makes extra profits from the add-on. If a policy reduces the profits from the add-on, serving many consumers becomes less profitable for the firm. Hence, the policy may increase the base-product price, which can be detrimental to consumer welfare. This effect would not arise under competition on the base product, however. The analysis is provided in the Supplementary Material.

2.5 Discussion

In this section, we discuss the effects and potential limitations of other policies—reminders, automatic terminations of subscriptions, regulating prices, and deadlines—in turn. We also discuss how our suggested policy (and other policies) can interact with other behavioral biases, such as forgetting or inattention to a switching opportunity.

[23] Here, the logic of the existence of the positive-profit equilibrium is close to that of Heidhues, Kőszegi and Murooka (2016c), although our model is dynamic and firms do not have an option to educate naive consumers.

[24] Nocke and Peitz (2003) analyze the durable-good market in the presence of sophisticated present-biased consumers.

Reminder A policymaker could send a reminder or provide more information for consumers. Note that if consumers have a self-control problem, merely providing additional information does not prevent their procrastination. Further, even if consumers also have other biases such as forgetting or inattention to their switching opportunities, our suggested policy—sending an email with a simple cancellation format upon increasing a price—would also work as a reminder. In this sense, our suggested policy is robust to such other behavioral biases.

Automatic Termination A policymaker may be able to impose an automatic termination of subscriptions of a service after a free-trial period (or, equivalently, employ an opt-in policy as a default). Although such an automatic termination policy would work in our basic model, consumers who want to keep using the current firm's service would have to sign up again, and hence the policy may generate unnecessary re-registration costs. In addition, if a service is automatically terminated, then present-biased consumers (or consumers who may forget to re-subscribe to the service) may fail to sign up again, which harms consumer and social welfare. By contrast, a policy decreasing consumers' switching costs does not have these drawbacks.

Price Regulation A policymaker could directly regulate the price for a service. Note that in our illustrative model, simply imposing $p_t^a = 0$ for all t maximizes social welfare, and setting any price ceiling below $\frac{\overline{k}}{\beta}$ would (weakly) increase consumer welfare. In practice, however, it is often hard for the policymaker to know the firm's cost function for the service. If the policymaker inaccurately estimates the cost function, then direct price regulation can decrease welfare. In summary, although imposing mild price regulation may prevent firms' exploitation and hence increase welfare, imposing a stringent price regulation would be difficult and questionable.

Deadline A policymaker may be able to impose a strict deadline for consumers' switching decisions. Indeed, if consumers may incur multiple add-on payments as in Section 2.4.1, then imposing such a deadline increases welfare. This finding is in line with the theoretical literature that analyzes the effects of imposing deadlines (O'Donoghue and Rabin 1999b, Herweg and Müller 2011). Unlike a policy decreasing switching costs, however, one should be cautious about imposing such a deadline in practice. Imposing a deadline may be harmful if the add-on values or switching costs change over time. In addition, if consumers are inattentive to or forget their switching opportunities, then imposing such a deadline would decrease consumer and social welfare. Furthermore, imposing a deadline might be infeasible if the firm can circumvent the deadline by (pretendedly) changing the product features of the add-on such that consumers receive extraordinary termination rights. The analysis is provided in the Supplementary Material.

2.6 Concluding Remarks

We investigate the welfare consequences of policies that reduce consumers' switching costs when a firm can change its strategy in response to a policy. We show that a conventional policy—reducing the switching cost when consumers are enrolled into a service—can decrease consumer and social welfare. We also show that an alternative policy—reducing the switching cost when a firm charges a higher price for the service—always (weakly) increases welfare compared with no policy or a conventional policy. Our welfare and policy implications shed light on the design of choice-enhancing and active-choice policies. The logic of our model and its policy implications seem applicable when rational consumers are more responsive to a change in the economic environment than consumers who have behavioral biases.

We conclude by discussing two important issues related to (but beyond the scope of) this chapter. First, how to detect consumer naivete and an adverse policy effect from market data is both theoretically and practically important. One difficulty—as briefly discussed in Section 2.5—is that an automatic enrollment itself may not harm consumer and social welfare. For example, naive consumers may procrastinate taking up a valuable additional service if there are costs for registration and no automatic enrollment. In such a case, the automatic enrollment itself is valuable, although it may allow the firm to exploit consumers as analyzed in this chapter. As a potential future direction, investigating usage as well as purchase data could help identify consumer naivete and exploitation.

Second, this chapter focuses on the present bias as a source of procrastination. Although present bias is one of the most prevalent behavioral biases and our policy implications seem applicable whenever rational consumers are more responsive to a policy than naive consumers, empirically identifying the type of consumer bias is an important issue. Further, designing an optimal policy depends on the types of consumer biases in general. Identifying the type of consumer biases from market data and investigating an optimal policy in a model with multiple sources of consumer biases are left for future research.

Chapter 3

Long-Term Employment Relations When Agents Are Present Biased[*]

3.1 Introduction

Numerous studies have documented that a substantial fraction of the population suffers from a present bias. This present bias affects decision making in diverse domains like retirement savings behavior, see Laibson (1997), health club attendance, see DellaVigna and Malmendier (2006), or credit card usage, see Shui and Ausubel (2005). All these settings have in common that immediate costs have to be traded off with delayed benefits and that their present bias leads people to make too "shortsighted" decisions. A domain that, so far, has not been studied with this focus but that also shares this feature – immediate costs, delayed rewards – is individuals' career decisions in the labor market.

Deciding which career to take and where to work is one of the most important decisions people make. It is an inherently long-term decision, involving inter–temporal trade-offs. Survey evidence suggests that most young employees want to make a "career", i.e., they strongly care about opportunities for career development that are offered by prospective employers[1] and the opportunities for career development are a top priority in job choice[2]. Hence, many firms prominently advertise their development programs[3] and attract and motivate employees with career prospects and development plans.

[*]This chapter is based on joint work with Florian Englmaier and Matthias Fahn.

[1]See, e.g., slide 28 of the "E&Y Studentenstudie 2014", `http://www.ey.com/Publication/vwLUAssets/EY_-_Acht_von_zehn_Studenten_sind_zufrieden/$FILE/EY-studentenstudie-2014-werte-ziele-perspektiven-prÃďsentation.pdf`, last accessed, May 29, 2015.

[2]See, e.g., slide 23 of the "E&Y Absolventenbefragung 2013", `http://www.ey.com/Publication/vwLUAssets/EY-Absolventenbefragung_2013/$FILE/EY-Absolventenbefragung-2013-Studie.pdf`, last accessed, May 29, 2015.

[3]To be a "Great Place to work" the Harvard Business Review asks firms to "Provide employees with ongoing opportunities and incentives to learn, develop and grow [...]"; see `https://hbr.org/2011/09/the-twelve-attributes-of-a-tru.html`, last accessed January 28, 2015.

However, apparently many employees fail to take up these options and actually do not "make a career". This has been documented by several studies in organizational psychology; see, e.g., Garavan and Morley (1997). In particular, Rindfuss, Cooksey and Sutterlin (1999) state that "no matter when expectations are measured, even as late as age 25, fewer than half of the young men and women actually achieve their occupational expectations." Moreover, this "failure" is not (only) due to firms defaulting on promised development options, but an explanation resides with a core feature of career choices, their cost/reward structure with immediate costs, but delayed benefits. Such a structure opens up room for time-inconsistency phenomena. It is well-known and documented that a substantial fraction of the population exhibits a present bias and consequently systematically fails to make optimal intertemporal choices.[4] Here we study how the design of optimal long-term employment contracts is affected by such a present bias.

We analyze how agents' present bias affects contracting in an infinite-horizon employment setting, abstracting from other agency problems such as moral hazard, adverse selection, or limited commitment.[5] Naive agents are offered a menu of contracts, consisting of a *virtual* contract – which they plan to choose in the future – and a *real* contract which they end up choosing. This virtual contract allows the principal to exploit the agent, because a major part of the agent's compensation is shifted from the real into the virtual contract. If the agent is protected by limited-liability, implemented effort can be inefficiently high from a social-planner-perspective. Moreover, it turns out that changes that in general should improve the agent's situation, like employment protection legislation, in fact hurt him here. Finally, considering the finite-horizon setting reveals that the degree of exploitation of naive agents decreases over the life-cycle, i.e., older (naive) workers are ceteris paribus better off.

In the main part of the analysis, we assume a setting where a risk neutral principal and a risk neutral agent interact over an infinite time horizon. Whereas the principal discounts the future exponentially, the agent discounts future utilities in a quasi-hyperbolic way. In any period, the agent can either work for the principal or not. If the agent works for the principal, he chooses effort which is associated with strictly increasing and convex effort costs. Effort generates a deterministic output which is consumed by the principal. The agent receives a fixed wage payment and a bonus. As effort today yields output today and effort and output are verifiable, a present bias per se does not represent an obstacle for efficient outcomes. A static contract could implement efficient allocations and allow the principal to extract the full surplus. Any distortion will thus be due to the fact that the principal, who is aware of the agent's bias, can design a dynamic contract to exploit the agent's self-control problem

[4]See e.g., Heidhues and Kőszegi (2010) for a literature overview on evidence for time-inconsistency in credit markets, DellaVigna (2009) for a more general review on field evidence for present bias in the population, or Kőszegi (2014) for a literature overview on behavioral contract theory.

[5]Below, in Sections 3.7.4 and 3.7.8, we show that the presence of a moral hazard or an adverse selection problem, respectively, leaves our main results unaffected.

and extract excess rents. Hence, in our setting, long–term contracts are strictly better for the principal than a sequence of spot contracts.[6]

A naive agent, who is not aware of his future self-control problems, is offered a menu of contracts, consisting of a virtual contract which the agent intends to choose in the future and a real contract which he ends up choosing. The virtual contract has a "probation phase" in the first period it is chosen. In this probation phase, the agent receives a low utility as he either has to work extra hard or receives only a very low compensation. This probation phase deters the present-biased agent from actually choosing the virtual contract as it entails immediate costs but only delayed gratification. After the probation phase, the virtual contract grants the agent very attractive compensation. From today's perspective, the prospect of this attractive compensation after the probation phase outweighs the low utility during the probation phase tomorrow and thus makes the agent willing to accept a lower compensation today. However, when tomorrow comes, the lower utility during the probation phase does not seem worth the future benefits anymore and the agent postpones taking up the virtual contract. Consequently, the agent always ends up choosing the real contract which gives him a utility below his outside option.

To come back to the above example of career choice, we find that firms can exploit naive present–biased employees by offering them a lucrative "career" as compensation for (presumably) short-run sacrifices. However, the naive present–biased agent fails to take up this option as he incrementally postpones making an extra investment to get onto the career trajectory. Note that the"immediate costs, delayed rewards"-structure is not a property of the environment but the result of the individually optimal contract design choice by the principal.

In fact we do observe that in many employment settings, the agent is offered the "carrot" of international mobility or lucrative career prospects. However, there is always some additional effort required like organizing a stay abroad or just "going the extra mile" in a new, challenging assignment. While employees think that they will collect these benefits and hence are motivated by them, many will never be in the position to consume them as they will procrastinate and indefinitely postpone making a career.[7] Nevertheless, the employee is willing to accept lower compensation today because of his misconceived effort (and career) expectations. As a consequence, the firm can provide incentives at substantially lower costs. When we analyze an extension of our model with a finite horizon, (young) agents who are at the beginning of their working lives will be exploited more by the principal as they can be more easily lured by career prospects: They work weakly more (strictly so, if there is a binding limited liability constraint) and are paid less. This broadly resembles actual labor market

[6]This differs from standard results in contract theory where, in general, whenever actions taken by the agent in period t affect the output of period t only, a sequence of short-term contracts is as efficient as a long-term contract; see, e.g., Malcomson and Spinnewyn (1988) or Fudenberg, Holmstrom and Milgrom (1990).

[7]Several studies in organizational psychology documented the frustrations of young employees with respect to their success in achieving their occupational expectations; see, e.g., Garavan and Morley (1997) or Rindfuss et al. (1999).

patterns where we find age earning profiles with relatively better paid but less productive older workers.[8]

Finally, note that how much the principal can exploit the agent with the real contract depends on how auspicious the virtual contract appears. Since the agent is promised the total surplus of the virtual contract, and this (discounted) value is substracted from the agent's real payments, anything that makes this surplus look more attractive from today's perspective – such as a higher "standard" discount factor or lower outside options – actually harms the agent. This gives rise to the additional interpretation of our model where a more stringent employment protection actually harms a naive present–biased agent: Note that the principal's outside option (which can also be negative) includes the cost the latter has to bear when firing the agent. Therefore, higher firing costs (where we assume that those are equivalent to a more stringent employment protection) increase the future virtual surplus and consequently reduce the amount the agent is paid in the real contract. Sophisticated present–biased agents, i.e., those who are aware of their future self-control problems, do not suffer from their self-control problems because it is optimal to compensate them for their effort at the end of a given period. Hence, their present bias is irrelevant and they are effectively in the same situation as agents without any present bias (because of the technological independence across periods).

Next to the extension of our model within a finite horizon, we analyze a number of additional extensions to document the robustness of our findings. We let the agent be protected by limited liability and allow for a weaker present bias in the monetary domain, as well as for moral hazard, for competition in the labor market and hence varying bargaining power of the agent and partial naivete of the agent. Furthermore, we allow the agent to learn about his type over time and consider the case of heterogeneous and unobservable agent types. None of these extensions changes our findings qualitatively. They rather offer additional insights:

If the agent is protected by limited liability, the principal cannot always fully exploit the agent when letting him work at the efficient effort level. Instead, once limited liability becomes binding and reducing the payment is no longer possible as the agent already receives a zero payment in the real contract, the principal resorts to extracting additional rent by inducing the agent to work harder than the efficient effort level. Therefore, seemingly contradicting the standard procrastination result, in our model, naive present–biased agents might work harder than agents without a present bias.

When we allow for a weaker present bias in the monetary domain, our results do not change qualitatively. Moreover, we find that a naive agent who has a present bias in the effort domain may benefit from a bias in the monetary domain because this would reduce the degree to which compensation could be shifted from the real into the virtual contract.

[8]This general pattern, overly hard work early on in an agent's working live and relatively less hard work towards the end of a career is also generated in models of career concerns; see Holmstrom (1999). However, there this pattern is driven by a signal-jamming logic.

Under the assumption of unobservable agent types, the principal generally still offers exploitative menus of contracts. In this case, only the extent of exploitation is lower than when types are observable in order to make each agent select the menu intended for his type.

Allowing for moral hazard, bargaining, the existence of partially naive types and the possibility of learning about one's type do not affect the structure of profit-maximizing contracts. Neither moral hazard nor bargaining change the basic structure of the optimal contract. A (partially) naive agent expects that in the future choosing the virtual contract will be strictly better than choosing the real contract, while in fact, tomorrow he will be just indifferent. For this decision, his true present bias is relevant, not his perception about its future extent. This result on partial naivete also implies that the offered contracts do not change when the agent is able to learn that he is time-inconsistent (unless he learns his bias perfectly and becomes fully sophisticated).

The paper is organized as follows. Following a literature overview in Section 3.2, Section 3.3 introduces our model. Section 3.4 deals with the benchmark cases of the principal facing a rational or a sophisticated agent. Section 3.5 sets up the principal's problem when facing a naive time-inconsistent agent and Section 3.6 presents our main results. Section 3.7 deals with a finite time horizon, limited liability, a weaker present bias in the monetary domain, moral hazard, competition and bargaining, partially naive agents, learning, and adverse selection. Section 3.8 concludes and Appendices C.1 and C.2 collect all proofs.

3.2 Related Literature

The paper relates to the literature on inconsistent time preferences, first formalized by Strotz (1955), who allows for an individual's discount rate between two periods to depend on the time of evaluation. Phelps and Pollak (1968) argue that since inconsistent time preferences affect savings, the possibility of individuals having a present bias should be incorporated into growth models. They develop the workhorse model to analyze inconsistent time preferences, the so-called $\beta - \delta$-model: An individual gives extra weight to utility today over any future period, but discounts all future instances in a standard exponential way. Laibson (1997) shows that given people have inconsistent time preferences, choices that seem suboptimal – for example the purchase of illiquid assets – can actually increase an individual's utility by binding future selves and hence providing commitment. O'Donoghue and Rabin (1999a) compare the welfare of so-called "sophisticated" and "naive" individuals, where the former are aware of their time inconsistency and the latter are not. They show that in case there is no commitment device an agent can take up, sophisticated agents are better off than naive ones if costs are immediate but rewards are delayed, and vice versa if rewards are immediate but costs are delayed.

Besides numerous theoretical contributions, there also is substantial evidence suggesting that people make decisions that are not consistent over time. For example, consider Shui and

Ausubel (2005) or DellaVigna and Malmendier (2004, 2006), who document present–biased behavior for credit card usage and health club attendance respectively. Kaur et al. (2010, 2015) provide evidence from a field experiment with data entry workers that self-control problems at work are important. The recent study by Augenblick, Niederle and Sprenger (2015) is particularly interesting for us as they document that subjects show a considerable present bias in effort (while they only find limited time inconsistency in monetary choices). This suggests to us that studying the role of present bias in the workplace and in workers' careers is a particularly relevant and promising topic of research.

There is also a literature focusing on optimal contracting choices when agents exhibit time-inconsistencies. O'Donoghue and Rabin (1999b) develop a model where a principal hires an agent with present–biased preferences in order to complete a task, but the agent's costs of doing so vary. If the latter is the agent's private information, it can be optimal to employ a scheme of increasing punishments if the task has not been completed yet. Whereas the interaction in O'Donoghue and Rabin (1999b) is basically one-shot, i.e., the relationship between principal and agent is terminated once the task has been completed, we show how repeated interaction can have a substantial impact on optimal contracts, by allowing the agent to choose among a menu of contracts (with history-dependent elements) in every period.

Similar to O'Donoghue and Rabin (1999b), Gilpatric (2008) analyzes a contracting problem between a principal and a (risk-neutral) agent where the latter's effort choice is observable. He shows that it might be optimal to let the agent slack-off after signing the contract (where the slacking-off is unpredicted by a naive agent), however requiring the agent to "compensate" the principal for that. We relate to Gilpatric (2008) in the sense that agents make different choices than they expected, the principal foresees that and can exploit it. However, the setting in Gilpatric (2008) is not fully dynamic as principal and agent only interact once.

Yılmaz (2013) analyzes a repeated moral hazard setting with a risk-neutral principal and a risk-averse agent, where the latter has β–δ-preferences. He shows that due to time-inconsistency, a lesser degree of consumption smoothing is optimal. However, he restricts the contracting space to consist of only one element, hence does not allow the agent to revise his observable actions in future periods.

From a modeling perspective, but analyzing a rather different environment, the paper closest to ours is Heidhues and Kőszegi (2010). They analyze contracting with time-inconsistent consumers in a competitive credit market and find that naive consumers are attracted by contracts characterized by cheap baseline repayment terms and large fines for delays. Naive consumers overborrow and end up paying fines and suffering welfare losses. These can be reduced by prohibiting excessive fines. The results in this chapter are driven by a related basic intuition. Nevertheless the application, institutional details, and interpretation of results in this chapter are different.

Furthermore, Eliaz and Spiegler (2006) analyze optimal contracts for consumers with different degrees of sophistication. The principal benefits monotonically from a lower degree of sophistication. They do not allow for repeated interaction, though, and hence do not aim to characterize dynamic contracts.

Being naive about one's future time preferences might also be perceived as a specific form of overconfidence. The following two papers look at overconfidence in a principal-agent setting with moral hazard. Both Gervais, Heaton and Odean (2011) and de la Rosa (2011) find that, for certain parameter values, the principal can incentivize the agent more cheaply as the agent overestimates the likelihood to succeed. In our model the agent's bias makes it cheaper for the principal to incentivize the agent as well, however, the structure of the optimal contract is entirely different.[9] Hoffman and Burks (2015) analyze a data set of a trucking firm with detailed information on drivers' beliefs about their future performance. Many drivers overestimate their future productivity and adjust their beliefs only slowly. The firm benefits from - as Hoffman and Burks (2015) interpret it - the drivers' overconfidence as drivers with a larger bias are less likely to quit.

3.3 Model Setup

Technology & Per-Period Utilities There is one risk neutral principal ("she") and one risk neutral agent ("he"). We consider an infinite time horizon where time is discrete and periods are denoted by $t = 0, 1, 2, ...$ In any period t, the agent can either work for the principal or not, which is described by $d_t \in \{0, 1\}$. If $d_t = 1$, the agent works for the principal and chooses effort $e_t \geq 0$. This choice is associated with strictly increasing and convex effort costs $c(e_t)$, where $c(0) = 0, c'(0) = 0$, and $\lim_{e_t \to \infty} c' = \infty$. Effort e_t generates a deterministic output $y(e_t) = e_t\theta$ which is consumed by the principal. Furthermore, the agent receives a fixed wage payment w_t and a bonus $b(e_t)$. Note that we do not impose a limited liability constraint on the agent, hence the payments can (and under some instances will) be negative (implying payments from agent to principal). We consider the case of limited liability below, in Section 3.7.2. The agent's payoff in period t when $d_t = 1$ is

$$w_t + b(e_t) - c(e_t)$$

whereas the principal receives

$$e_t\theta - b_t(e_t) - w_t.$$

If $d_t = 0$, i.e., the agent does not work for the principal in period t, he receives his outside option $\overline{u}$. The principal's outside option in this case is denoted by $\overline{\pi}$ (where both, $\overline{u}$ and $\overline{\pi}$ are not restricted to non-negative values).

[9]One could conceive of a specific model based on overconfidence that would generate similar dynamic patterns as our setting based on present bias. However, such a model would have to be structurally very close to our model where we assume that the present bias makes the agent expect to be better off tomorrow.

The effort level maximizing total surplus if the agent works for the principal, denoted by e^{FB}, is implicitly defined by

$$\theta - c'(e^{FB}) = 0.$$

For the remainder of the paper, we assume $\theta e^{FB} - c(e^{FB}) - \overline{u} - \overline{\pi} > 0$, i.e., the employment relationship is socially efficient.

Time Preferences The principal discounts the future exponentially with a constant factor $\delta \in (0, 1]$, whereas the agent discounts future utilities in a quasi-hyperbolic way according to Phelps and Pollak (1968) and Laibson (1997): While current utilities are not discounted, future (period t) utility is discounted with a factor $\beta\,\delta^t$, with β and $\delta \in (0, 1]$ (which agent and principal have in common). Hence, the agent is present–biased and his preferences are dynamically inconsistent. Concerning the agent's belief about his future preferences for instant gratification, we follow O'Donoghue and Rabin (2001) and their description of partial naivete. While an agent discounts the future with the factor β, he thinks that in any future period, he will discount the future with the factor $\hat{\beta}$, where $\beta \leq \hat{\beta} \leq 1$. In other words, the agent may be aware of his present bias but expects it to be weaker than it actually is. We will mainly focus on two extreme cases, $\hat{\beta} = 1$ and $\hat{\beta} = \beta$. The first case describes a fully naive agent who – in every period – thinks that from tomorrow on, his present bias will disappear and he will discount the future exponentially. The second case describes a sophisticated agent who is fully aware of his persistent present bias. In Section 3.7.6 we explicitly allow for partial naivete, i.e., $\hat{\beta} \in (\beta, 1)$ and show that the outcome in this case is *exactly* the same as with a fully naive agent. Furthermore, in Section 3.7.7 we also allow for learning in the sense that $\hat{\beta}$ decreases over time and show that results are robust as long as learning is not perfect. Finally, note that we assume the agent to be equally present–biased regarding money and effort.

Perceptions Here, we have to distinguish between intra- and inter-player perceptions. Concerning the first, we assume the agent's beliefs to be dynamically consistent as defined by O'Donoghue and Rabin (2001) (p. 129), i.e., the agent's belief of what he will do in period τ must be the same in all $t < \tau$.

Concerning inter-player perceptions, we assume common knowledge concerning the principal's time preferences. Furthermore, the principal is aware of the agent's time preferences and knows his values β and $\hat{\beta}$. This implies that for a (partially) naive agent, the principal correctly predicts any (potential) discrepancy between intended and realized behavior. Finally, we assume that the agent believes the principal to share his own perception of himself. A (partially) naive agent hence is convinced that the principal also perceives the agent's future present bias to be characterized by $\hat{\beta}$. In Section 3.7.8 we explicitly allow for unobservable types and derive optimal screening contracts.

Contractability, Timing, and Histories To isolate the effect of the agent's present bias on the structure of the employment relationship, we abstract from any other potential agency problem. Hence, the agent's effort as well as wage and bonus payments are verifiable and the principal can commit to long-term contracts.[10] The principal's commitment is limited by the firm value, though, and she can always escape her obligations by declaring bankruptcy and proceeding to consume her outside option $\overline{\pi}$ in every subsequent period.[11]

We do not allow the agent to commit to long-term contracts. Note that this assumption turns out to be without loss of generality: Since the principal benefits from the (partially) naive agent's misperceptions about his future choices, she actually does not want to grant the agent the opportunity to commit to a long-term contract. Moreover, if the agent is either not present–biased or sophisticated about his present bias, any long-term commitment on his side will not increase profits. Hence, we can restrict attention to situations where the agent is free to leave in any period.[12]

In the first period of the game, in $t = 0$, the principal makes a take-it-or-leave-it long-term contract offer, denoted by $\boldsymbol{C}$, to the agent. This offer consists of a menu of contracts (with finitely many elements) for every future period, contingent on any possible history. Each of these contracts contains a fixed wage payment and a bonus for every potential effort level. The principal is fully commited to this long-term contract and could only walk away from her obligations by declaring bankruptcy.

The menu of contracts offered in period t is denoted by C_t. It has I_t elements, where a single element is indexed $i_t \in \{0, .., I_t\}$, i.e., $C_t = \left\{C_t^{i_t}\right\}_{i_t=1}^{I_t} = \left\{w_t^{i_t}, b_t^{i_t}(e_t)\right\}_{i_t=1}^{I_t}$, where $w_t^{i_t} \in (-\infty, \infty)$ for each i_t, and $b_t^{i_t}(e_t) \in (-\infty, \infty)$ for each i_t and each e_t. At the beginning of every period t, the principal first decides whether to declare bankruptcy or whether to offer C_t. This choice is denoted by $d_t^P \in \{0, 1\}$, where $d_t^P = 0$ indicates a bankruptcy and implies that $C_t = \{\emptyset\}$.

Declaring bankruptcy is an irreversible decision (hence, $d_{t+1}^P = 0$ automatically follows from $d_t^P = 0$), inducing principal and agent to consume their outside utilities $\overline{\pi}$ and $\overline{u}$ in every subsequent period. Given $d_t^P = 1$, the agent chooses whether to work for the principal or not, hence selects $d_t \in \{0, 1\}$. If $d_t = 0$, principal and agent consume their outside utilities in the respective period. If $d_t = 1$, the agent selects one element out of C_t, where this choice is denoted by $\hat{i}_t \in I_t$. Then, the agent receives $w_t^{\hat{i}_t}$ and makes his effort choice, triggering the (automatically enforced) bonus payment $b_t^{\hat{i}_t}(e_t)$. Finally, the principal consumes the output $e_t\theta$ and the game moves on to the next period.

The publicly observable events during periods $t \geq 0$ are $h_t = \left\{d_t^P, C_t, d_t, \hat{i}_t, w_t, e_t, b_t(e_t)\right\}$. The history of publicly observable events at the beginning of period t is $h^t = \cup_{\tau=1}^{t-1} h_\tau \cup \{\boldsymbol{C}\}$,

[10] This strong commitment could be endogenized in a setting with many agents where the principal's behavior can be observed by everyone, and where she cares about her reputation to keep her promises.

[11] Alternatively, we can assume that the principal can fire the agent at some cost – for example reflecting the degree of employment protection – implying that $\overline{\pi}$ could also be negative. Bankruptcy or firing costs are then given by $-\frac{\overline{\pi}}{1-\delta}$.

[12] Again, any costs for the agent to leave his current occupation could be captured by an appropriate choice of $\overline{u}$.

with $h^0 = \{\emptyset\}$. Furthermore, H^t is the set of histories until t and $H = \cup_t H^t$ the set of histories.

Strategies Following O'Donoghue and Rabin (1999b), we use the phrase *perception-perfect strategy* to describe players' strategies. Such a strategy specifies a player's behavior given dynamically consistent beliefs about future behavior. Whereas a time-consistent or a sophisticated agent correctly predicts his future behavior, the same is not necessarily true for a (partially) naive agent who has wrong beliefs concerning his future time preferences.

The principal's strategy is denoted by σ_P. In period $t = 0$, it determines the long-term contract $\boldsymbol{C}$. In every period $t \geq 0$, $\boldsymbol{C}$ maps the history $h^t \in H^t$ into an offered menu of contracts, C_t. Following any history $h^t \in H^t$, $t \geq 0$, σ_P also determines whether the principal follows $\boldsymbol{C}$ or terminates the contract by declaring bankruptcy.

An agent's strategy is denoted by σ_A and – given $h^t \cup \{d_t^P\} \cup \{\boldsymbol{C}\}$ – determines d_t and eventually $\hat{i}_t$ and e_t.

Real and Virtual Contract Without loss of generality, we can restrict C_t to either consist of one or two elements, depending on the agent's naivete. If the agent is sophisticated or not present–biased, he does not have misperceptions concerning his future behavior, and C_t consists of exactly one element. Given that a profit-maximizing contract is included in the menu, adding additional contracts will have no (or even adverse) effects on profits, as these other elements are either never chosen by the agent or – if they are chosen – yield (weakly) lower profits than the profit-maximizing one.

If the agent is (partially) naive, the principal finds it optimal to let C_t consist of exactly two elements: the element that the agent believes to choose in the future, and the element the agent is actually going to choose (which is correctly anticipated by the principal). We call the former the virtual contract and describe the respective components (as well as the effort level the agent expects to choose when selecting the virtual contract) with a superscript "v". The latter is called real contract and its components (as well as the effort level the agent chooses when selecting the real contract) are described with the superscript "r".

Payoffs The agent's actually realized utility stream at the beginning of any period t in a setting where the principal never declares bankruptcy is

$$U_t^r = d_t^r\,(b_t^r + w_t^r - c(e_t^r)) + (1 - d_t^r)\,\overline{u} + \beta \sum_{\tau=t+1}^{\infty} \delta^{\tau-t}\,[d_\tau^r\,(b_\tau^r + w_\tau^r - c(e_\tau^r)) + (1 - d_\tau^r)\,\overline{u}].$$

For naive types, this real utility may be different from their perceived virtual payoff the agent believes to receive. If an agent prefers to choose the virtual contract rather than the real contract in the future, then this payoff includes real current and virtual future payoff and is indicated with the superscript "rv".

$$U_t^{rv} = d_t^r\,(b_t^r + w_t^r - c(e_t^r)) + (1 - d_t^r)\,\overline{u} + \beta \sum_{\tau=t+1}^{\infty} \delta^{\tau-t}\,[d_\tau^v\,(b_\tau^v + w_\tau^v - c(e_\tau^v)) + (1 - d_\tau^v)\,\overline{u}].$$

The principal's expected (which coincides with the actually realized) payoff in any period t in a setting where the principal never declares bankruptcy is

$$\Pi_t^r = \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left[d_\tau^r \left(e_\tau^r \theta - b_\tau^r - w_\tau^r\right) + \left(1 - d_\tau^r\right) \overline{\pi}\right].$$

If an agent prefers to choose the virtual contract rather than the real contract in the future, though, then he expects the principal to maximize

$$\Pi_t^{rv} = d_t^r \left(e_t^r \theta - b_t^r - w_t^r\right) + \left(1 - d_t^r\right) \overline{\pi} + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left[d_\tau^v \left(e_\tau^v \theta - b_\tau^v - w_\tau^v\right) + \left(1 - d_\tau^v\right) \overline{\pi}\right].$$

For later use, we further define perceived future virtual payoff for the agent

$$U_t^v = \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left[d_\tau^v \left(b_\tau^v + w_\tau^v - c(e_\tau^v)\right) + \left(1 - d_\tau^v\right) \overline{u}\right]$$

and the principal

$$\Pi_t^v = \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left[d_\tau^v \left(e_\tau^v \theta - b_\tau^v - w_\tau^v\right) + \left(1 - d_\tau^v\right) \overline{\pi}\right].$$

Equilibrium We define a *perception-perfect equilibrium* to be one that maximizes players' payoffs, given each player's perception of their own future behavior as well as of the other's future behavior.

The principal hence chooses $\boldsymbol{C}$ in order to maximize Π_0^r, and d_t^P in order to maximize Π_t^r, taking into account the agent's actual, i.e., "real", behavior. The (partially) naive agent, though, expects the principal to choose d_t^P in order to maximize Π_t^{rv} or Π_t^v, respectively.

The (partially) naive agent makes optimal choices given her current preferences and her perceptions of future behavior, i.e., expecting to have different time preferences in the future. Hence, he chooses d_t, $\hat{i}_t$ and e_t in order to maximize U_t^{rv}, given that an agent believes to choose the virtual contract rather than the real contract in the future.

3.4 Benchmarks: Time-Consistent and Sophisticated Agent

We first derive two benchmarks, profit-maximizing menus of contracts with time-consistent and with sophisticated present-biased agents.

Time-Consistent Agent Consider the case of a time-consistent agent, i.e., an agent who has $\beta = \hat{\beta} = 1$. Since the agent's effort is verifiable, effectively there is no agency problem that must be addressed. In this case, the principal does not need to make use of her ability to make long-term commitments. Instead, she can use a series of spot contracts in order to make the agent choose surplus-maximizing effort in every period, and collect the whole surplus for herself.

One possibility to generate such an outcome is to offer the following contract in every period t: $w = \overline{u}$, $b(e^{FB}) = c(e^{FB})$, and $b(\tilde{e}) = 0$ for $\tilde{e} \neq e^{FB}$. The agent always accepts such a contract and exerts effort e^{FB}. Since the principal extracts the whole surplus, she has no incentives to declare bankruptcy. The optimal menu of contracts with a time-consistent agent hence maximizes the surplus and holds the agent down to his outside option.

Sophisticated Present-Biased Agent If the agent is present-biased and sophisticated (i.e., $\hat{\beta} = \beta$), he is aware of his future preferences and hence of the choices he is going to make in the future. Then, as with a time-consistent agent, it is sufficient to let $\boldsymbol{C}$ consist of only one element in every period. Generally, the principal has the opportunity to reward the agent for effort with a bonus paid at the end of the period, or with the promise of a higher continuation payoff. Since $\beta < 1$, though, the agent effectively discounts the future at a higher rate than the principal does. Therefore, it is (now strictly) optimal for the principal to also offer the following contract in every period t: $w = \overline{u}$, $b(e^{FB}) = c(e^{FB})$, and $b(\tilde{e}) = 0$ for $\tilde{e} \neq e^{FB}$. This contract makes the agent accept the contract in every period, induces him to choose the surplus-maximizing effort level, and allows the principal to extract the whole surplus. Note that this contract cannot be improved upon by taking into account the agent's effective lower discount factor and shifting payments to period 0. In this case, the agent would simply walk away after this period. Concluding, time-inconsistencies have no impact on the optimal contract if the agent is sophisticated about his present bias. This is because the production technology is effectively static in the sense that there are no technological linkages across periods, because the agent can immediatetly be compensated for his effort, and because effort is verifiable.

3.5 The Principal's Problem with a Naive Agent

This section considers the principal's problem when facing a naive present–biased agent, i.e., whose $\hat{\beta} = 1$. In Subsection 3.7.6, we show that the results derived here remain unaffected if $\hat{\beta} \in (\beta, 1)$, i.e., for the case of partial naivete. First of all, note that as with a sophisticated agent, the possibility of writing formal spot contracts to motivate the agent implies that the existence of a present-bias does not automatically trigger inefficiencies. The same spot contract optimally offered to a sophisticated agent could also be offered to naive present-biased agents, yielding exactly the same outcome. However, now the principal can design a menu of long-term contracts, which can exploit the naive agent's misperception of his future behavior. She can include elements into $\boldsymbol{C}$ that seem optimal for the agent from the perspective of earlier periods – but that are not chosen once the agent actually faces the respective choice. As discussed above, we can without loss of generality restrict $\boldsymbol{C}$ to consist of exactly two elements in every period: The contract the agent actually chooses and the contract the agent had planned to choose from the perspective of earlier periods. We call the former contract the real *contract* and add the superscript "r" to all its components and the latter contract the virtual *contract* and add the superscript "v" to all its components. Both contracts can be contingent on the full history of the game, h^t, i.e., $C(h^t) = \{C^r(h^t), C^v(h^t)\} = \{(w^r(h^t), b^r(e^r_t, h^t)), (w^v(h^t), b^v(e^v_t, h^t))\}$. Hence, $\boldsymbol{C}$ also gives values $U^r(h^t)$, $U^{rv}(h^t)$, $U^v(h^t)$, $\Pi^r(h^t)$, $\Pi^{rv}(h^t)$ and $\Pi^v(h^t)$ for every history h^t.

In the following, we focus on contracts that, on the equilibrium path, have $d_t^r = d_t^v = d_t^P = 1$ in all periods t. Hence, the principal never declares bankruptcy, and the agent always accepts the employment offer. Any (temporary or permanent) termination of the relationship can never be optimal - simply because the principal could always include the optimal contract for time-consistent and sophisticated agent in the long-term contract offered to the naive agent, which the latter would accept if the alternative was $d_t = 0$.

The elements of the long-term contract $\boldsymbol{C}$ must satisfy several constraints. Different from a "standard" contracting problem, certain constraints must also hold for the virtual contract. Furthermore, it must be ensured that the agent actually chooses the real contract in every period, but plans to choose the virtual contract in all future periods.

This gives rise to four classes of constraints, where the first three are standard in contracting problems: 1) **Individual rationality constraints** for the agent (IRA) which make him accept one of the offered contracts (compared to reject all of them for the respective period and consume his outside option instead). 2) **Individual rationality constraints** for the principal (IRP) that keep her from declaring bankruptcy. 3) **Incentive compatibility constraints** for the agent (IC) that induce him to select the principal's preferred effort level. These constraints must also hold for the virtual contracts, i.e., for future histories that never materialize. 4) **Selection constraints** ensure that the agent keeps **choosing** the real contract in every period, while still **intending** to choose the virtual contract in all future periods.

Individual rationality constraints for the agent For every history h^t, it must be optimal for the agent to accept the real contract (expecting to choose equilibrium effort e_t^r and to select the virtual contract in all future periods), compared to rejecting any contract and consuming $\overline{u}$ in the respective period:

$$w^r(h^t) + b^r(h^t, e_t^r) - c(e^r) + \beta\delta U^v\left(h^t \cup \{d_t = 1\}\right) \geq \overline{u} + \beta\delta U^v\left(h^t \cup \{d_t = 0\}\right). \quad \text{(rIRA)}$$

Note that when stating $U^v(\cdot)$, we only include the elements that are relevant for the respective constraint (and do the same when describing other constraints), assuming that all other elements are chosen as prescribed by play on the equilibrium path.

Furthermore, $\boldsymbol{C}$ has to be such that the agent expects to accept a contract in all future periods, i.e., for all h^t,

$$U^v\left(h^t \cup \{d_t = 1\}\right) \geq U^v\left(h^t \cup \{d_t = 0\}\right) \quad \text{(vIRA)}$$

has to hold.[13] However, an individual rationality constraint is not needed for U^r because the agent does not expect to select the real contract in future periods. I.e.,

$$w^r(h^t) + b^r(h^t,\, e_t^r) - c(e^r) + \beta\delta U^r\left(h^t \cup \{d_t = 1\}\right) \geq \overline{u} + \beta\delta U^r\left(h^t \cup \{d_t = 0\}\right)$$

[13]Note that vIRA also implies $U^v(h^t) \geq \frac{\overline{u}}{1-\delta}$ for all histories.

does *not* have to hold. In fact, this constraint is generally violated and the agent effectively will be exploited and receive less than his outside option.

Individual rationality constraints for the principal For every history h^t, the following constraints must be satisfied for the principal:

$$\Pi^{rv}(h^t) \geq \frac{\overline{\pi}}{1-\delta} \tag{vrIRP}$$

$$\Pi^{v}(h^t) \geq \frac{\overline{\pi}}{1-\delta} \tag{vIRP}$$

If either of these constraints was not satisfied, the agent would expect the principal to not honor her obligations in the virtual contract and instead shut down. A real IR constraint for the principal, $\Pi^r(h^t) \geq \overline{\pi}/(1-\delta)$, is not required because of the agent's misperceptions about his own behavior. However, it is obvious that shutting down will never be optimal for the principal because she could at any point offer the optimal contract for the time-consistent or sophisticated agent and thereby collect a rent that is strictly positive.

Incentive compatibility constraints It has to be in the agent's interest to choose equilibrium effort e_t^r, given the compensation he receives today, and given his expectation of future (virtual) payoffs. This gives rise to a real incentive compatibility constraint,

$$-c(e_t^r) + b^r(h^t, e_t^r) + \beta\delta U^v\left(h^t \cup \{e_t^r\}\right) \geq -c(\tilde{e}_t^r) + b^r(h^t, \tilde{e}_t) + \beta\delta U^v\left(h^t \cup \{\tilde{e}_t\}\right), \tag{rIC}$$

which has to hold for all histories h^t and effort levels $\tilde{e}_t$ in a given contract menu.

The agent also has to expect to select on-path effort levels in the virtual contract, which for all histories h^t and effort levels $\tilde{e}_t$ gives

$$-c(e_t^v) + b^v(h^t, e_t^v) + \beta\delta U^v\left(h^t \cup \{e_t^v\}\right) \geq -c(\tilde{e}_t^r) + b^r(h^t, \tilde{e}_t) + \beta\delta U^v\left(h^t \cup \{\tilde{e}_t\}\right). \tag{vIC}$$

Selection constraints Finally, the agent has to select the real contract in every period, however expect to select the virtual contract in the future. For every history h^t, this yields the constraints

$$w^r(h^t)+b^r(h^t, e^r)-c(e_t^r)+\beta\delta U^v\left(h^t \cup \left\{\hat{\imath}_t = r\right\}\right) \geq w^v(h^t)+b^v(h^t, e^v)-c(e_t^v)+\beta\delta U^v\left(h^t \cup \left\{\hat{\imath}_t = v\right\}\right) \tag{rC}$$

and

$$w^v(h^t)+b^v(h^t, e^v)-c(e_t^v)+\delta U^v\left(h^t \cup \left\{\hat{\imath}_t = v\right\}\right) \geq w^r(h^t)+b^r(h^t, e^r)-c(e_t^r)+\delta U^v\left(h^t \cup \left\{\hat{\imath}_t = r\right\}\right). \tag{vC}$$

Note that due to naivete, β does not feature in (vC).

Objective The principal's objective is to offer a long-term menu of contracts $\boldsymbol{C} = \{C^r(h^t), C^v(h^t)\}$ for all $h^t \in H$ that maximizes $\Pi^r(h^0)$, subject to (rIRA), (vIRA), (rIRP), (vIRP), (rIC), (vIC), (rC) and (vC) that must hold for any potential history. In the following, we first simplify the problem and then characterize an optimal long-term menu of contracts, $\boldsymbol{C}$.

3.6 Results

3.6.1 Preliminaries

First, we show that the real contract can be stationary without loss of generality, hence its components are independent of the (on-path) history of the game. The same is true for the virtual contract, with the exception of the first period where it is (expected to be) chosen.

Lemma 3.1. *The real contract is stationary, hence independent of the history of the game. The virtual contract is independent of calendar time and is of the form C^v_τ, where τ counts the virtual contract's number of periods after it has first been chosen. Furthermore, the virtual contract is stationary for all $\tau \geq 2$.*

The proofs for all lemmas and propositions are collected in Appendices C.1 and C.2.

In the remainder of this article, we use the subscript 1 for the first period of the virtual contract. For all subsequent periods, we omit time subscripts. The real contract hence consists of e^r, w^r, b^r, the virtual contract of e^v_1, w^v_1, b^v_1 for the first and e^v, w^v, b^v for all subsequent periods after it has been selected.

Stationarity of the real and later periods of the virtual contract is straightforward, as the game is stationary and principal and agent are both risk-neutral. However, in order to be able to exploit the agent's present bias and keep him below his outside option, the virtual contract must be different in the first period it is chosen: In order to extract rents from the agent in the real contract – the contract that is actually selected in every period (and hence triggers real transfers) – the principal shifts as much as possible of the compensation promised to the agent into the virtual contract. However, the principal has to ensure that the virtual contract is never selected by the agent. This is achieved by designing the first period of the virtual contract to be sufficiently unattractive for the agent – but sufficiently attractive to expect him to still opt for C^v in future periods. In our first main result, we show that the offered menu of contracts effectively harms the agent.

Proposition 3.1. *If the agent is naive and has $\beta \in (0, 1)$, then in the profit-maximizing menu of contracts, $w^r - c(e^r) + b^r < \overline{u}$.*

The principal uses the promise of a virtual future surplus to keep the agent below his outside option while still accepting the contract – which is perceived to be optimal by the agent as he expects to earn a rent in the future from choosing the virtual contract – and working for the principal in every period.[14]

In Appendix C.1, we show that the profit-maximizing structure of $\boldsymbol{C}$ allows us to simplify the problem by eliminating several constraints: Because the agent receives a rent in the virtual contract, the respective (IR) constraints on the agent's side can be omitted. Furthermore, it is optimal to make the virtual contract attractive enough for the agent to always expect to choose it in future periods. It only remains for the principal to make sure that the agent is not selecting it when actually facing the choice in every period. Note that, unless $\beta = 1$ or the agent is always kept at his outside option, (vC) and (rC) cannot bind simultaneously: Either the agent believes that he will be indifferent in the future or he will actually be indifferent. Because the principal could profitably make the real contract less attractive if the agent were not indifferent when choosing between the contracts, (rC) binds and (vC) does not. Furthermore, we can set $w^r = w_1^v = w^v = 0$ without loss of generality, because the principal can arbitrarily substitute between wages and bonus payments as effort is verifiable. In Appendix C.1 we show that the principal's objective can be expressed as the problem of maximizing

$$\Pi^r = \frac{e^r\theta - b^r}{1-\delta},$$

subject to

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}\left(b^v - c(e^v) - \overline{u}\right)\right] \geq 0 \tag{rIRA}$$

$$e^v\theta - b^v \geq \overline{\pi} \tag{vIRP}$$

$$\begin{aligned} -c(e^r) + b^r - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}\left(b^v - c(e^v) - \overline{u}\right)\right] \\ \geq (b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta\beta}{1-\delta}\left(b^v - c(e^v) - \overline{u}\right). \end{aligned} \tag{rC}$$

3.6.2 Profit-Maximizing Contract

In this section, we fully characterize a profit-maximizing menu of contracts $\boldsymbol{C}$.

Lemma 3.2. *A profit-maximizing menu of contracts has the following features:*

- *The constraints (rIRA), (rC) and (vIRP) hold with equality*
- $e^r = e^v = e^{FB}$
- $b^r = c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right).$

[14]That the principal can exploit the agent because the agent does not stick to his planned action is not surprising given the literature (see e.g., DellaVigna and Malmendier, 2004, Eliaz and Spiegler, 2006, and Heidhues and Kőszegi, 2010).

The first bullet point of Lemma 3.2 indicates that the principal promises the whole virtual rent to the agent (binding (vIRP)), with the exception of the first virtual period which is constructed to be just sufficiently unattractive to be never selected by the agent (binding (rC)). Here, we also see that if the principal was able to credibly make arbitrary promises and not able to shut down (i.e., if she did not face (IR) constraints), we would not have an equilibrium. In this case, the principal would promise infinitely high payments to the agent in the virtual contract, and extract infinitely high payments from the agent in the real contract. Furthermore, even though the agent expects to get a rent in the future, his expected rent from today's perspective (including today's payoffs from the real contract) is zero (binding (rIRA)).

First-best effort levels in the real and virtual contract (with the exception of e_1^v) not only maximize the surplus that the principal can reap in any case. $e^v = e^{FB}$ also maximizes the future rent the agent expects, making it possible to reduce real payments by the highest feasible amount.

Finally, the real bonus b^r captures the link between real and virtual compensation. The first elements, $c(e^{FB}) + \overline{u}$, capture the agent's effort costs and outside option, hence would constitute his "fair" compensation. The last term characterizes the principal's savings compared to this fair compensation. It amounts to the total expected and discounted rent the agent expects from chosing the virtual contract in the future, which is supposed to serve as the reward for *today's* effort. If this value is rather high, b^r can actually be negative - in this case, the agent has to make a payment to the principal in every period in order to keep the option of receiving future virtual rents. Note that, as we show below in Section 3.7.2, even in the arguably more realistic case of limited liability the agent would not be better off. In such a setting, the principal will respond to a binding limited liability constraint by requiring real effort to be above e^{FB}. While this leaves the agent's real utility unaffected, it entails an allocative inefficiency and allows the principal to extract only a lower rent from the agent compared to the case absent the binding limited liability constraint.

Real payoffs of principal and agent are characterized in Proposition 3.2.

Proposition 3.2. *Real net per-period payoffs of principal and agent are*

$$\pi^r - \overline{\pi} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left(1 + \beta(1-\beta)\frac{\delta^2}{1-\delta}\right)$$

and

$$u^r - \overline{u} = -\beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right),$$

respectively.

Proposition 3.2 implies that agents with levels of β close to 0 or 1 can hardly be exploited by the principal. The real loss for the agent is maximized for intermediate values of β (see

Figure 3.1). Put differently, agents with a medium inclination to procrastinate can be most effectively exploited. Hence, these types are preferred by the principal (see Figure 3.2). The reason for this is that agents with a β close to 1 do not depart from time-consistency very much and agents with a β close to 0 do not care very much about the future, so the virtual future surplus cannot have a big impact on today's choices.

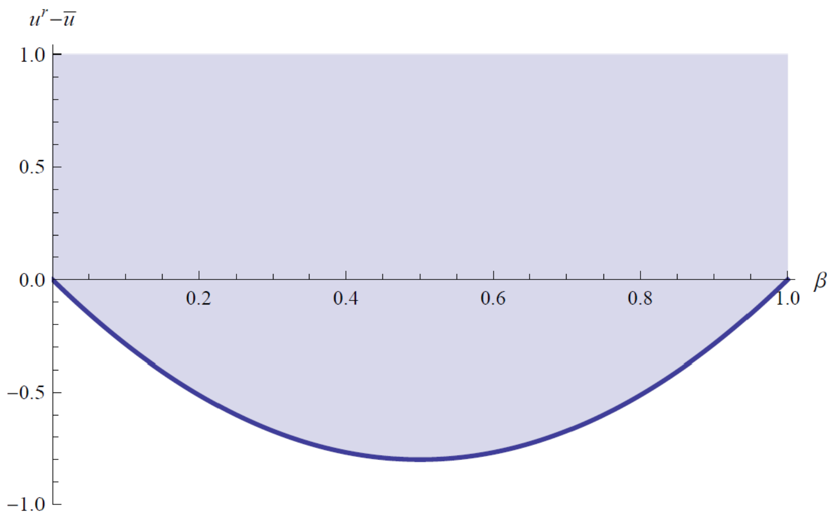

Figure 3.1: Utility as Function of β

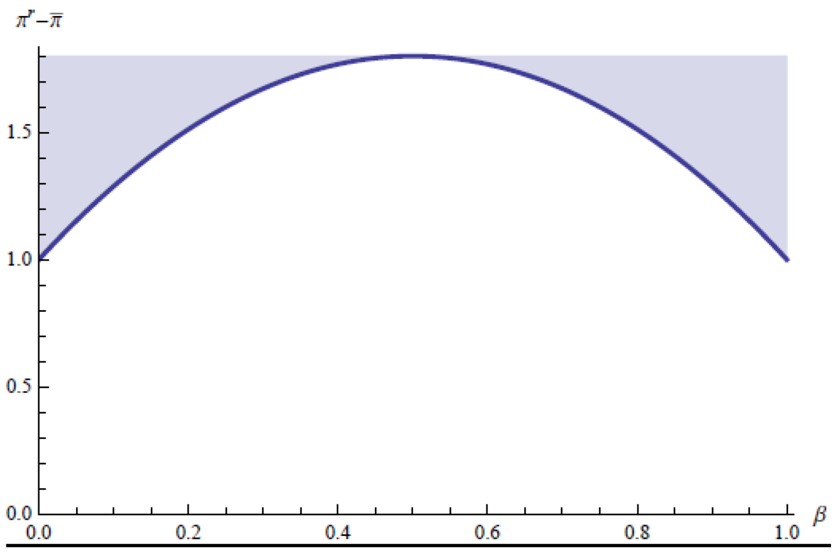

Figure 3.2: Profit as Function of β

Consequently, our model predicts that firms benefit from offering long-term contracts – as compared to a sequence of spot contracts – to agents with an intermediate time-inconsistency, while there are no benefits from offering long-term contracts to other agents. If there were some additional costs of long-term contracts, e.g., because they are more complicated to write, then agents with an intermediate degree of time-inconsistency would be more likely to receive long-term contracts than other agents.

Finally, we take a brief look at the first period of the virtual contract, which is constructed as something like a probation phase in order to deter the agent from ever selecting C^v. It only depends on the difference $c(e_1^v) - b_1^v$, without the exact values of b_1^v and e_1^v being relevant. Note that this also implies that if the agent's present bias would only manifest in one domain, i.e., either in monetary payments or effort, our results would (qualitatively) be the same as we discuss in Section 3.7.3. Using the respective binding constraints gives

$$c(e_1^v) - b_1^v = \frac{\delta\beta}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) - \overline{u}.$$

A policy of subsidizing career-development (which could be captured by a subsidy to effectively increase b_1^v) would not help, since the principal would adjust the required level of e_1^v accordingly in order to keep the agent from selecting the virtual contract.

Figure 3.3 depicts the path of effort over time and Figure 3.4 the path of utility over time.

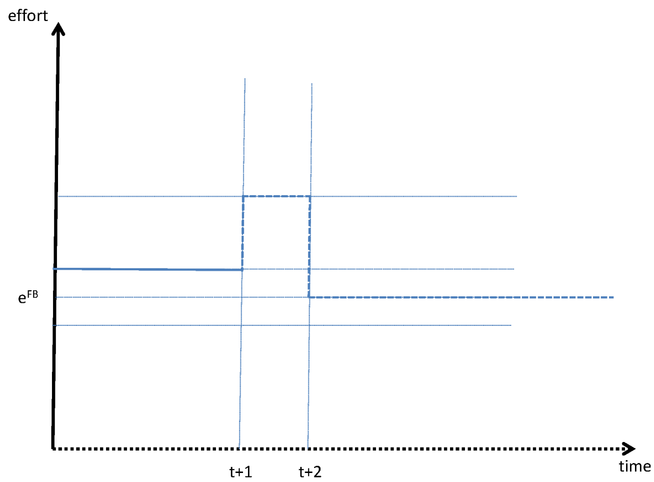

Figure 3.3: Effort Path in Virtual Contract

Note that a high relationship surplus, as well as much "standard" patience (high δ), is bad for the naive agent. Put differently, if the future is not valued very much, then the agent does not want to trade today's payment against future benefits and hence cannot be exploited as much. The agent also benefits if the principal's outside option is higher, which contrasts our analysis to Nash bargaining and relational contracting models. If the principal's outside option is higher, then she can only commit to lower future benefits for the agent. As the agent expects to earn less in the future, he wants to earn more today and can be exploited less. Somewhat counter-intuitively, this argument implies that employment protection, associated with increased layoff costs (corresponding to a reduced outside option), might hurt the agent. This makes it easier for the principal to commit to future employment and hence improves the scope for exploitation.

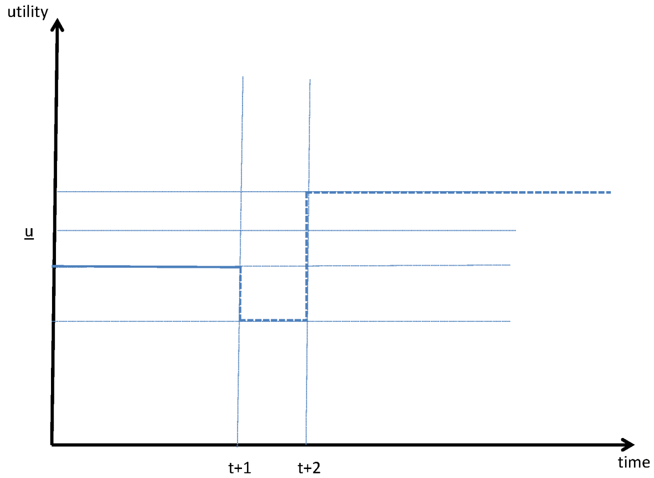

Figure 3.4: Utility Path in Virtual Contract

Corollary 3.1. *The agent's utility is increasing in $\overline{\pi}$, the principal's outside option: Any policy that lowers $\overline{\pi}$, e.g., employment protection, hurts the agent.*

3.7 Extensions

3.7.1 Finite Time Horizon

While employment relations are in general long term, they still might have a pre-defined (maximum) tenure. In particular, in many markets and countries there exists a mandatory retirement age. Here we document that our results are qualitatively robust when we consider a model with a finite time horizon.

Proposition 3.3. *For a finite time horizon T, a profit-maximizing contract has the following features for all periods $t \leq T$:*

- $e_t^r = e_t^v = e^{FB}$.
- $b_t^r = c(e_t^r) + \overline{u} - \beta(1-\beta)\delta^2 \sum_{j=0}^{T-t-2} \delta^j \Big(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\Big), with \sum_{j=0}^{k} x_j := 0 \, \forall k < 0.$
- $u^r - \overline{u} = -\beta(1-\beta)\delta^2 \sum_{j=0}^{T-t-2} \delta^j \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$, *with* $\sum_{j=0}^{k} x_j := 0 \; \forall k < 0$.

Note that for $T \to \infty$, this expression approaches the result of the infinite-horizon case. The more periods there are left, the larger are the total future benefits the principal can ("virtually") promise to the agent. Hence, the worker is willing to accept lower payments today, although he does not actually choose the virtual contract in the future.

At later stages of a career, i.e., in periods closer to retirement, less future benefits can be promised to the agent. In fact, this protects the agent as he cannot be exploited as much.

In combination with a limited liability constraint, which we will analyze in Section 3.7.2[15], younger workers are also more likely to work harder. This finding broadly resembles actual labor market patterns where we find age earning profiles with relatively better paid but less productive older workers.

Corollary 3.2. *For a finite time horizon T, the optimal b^r_t is decreasing in $T-t$. I.e., the principal can exploit the agent more if he is in the early stages of his career.*

3.7.2 Limited Liability

We have shown that without a lower bound on payments, b^r can actually be negative, indicating effective payments from agent to principal *in addition* to the agent's effort. In many cases, though, payments are restricted by some lower bound. In this section, we assume that the agent is protected by limited liability, i.e., payments can not be negative ($b^r \geq 0$). Interestingly (and different from moral hazard problems with limited liability where the agent generally gets a rent), the agent is not better off than before but receives exactly the same level of real utility. This is because the principal optimally responds to a binding limited liability constraint by letting the agent work harder and setting e^r above e^{FB} to extract additional rents. Hence, we might face a situation where the present–biased agent works harder than a sophisticated or time-consistent one. This seems at odds with a popular interpretation of present bias as a source of procrastination and laziness. However, the agent still procrastinates in our setting by pushing off a seemingly even more unattractive *combination* of effort and bonus (the first period of the virtual contract) over and over again.

Summing up, the profit-maximizing contract when payments have to be non-negative is characterized in Proposition 3.4.

Proposition 3.4. *Assume $b, w \geq 0$ in every period t. Then, a profit-maximizing contract $\boldsymbol{C}$ has*

$$e^r \geq e^{FB},\ e^v = e^{FB} \text{ and } b^r = max\Big\{0,\ c(e^{FB}) + \overline{u} - \beta(1-\beta)\tfrac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)\Big\}.$$

Moreover,

- *if $c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) < 0$ holds, the limited liability constraint for the agent's real bonus binds. Then $e^r > e^{FB}$ and is chosen such that $b^r = c(e^r) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) = 0$.*

- *If $c(e^{FB}) + \overline{u} - \beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) \geq 0$, the limited liability constraints do not bind and hence do not affect $\boldsymbol{C}$.*

- *In either case, the agent's real payoff is $u^r - \overline{u} = -\beta(1-\beta)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$.*

[15] A binding limited liability constraint induces the principal to optimally implement an inefficiently high effort level in order to exploit the agent as much as possible.

The agent's real payoff is unaffected by him being protected by limited liability. Hence, the full burden of the inefficient outcome is borne by the principal: Because an effort above e^{FB} reduces total surplus, she can extract less from the agent than when payments are not restricted. While under a binding limited liability constraint for the agent's real bonus, the principal's net profit is $\pi^r - \overline{\pi} = e^r\theta - \overline{\pi}$, it is $\pi^r - \overline{\pi} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left(1 + \beta(1-\beta)\frac{\delta^2}{1-\delta}\right)$ if no limited liability constraints bind.

3.7.3 Weaker Present Bias in the Monetary Domain

Augenblick et al. (2015) find in an experiment that people have a stronger present bias in effort than in monetary choices. Our results are qualitatively the same if the agent has a weaker or no present bias in the monetary domain. Moreover, we find that having a weaker or no present bias in the monetary domain even hurts a naive present-biased agent.
Let β be the discount factor for all future periods as before but now restricted to effort. Let $\beta' \in (\beta, 1]$ be that discount factor for money and the outside option. Morever, let b' and e' be the bonus and effort levels in these situations. By the same arguments as in our basic model, the equivalents of the contraints (rIRA) and (rC) hold with equality, i.e.,

$$b^{r\prime} = c(e^{r\prime}) + \overline{u} + \frac{\beta'\delta}{1-\delta}\overline{u} - \delta(\beta' b_1^{v\prime} - \beta c(e_1^{v\prime})) - \frac{\delta^2}{1-\delta}\left(\beta' b^{v\prime} - \beta c(e^{v\prime})\right) \tag{3.1}$$

and

$$b^{r\prime} = c(e^{r\prime}) + b_1^{v\prime} - c(e_1^{v\prime}) + \delta\beta'(c(e^{v\prime} - c(e_1^{v\prime}).$$

These equations together imply

$$b_1^{v\prime} - c(e_1^{v\prime}) = \overline{u} + \frac{\delta\beta}{1-\delta}c(e^{v\prime}) - \frac{\delta\beta'}{1-\delta}(b^{v\prime} - \overline{u}). \tag{3.2}$$

Unlike in our basic model, where solving for the optimal first-period virtual contract only pinned down the difference between b_1^v and $c(e_1^v)$, the optimal $b_1^{v\prime}$ and $c(e_1^{v\prime})$ are unique. As one can see from equation (3.1), the principal benefits from choosing them as high as possible. Consequently,

$$b_1^{v\prime} = \theta e_1^{v\prime} - \overline{\pi}.$$

Furthermore, as the maximal future per-period surplus from the agent's perspective is determined by maximizing $\beta'(\theta e - \overline{\pi} - \overline{u}) - \beta c(e)$, the optimal effort for the later periods of the virtual contract is higher than in our basic model, i.e., $e^{v\prime} > e^v$. Together with equation (3.2) we can conclude that

$$b^{r\prime} < b^r - (\beta' - \beta)\delta\left(\frac{\delta}{1-\delta}b^{v\prime} - \frac{1}{1-\delta}\overline{u} + b_1^{v\prime}\right) < b^r.$$

This implies that the principal can exploit a naive agent whose present bias is more intense in the effort domain than in the monetary domain more than an an agent whose present bias is equally intense in both domains. The reason for this is that an agent with a less severe present bias in the monetary domain discounts the high bonus of the future virtual contract less, which lets the future virtual contract appear even more attractive to him.

3.7.4 Moral Hazard

In this extension, we show that our main results do not depend on the assumption that contracts can be based on the agent's effort, but also hold under moral hazard. In this case, the naive agent is still mainly incentivized by rents provided by the virtual contract. However, the agent can only select the virtual contract if output has been high. Otherwise he must stick to the real contract for another period. In words, this means that only an employee who has been successful in doing his job has the option to take the next step and go for the virtual contract.

More precisely, consider the following adjustment to our model setting: Effort e_t is the agent's private information. Output remains verifiable, but now amounts to $y_t \in \{0, \theta\}$, with $\text{Prob}(y_t = \theta) = e_t$. Hence, output is not deterministic anymore but can be either high or low. Effort determines the probability with which high output is realized. Moreover, in order to guarantee an interior solution we assume $\lim_{e_t \to 1} c' = \infty$. This implies that first-best effort e^{FB} is still characterized by $\theta - c'(e^{FB}) = 0$. For concreteness we also set $\overline{u} = \overline{\pi} = 0$.

First, note that the optimal contract for a sophisticated agent is the same as for a time-consistent agent: In every period, the agent receives the bonus if output has been high. This bonus is set such that the agent chooses first-best effort, i.e., $b = \theta$. Furthermore, the wage is used to extract the generated rent, hence $w = \overline{u} + c(e^{FB}) - e^{FB}\theta$. Then, the principal collects the full surplus, and the optimal contract is stationary. This is different from the repeated moral hazard settings in Rogerson (1985) or Spear and Srivastava (1987), where the optimal contract contains memory. This is driven by the agent's risk aversion, though, which does not make it optimal to effectively sell the firm to the agent in every period (as is the case with a risk neutral agent).[16]

The structure of the optimal contract menu for the naive agent is the same as without moral hazard, with the following exceptions: First, the bonus is only paid if output has been high (note that because we can freely choose w^r, it is without loss of generality to set the bonus after a low output realization to zero). Second, starting with the real contract, the agent only has a choice to select the virtual contract in the next period if output has been high. This is because only in this case, the rents in the virtual contract provide incentives (recall that in the setting without moral hazard, the agent can choose the virtual contract in case he has exerted equilibrium effort, which however happens with certainty on the equilibrium path).

[16] Even with a risk neutral agent the optimal contract might contain memory, namely if the agent were protected by limited liability. Then, a profit-maximizing contract would provide incentives via a combination of bonus payments and on-path termination threats. See Fong and Li (2016), who derive a profit-maximizing dynamic contract in case output is not verifiable.

Once the agent has opted for the virtual contract, though, he can remain there independent of output realizations. This arrangement maximizes the agent's rent in the virtual contract and hence allows for the largest possible reduction of his real rent. Therefore, the naive agent's expected utility in a given period amounts to

$$U^{rv} = w^r - c(e^r) + e^r\left(b^r + \beta\delta U_1^v\right) + (1-e^r)\beta\delta\tilde{U}^{rv},$$

with $\tilde{U}^{rv} = w^r - c(e^r) + e^r\left(b^r + \delta U_1^v\right) + (1-e^r)\delta\tilde{U}^{rv}$. Hence, the agent takes into account that he might be stuck with the real contract in the subsequent period, namely if output has been low. Furthermore, the agent's virtual payoffs amount to $U_1^v = w_1^v - c(e_1^v) + e_1^v b_1^v + \delta U^v$ and $U^v = \left(w^v - c(e^v) + e^v b^v\right)/(1-\delta)$.

Furthermore, the principal now must use wage payments in the real and first period of the virtual contract to finetune the arrangement, i.e., extract real rents and in the end ensure that the agent does *not* choose the virtual contract. Without moral hazard, the principal could also set e_1^v in a way to prevent the agent from choosing the virtual contract. Now, effort is automatically pinned down by bonus payments and future rents. More precisely, the agent's real effort is given by

$$e^r \in \operatorname{argmax}\left[e^r b^r - c(e^r) + \beta\delta\left(e^r U_1^v + (1-e^r)\tilde{U}^{rv}\right)\right]. \tag{rIC}$$

Note that, although e^r also enters $\tilde{U}^{rv}$, its future values are taken as given by the agent when making his effort choice today. Therefore, e^r is characterized by $b^r - c'(e^r) + \beta\delta\left(U_1^v - \tilde{U}^{rv}\right) = 0$. Virtual effort levels are characterized by $b_1^v - c'(e_1^v) = 0$ and $b^v - c'(e^v) = 0$, respectively.

It is straightforward to show that the agent still receives the full virtual rent from the second virtual period on, and that it remains optimal to maximize this rent. Hence, $e^v = e^{FB}$ (i.e., $b^v = c'(e^{FB})$, and w^v is set such that $u^v = e^{FB}\theta - c(e^{FB})$). Furthermore, given U_1^v and $\tilde{U}^{rv}$, b^r is set such that $e^r = e^{FB}$ as well (since we do not impose a limited liability constraint in this section, the bonus b^r can potentially be negative). Still, the first virtual period is designed in a way that the agent does not select the virtual but sticks to the real contract.

As without moral hazard, it is optimal to have (rIRA) and (rC) constraints hold as equalities (we will check ex-post that the respective solution satisfies the (vC) constraint).

These constraints are $U^{rv} \geq 0$ (rIRA) and $U^{rv} \geq u_1^v + \beta\delta U^v$ (rC), where $u_1^v = w_1^v + e_1^v b_1^v - c(e_1^v)$, and can be rewritten as

$$\left(w^r - c(e^{FB}) + e^{FB}b^r\right)\left(\frac{1-\delta\left(1-\beta\right)\left(1-e^{FB}\right)}{1-\delta(1-e^{FB})}\right) + \frac{\beta\delta e^{FB}}{1-\delta(1-e^{FB})}U_1^v \geq 0 \tag{rIRA}$$

and

$$\left(w^r - c(e^{FB}) + e^{FB}b^r\right)\left(\frac{1-\delta\,(1-\beta)\,(1-e^{FB})}{1-\delta(1-e^{FB})}\right) + \frac{\beta\delta e^{FB}}{1-\delta(1-e^{FB})}U_1^v \geq u_1^v + \beta\delta U^v. \quad \text{(rC)}$$

Having both constraints bind yields

$u^r = -\beta\,(1-\beta)\,\delta^2 e^r \frac{\left(e^{FB}\theta - c(e^{FB})\right)}{(1-\delta)(1-\delta(1-\beta)(1-e^{FB}))}$ and $u_1^v = -\beta\delta\frac{\left(e^{FB}\theta - c(e^{FB})\right)}{(1-\delta)}$. For these values, the (vC) constraint holds as well.

Concluding, the agent still is exploited in case moral hazard is present, with (qualitatively) equivalent comparative statics. Furthermore, the first period of the virtual contract can be set such that the agent never selects it in the end.

3.7.5 Bargaining

So far, we have assumed that the principal has full bargaining power and can hence determine the terms of the employment relationship. Then the agent accepts any contract that at least gives him his outside option. However, in many labor market settings firms compete for agents. This will generally allow agents to extract a share of the relationship surplus. In this extension, we explore whether our results hold up if the principal does not have full bargaining power. We assume that players (potentially) bargain about how to allocate the relationship surplus at the beginning of every period. We do not explicitly model the bargaining process, but assume that players arrive at a Nash bargaining outcome where the principal keeps a share α and the agent a share $1-\alpha$ of the relationship surplus. More precisely, the agent accepts any offer that in expectation leaves him with at least $1-\alpha$ of the per-period relationship surplus. We use this approach and do not assume that players bargain about the total relationship surplus (like, for example, Ramey and Watson, 1997, or Miller and Watson, 2013) because the different effective discount factors also trigger different valuations of the future surplus stream.

We show that unless $\alpha = 0$, i.e., the agent has full bargaining power, the structure of the optimal menu of contracts for a naive agent remains unaffected. Therefore, competition on the labor market does not cause our results to disappear. This is only the case if competition for the agent's labor is perfect and frictionless. Naturally, the agent is better off for lower values of α, but is still exploited in the sense that his real share of the relationship surplus is lower than $1-\alpha$.

Importantly, because the principal can commit to long-term contracts, she is able to backload the agent's compensation and hence can commit *not* to renegotiate the virtual contract (different from e.g., Miller and Watson, 2013, and Fahn, 2016, where the principal's inability to commit not to renegotiate any agreement has substantial negative consequences on the efficiency of a long-term employment relationship). However, because the agent can *not* commit not to renegotiate an agreement, a front-loading of the agent's compensation is not feasible – although it might be optimal under some instances given the agent's lower effective discount factor.

Now, the (net) surplus in a given period t amounts to $e_t\theta - c(e_t) - \overline{u} - \overline{\pi}$. In order to characterize payments, we also have to specify what happens if bargaining fails. We assume that in this case, players do not enter an employment relationship in the given period and consume their respective outside options. Concluding, and already taking into account that the agent expects to select the virtual contract in future periods, the agent is willing to accept any offer that gives him at least an equivalent to $\overline{u} + (1-\alpha)\left(e^r\theta - c(e^r) - \overline{u} - \overline{\pi}\right) + \beta\delta\left\{\overline{u} + (1-\alpha)\left(e_1^v\theta - c(e_1^v) - \overline{u} - \overline{\pi}\right) + \delta\left[\overline{u} + (1-\alpha)\left(e^v\theta - c(e^v) - \overline{u} - \overline{\pi}\right)\right]/(1-\delta)\right\}$.

As before, the principal will optimally offer a contract menu that shifts a major part of the compensation into the virtual contract. Therefore, it remains optimal to promise the agent the full virtual surplus (from the second virtual period on), and consequently reduce real payments. Furthermore, $e^r = e^v = e^{FB}$. In addition, the agent's utility in the first virtual period, u_1^v, must be sufficiently low such that the agent will eventually not go for the virtual contract. There, we set $e_1^v = e^{FB}$, which is without loss of generality because what matters in the first period of the virtual contract is the value u_1^v, not the exact specifications of e_1^v and w_1^v.

Still, the relevant constraints which pin down payments and utilites are (rIRA) and (rC). We also have to take into account that the agent's real (IC) constraint looks different because once he deviates from selecting equilibrium effort, he cannot be punished by receiving his outside option from then on. Instead, he would be able to negotiate a new contract, expecting a share $1-\alpha$ of the relationhsip surplus. Therefore, when deviating, the agent only sacrifices this period's bonus payment and the option to select the virtual period in the subsequent period. However, it can be shown that the (rIC) constraint is automatically implied by the (rIRA) constraint.

This yields the following (relevant) constraints:

$$\begin{aligned}
&w^r + b^r - c(e^{FB}) + \beta\delta\left(w_1^v + b_1^v - c(e^{FB}) + \delta\frac{w^v + b^v - c(e^{FB})}{1-\delta}\right) \\
&\geq \overline{u} + (1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right) + \beta\delta\frac{\left[\overline{u} + (1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\right]}{1-\delta}
\end{aligned} \tag{rIRA}$$

and

$$\begin{aligned}
&w^r + b^r - c(e^{FB}) + \beta\delta\left(w_1^v + b_1^v - c(e^{FB}) + \delta\frac{w^v + b^v - c(e^{FB})}{1-\delta}\right) \\
&\geq w_1^v + b_1^v - c(e^{FB}) + \frac{\delta\beta}{1-\delta}\left(w^v + b^v - c(e^{FB})\right).
\end{aligned} \tag{rC}$$

Those will bind for the same reasons as above and pin down the values of the agent's real utility, as well as his utility in the first period of the virtual contract:

$$u^r - \overline{u} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left((1-\alpha) - \frac{\alpha\,(1-\beta)\,\beta\delta^2}{1-\delta}\right)$$

and

$$u_1^v - \overline{u} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left((1-\alpha) - \frac{\alpha\beta\delta}{1-\delta}\right).$$

Therefore, the agent receives less than his fair share, which would amount to $\overline{u} + (1-\alpha)\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)$ in every period. Furthermore, the agent's real payoff is decreasing in α. For $\alpha = 1$, the outcome is just like the one we derived in the main part. Only for $\alpha = 0$, the agent cannot be exploited. Then, $u^r = u_1^v = u^v$, and the distinction between real and virtual contract becomes immaterial. To conclude, the structure of the optimal menu of contracts is not affected by different distributions of bargaining power. Only if the agents have full bargaining power ($\alpha = 0$), a situation akin to frictionless competition on the labor market, agents are not exploited.

3.7.6 Partial Naivete

So far we only considered the extreme cases of completely naive and fully sophisticated agents. Now we relax this assumption and show that a partially naive agent receives exactly the same contract as a fully naive agent. A partially naive agent thinks that in any future period, he will discount the future with a factor $\hat{\beta} \in (\beta, 1)$. A fully sophisticated agent has $\hat{\beta} = \beta$, whereas a completely naive agent has $\hat{\beta} = 1$.

The principal's maximization problem is very similar to the problem when she faces a completely naive agent. (v1IC), (vIC), and (vC) are changed as these constraints involve the agent's expectations about his future self:

$$-c(e_1^v) + b_1^v + \hat{\beta}\frac{\delta}{1-\delta}\left(w^v + b^v - c(e^v) - \overline{u}\right) \geq 0, \tag{v1IC}$$

$$-c(e^v) + b^v + \hat{\beta}\frac{\delta}{1-\delta}\left(w^v - c(e^v) + b^v - \overline{u}\right) \geq 0, \tag{vIC}$$

and

$$\begin{aligned} &(w_1^v + b_1^v - c(e_1^v)) + \hat{\beta}\frac{\delta}{(1-\delta)}\left(w^v + b^v - c(e^v)\right) \\ \geq\ &(w^r + b^r - c(e^r)) + \hat{\beta}\delta\left[(w_1^v + b_1^v - c(e_1^v)) + \frac{\delta}{(1-\delta)}\left(w^v + b^v - c(e^v)\right)\right]. \end{aligned} \tag{vC}$$

The analysis is analogous to the one with the naive agent and by the same arguments we can omit several constraints. Thus we are left with the following simplified problem, maximizing

$$\Pi^r = \frac{e^r\theta - b^r}{1-\delta},$$

subject to

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right] \geq 0, \tag{rIR}$$

$$\begin{aligned} -c(e^r) + b^r - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right] \\ \geq (b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta\beta}{1-\delta}(b^v - c(e^v) - \overline{u}), \end{aligned} \tag{rC}$$

$$(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}(b^v - c(e^v) - \overline{u})\,\hat{\beta}\frac{(1-\delta)}{\left(1-\hat{\beta}\delta\right)} \geq \frac{(b^r - c(e^r) - \overline{u})}{\left(1-\hat{\beta}\delta\right)}, \tag{vC}$$

$$e^v\theta - b^v \geq \overline{\pi}. \tag{vIRP}$$

We show in Appendix C.1 that (vC) never binds, which implies that the principal optimally does not differentiate between a fully and a partially naive agent. The reason is exactly the same as in the case with a fully naive agent: Unless $\beta = \hat{\beta}$ or the agent is always kept at his outside option, (vC) and (rC) cannot bind simultaneously.

Proposition 3.5. *The partially naive agent receives the same contract as the fully naive agent.*

A (partially) naive agent thinks that in the future choosing the virtual contract is strictly better than choosing the real contract. In fact, tomorrow he will be just indifferent, though. For this decision, his true β is relevant but his belief $\hat{\beta}$ is not. Hence, for a given β, the principal will only offer two different sets of contracts to all agents: the virtual and the real contract for the naive agents, no matter whether they are fully or only partially naive, and the first-best contract for the sophisticated and the rational agents.[17] Such a discontinuity in the form of contracts is a common feature among papers that look at different degrees of sophistication, whether all but the fully sophisticated types receive exactly the same contract (as in this chapter; also see Heidhues and Kőszegi, 2010) or very similar contracts (e.g., DellaVigna and Malmendier, 2004) or rather naive types receive one contract and rather sophisticated types another (e.g., Eliaz and Spiegler, 2006).

[17]That is, the degree of naivete is not important. Yılmaz (2015) presents a parallel result, showing that degree of naivete does not matter and the principal has no exploitation rents, in a repeated moral hazard problem using artificial and effective effort levels.

3.7.7 Learning

In our main model, we assume that the agent fails over and over again: In every period he selects the real contract, although before he planned to take up the virtual contract in the current period. This assumption seems quite strong. One might expect that at least after a couple of failed attempts to actually select the virtual contract, the agent should realize that he has a present bias (for fully naive agents) or that his present bias is stronger than he thought (for partially naive agents). In the following, we show that even if the agent becomes aware of this, as long as he does not become fully sophsiticated, the principal will offer exactly the same contract.

Take any learning process characterized by a sequence of beliefs where the agent starts with a belief $\hat{\beta}_1$ about his present bias. Whenever he fails to accept the virtual contract, he learns that his beliefs $\hat{\beta}_{t-1}$ about his present bias must have been wrong and adjusts his belief to $\hat{\beta}_t \geq \beta$.

Corollary 3.3. *Consider an arbitrary learning process in which $\hat{\beta}_t > \beta$ for all t, i.e., the agent never learns his true β. Then learning about the present bias does not affect the optimally offered contracts.*

This immediately follows from Proposition 3.5.

If the agent has the chance to learn his true β, he will choose the first-best contract from then on, which can always be added to the menu of contracts by the principal, but will only be chosen by a sophisticated agent.

Corollary 3.4. *If the learning process allows the agent to learn his true β (or leads to belief $\hat{\beta} < \beta$) with a positive probability, then it is optimal to add the first-best contract to the menu of contracts.*

This result is straightforward: When adding the first-best contract to the menu of contracts, an agent who has become sophisticated will from then on choose the first-best contract. A (fully or partially) naive agent does not have an incentive to select the first-best contract as he is indifferent between the first-best contract and the contract intended for him.

3.7.8 Unobservable Agent Types

So far we have assumed that the principal can tailor her offer to the agent's type. However, in real world contexts it is not clear whether the principal necessarily knows an agent's type. We show that our results are robust to considering the case of unobservable agent types: Agents are optimally separated, and each receives a menu tailored to his type. Moreover, menus generally are still exploitative in the same vein as before, only the extent of exploitation is reduced for some, in order to prevent one type from mimicking another one. Intuitively, the principal will focus on exploiting types she is more likely to face, and hence induces a separation by reducing the exploitative rents she collects from less frequent types.

Formally, assume there are two types of agents, $i \in \{1, 2\}$ with different values β_i. For simplicity, we set $\hat{\beta}_1 = \hat{\beta}_2 = 1$. Relaxing this would not affect the results. Without loss of generality, assume $\beta_1 < \beta_2 \leq 1$. This allows the principal to screen agents as they value the future differentially despite having the same belief about their future β. Let s_1 be the share of agents in the population with β_1, and $1 - s_1$ the share of agents with β_2. The principal cannot observe the agent's type, but only knows the distribution of types. As we assume that the principal can fully commit to the long-term contract contingent on the history of the game, she can preclude an agent, after the initial contract choice, to switch from one contract to the contract intended for the other type.[18] Hence, if different menus are offered, we need to make sure that it is optimal for agent i to select the respective contract intended for type i, which we denote $\tilde{\boldsymbol{C}}_i$.

Now, let $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$ be the profit-maximizing contracts derived in Subsection 3.6.2, with the slight modification that the virtual contract is only offered from $t = 1$ on, i.e., the agent can only choose the real contract in $t = 0$.[19] If these menus were also offered for unobservable types, agent 2 would actually go for $\boldsymbol{C}_2$. This is because agent 1 is just indifferent between taking up the contract and choosing his outside option. As $\beta_1 < \beta_2$, agent 2 values the future benefits of the virtual contract more than agent 1 and hence expects to receive a higher utility from $\boldsymbol{C}_1$. Furthermore, since $\boldsymbol{C}_1$ is designed in such a way that agent 1 is just indifferent between the virtual contract and the real contract, agent 2 would actually choose the virtual contract after selecting $\boldsymbol{C}_1$.

Hence, either $\tilde{\boldsymbol{C}}_2$ must be constructed in a way that makes it optimal for agent 2 to choose it (keeping $\tilde{\boldsymbol{C}}_1 = \boldsymbol{C}_1$), or $\tilde{\boldsymbol{C}}_1$ must be made sufficiently unattractive for 2 (keeping $\tilde{\boldsymbol{C}}_2 = \boldsymbol{C}_2$). How exactly the menus are adjusted is described in the Appendix, the main result is given in Proposition 3.6. We say that an agent is fully exploited when he receives the same utility as in our main analysis with one type of naive present-biased agent. We say that an agent is not exploited when he receives the utility of his outside option in every period. We say that an agent is partially exploited if an agent receives a utility in between.

Proposition 3.6. *For all $\beta_1, \beta_2, \delta > 0$, with $\beta_1 < \beta_2 < 1$, there exists a threshold $\underline{s} \in \left(0, \frac{\beta_2 - \beta_1}{\beta_2}\right)$ such that for all $s_1 \leq \underline{s}$ it is optimal to offer a menu of contracts such that agent 1 is not exploited and agent 2 is fully exploited.*
Furthermore, it is optimal to offer two different contracts to the agents which both exploit the agents, but only partially for all $s_1 \in [\underline{s}, 1 - \beta_1]$. For all $s_1 \geq 1 - \beta_1$, it is optimal to fully exploit agent 1 and to partially exploit agent 2.

[18] Note that it turns out to be in fact optimal for the principal to preclude these switches. If it were optimal to let the agent switch the principal could have just amended the original contract by the respective components of the alternative contract.

[19] This does not change anything for agent 1, but makes $\boldsymbol{C}_1$ less attractive for agent 2.

Therefore, one type is still exploited in exactly the same way as with observable types, the other type is less exploited in order to induce a separation (note that we also show in the Appendix that a separation is always strictly optimal for $\beta_2 < 1$. If $\beta_2 = 1$ and $s_1 > 1 - \beta_1$, the principal is indifferent between inducing a separation or just letting agent 2 select $\tilde{\boldsymbol{C}}_1 = \boldsymbol{C}_1$ who would then go for the virtual contract. If writing different menus was associated with some small costs, we would predict only one menu, with time-consistent agents actually "making a career"). Therefore, our results are qualitatively unaffected if types cannot be observed, only the extent of exploitation has to be reduced for one type.[20]

3.8 Conclusion

In this chapter we have shown how a principal should optimally contract with a present–biased agent in a long-term relationship absent moral hazard. The principal can take advantage of a naive present–biased agent and push him even below his outside option by offering a menu of contracts: a virtual contract which consists of a relatively low compensation in the initial period but promises high future benefits and a real contract which keeps the agent below his outside option. The agent expects that he will choose the virtual contract from the next period on and therefore accepts a lower compensation today. However, he always ends up choosing the real contract and hence never actually gets to enjoy the generous benefits from the virtual contract. These findings are robust to imposing limited liability on the agent, taking moral hazard into account, giving the agent some bargaining power, letting the agent have differently strong present biases in different domains, considering the case of partial naivete, allowing the agent to learn about his bias, and considering a finite time horizon.

A number of our results appear at odds with usual findings or basic intuition in models with time-inconsistent agents. First, in our model, the time-inconsistent agent might work harder than his time-consistent counterpart, while in many other settings previously studied in the literature he appears more lethargic and lazy. This is driven by the limited liability constraint. As the agent cannot receive a negative compensation, the only way to exploit him further is to let him work inefficiently hard. Note, however, that the agent still procrastinates as he indefinitely postpones the investment to take up the virtual contract and actually "make a career".

Second, a higher outside option for the principal actually benefits the agent, a result contrary to relational contracting or Nash-bargaining intuitions. The reason for this is that a higher outside option makes it harder for the principal to credibly commit to employing the agent in the future and providing him with generous benefits. Therefore, the promised

[20] Other papers analyze screening time-inconsistent agents: First, Heidhues and Kőszegi (2010)'s results are completely unchanged, but this relies on a condition that is violated in our framework. Second, Yan (2011) does not look at naive types with differently strong present biases. Finally, Galperti (2015) has similar trade-offs but a different form of bias.

future benefit must be lower, making the virtual contract appear less desirable and thus the agent can be exploited less today.

Third, our results suggest that employment protection regulations might in fact hurt agents. Inverting the preceding argument, employment protection increases the principals commitment to future employment, allowing her to credibly promise a higher compensation in the future – which she will never have to pay as the agent will always choose the real over the virtual contract. Hence, the agent is willing to accept a lower payment or work harder today.

Chapter 4

The Impact of Social Media on Belief Formation

4.1 Introduction

Social media are becoming an increasingly important part of many people's lives and our society. The number of Facebook's active users worldwide rose from 431 million at the beginning of 2010 to 1.71 billion in the second quarter of 2016.[1] On average, users spend more than 50 minutes per day on Facebook, Instagram, and Facebook messenger.[2] This is changing the way people communicate, how they acquire information, and how they form beliefs.[3] Social media advertisement revenue in the U.S. is estimated to exceed twelve billion USD, and the traffic to news sites from Facebook surpassed the traffic from Google in early 2015, while it had been less than half of the traffic from Google as late as 2013.[4] In fact, 48% of Americans with Internet access received news about politics from Facebook in 2014.[5] At the same time, research has found that people tend to communicate with similar people (McPherson, Smith-Lovin and Cook 2001, Gentzkow and Shapiro 2011). Moreover, the algorithms of some social media platforms make it easier to see what some individuals post and filter out other content. Sunstein (2001, 2009), Van Alstyne and Brynjolfsson (2005),

---

[1] See `http://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/` (accessed Sept.1, 2016).

[2] See `http://www.nytimes.com/2016/05/06/business/facebook-bends-the-rules-of-audience-engagement-to-its-advantage.html` (accessed Sept.1, 2016).

[3] See e.g., Bailey, Cao, Kuchler and Stroebel (2016) for geographically-distant friends' influence on real estate purchasing behavior.

[4] See `http://www2.biakelsey.com/webinars/biakelsey-u-s-local-advertising-forecast-for-2016-key-findings.pdf` and `http://fortune.com/2015/08/18/facebook-google/` (both accessed Sept.1, 2016); media outlets such as Wired, The Atlantic, Reuters, and The Daily Telegraph belong to Parse.ly's clients.

[5] See `http://www.journalism.org/2014/10/21/section-2-social-media-political-news-and-ideology` (accessed Oct.2, 2016).

and Pariser (2011) argue that social media and the Internet in general allow people to mostly communicate with similar individuals, which would expose these people to only certain kinds of information and may lead to increased segregation and polarization. Indeed, Americans disagree about Obama's religion, the existence of weapons of mass destruction in pre-2003 Iraq, and the reality of global warming.[6]

To answer the question of how social media impact belief formation, this chapter combines a model of strategic network formation with learning: Agents have different types, live in different regions and form links to acquire additional information from other agents in order to make a decision in the near future. The cost of links to other agents depends on two dimensions: The costs are lower if (i) the agents have more similar types and (ii) if the agents live in the same region. If they do not live in the same region, agents can form inter-regional links, for example by interacting online. While the links are formed for decisions in the near future, information that agents gather through these links and over time also more and more through indirect links influences agents' opinions in the long-run as well. This allows me to analyze the opinion formation in the long-run. I show that the extent and the speed of information aggregation for opinion formation follows an inverse-U-shape with respect to the cost of inter-regional links, i.e., up to a certain point, the increasing importance of social media leads to more and faster information, but beyond that point it leads to less and slower information aggregation. This implies that the growing influence of social media and progress in information and communication technology in general decreases disagreement in a society first, but facilitates disagreement if it grows even further. Social algorithms like Facebook's Top Stories accelerate and amplify these effects. At the same time, social media make people more likely to connect with similar people and let them disagree with different people for longer (or possibly forever). This may increase tensions between different parts of a society and prevent people from learning from each other. People who have more incentives to acquire information or who find it more difficult to find similar people in their own region form a higher fraction of links via social media. In addition, I explain how social media can help previously isolated people to form links with other people, but also how social media may make them more likely to hold extreme beliefs.

By incorporating correlation neglect (i.e., the unawareness that information may be correlated) in a simple portable way, I demonstrate how social media can hurt the decisions

[6] See `https://d3n8a8pro7vhmx.cloudfront.net/aai/pages/11126/attachments/original/1450651184/2015_American_Attitudes_Toward_Arabs_and_Muslims.pdf?1450651184`, `http://view2.fdu.edu/publicmind/2015/150107/`, and `www.eenews.net/assets/2016/04/27/document_cw_01.pdf` (all accessed Sept. 2nd, 2016).

of people with correlation neglect and how social media can lead to more disagreement within a country, but less disagreement across countries. While rational people always benefit from a decrease in inter-regional linking cost, people with correlation neglect might be worse off.

The intuition for why disagreement in a society is U-shaped with respect to the importance of social media is as follows. If inter-regional communication becomes less costly, agents will form links to agents of the same type in other regions and will form fewer links within their own region. This has two effects: First, more links to other regions allow an agent to observe signals in other regions earlier and therefore reduce the heterogeneity of beliefs across regions. Second, fewer links within an agent's own region mean that the agent forms fewer links to different types. So it takes the agent longer to observe the signals of very different types and therefore the heterogeneity of beliefs across types increases. At a high inter-regional linking cost, when not many links to agents in other regions are formed, the first effect dominates and disagreement decreases. Only when there exist many links across regions and the cost of interacting with agents in other regions is further reduced, the second effect dominates and leads to more disagreement. Along with the disagreement, the inaccuracy of opinions in the long-run follows a U-shape with respect to the inter-regional linking cost. Interestingly, disagreement arises despite rational updating and ex ante symmetry with respect to agents' beliefs and their expected signals.

Social algorithms like Facebook's Top Stories accelerate and amplify the above effects. These algorithms make information of similar types more visible and information of different types (whom people might know from their region but also communicate with via Facebook) less visible. This is equivalent to further reducing the inter-regional linking cost and simultaneously increasing the cost of forming links to different types. As a consequence, agents form more inter-regional and fewer domestic, i.e., intra-regional, links. This reduces disagreement and makes opinions more accurate if only few inter-regional links were formed without social algorithms, but intensifies disagreement and inaccuracy of opinions if agents obtain most of their information via similar types even without social algorithms. Similarly, when disagreement between types leads to a higher cost of communicating with other types, a vicious circle may emerge that leads to more segregation and disagreement: More disagreement leads to a higher cost of inter-type communication, and this in turn leads to more disagreement. Indeed, evidence suggests that communication costs with other types might have increased: According to the Pew Research Center (2014), the fractions of Democrats and Republicans who see the other party as "a threat to the nation's well-being" have increased to 27% and

36%, respectively, and 15% of Democrats and 17% of Republicans would be unhappy if a family member married a person of the other party.

Sorting by types increases with decreasing inter-regional linking cost, because agents form more inter-regional links to agents of their own type and fewer links with other types. Then, agents observe signals from similar types more often and signals from different types less often. If types reflect different ideologies and agents receive type-specific signals, then agents' beliefs become more consistent with an ideology, which is what the Pew Research Center (2014) finds: The share of consistent liberals rose from 5% in 1994 to 23% in 2014, and the share of consistent conservatives rose from 13% to 20%.

For agents with high stakes such as professionals sorting is more pronounced, because they have higher incentives to acquire information and form more links.[7] At least some of these additional links are likely to be inter-regional, because linking to different types becomes increasingly expensive for more distant types, while inter-regional links are available at a uniform cost. Similarly, agents who are very different from most other agents make use of inter-regional links at a higher inter-regional linking cost than agents with moderate types and they form a higher fraction of inter-regional links. The reason is that for agents with very different types, it is more expensive to talk to agents in their own region, because the types of those agents are further away on average. As a real-world example, 47% of consistently conservatives and 32% of consistently liberals say that posts about politics on Facebook are mostly or always in line with their own views, while among all respondents only 23% say so.[8] Relatedly, agents with extreme types might form no links at all, because the number of domestic links required to potentially change their decision is too expensive. When the inter-regional linking cost is low enough for them to form links, their decision-making is better most of the time. However, there is a chance that they observe a higher number of wrong signals, which leads to confident wrong beliefs and thus makes them more likely to take extreme actions, with potentially negative consequences for other people.[9]

---

[7] See e.g., Poole and Rosenthal (1984, 1985, 1997), Layman and Carsey (2002), Harbridge and Malhotra (2011), Noel (2014) and McCarty, Poole and Rosenthal (2016) for evidence of the polarization of and sorting in Congress. Furthermore, Rosenblat and Möbius (2004) find that the collaboration between distant similar economists has increased from the 80s to the 90s, while the collaboration between close dissimilar economists has decreased.

[8] See `http://www.journalism.org/2014/10/21/political-polarization-media-habits` (accessed Oct.2, 2016).

[9] For example, in many past mass shooting incidents, the perpetrators radicalized themselves by interacting with similar people in online forums, and ISIS is well-known for recruiting in Western societies online. Social media are not the only reason for this development, of course, but they probably are a promoting factor.

An extension of the model with correlation neglect predicts that social media can lead to an increased disagreement within countries, but reduced disagreement across countries. If people watch a lot of news on national TV, the correlation of information in a given city is probably lower than the correlation of information in a political movement, i.e., there are fewer region-specific signals than type-specific signals, and so increased segregation by types will lead to more disagreement. At the same time, the correlation of information in a given country may be higher than the correlation of information of people with a certain political preference, so reduced segregation by regions will lead to less disagreement. Another prediction is that when type-specific correlation is strong, agents with correlation neglect can be worse off when the inter-regional link formation cost decreases, while rational agents are better off. As a consequence, the difference in utilities between people with correlation neglect and rational people increases.

Long-run beliefs play an important role for voting: If the society does not aggregate all available information, it is more likely to make suboptimal choices. Even though forming additional links is too costly for an individual because each individual is unlikely to be pivotal, the aggregate effect of incomplete social learning and disagreement for the society might be large. Beyond politics and voting behavior, social learning and disagreement matter in a variety of other settings. First, if agents hold very different beliefs about the value of an asset, they may engage in belief-neutral inefficient trading according to Brunnermeier, Simsek and Xiong (2014). Moreover, in such a setting the realization of an outcome may lead to more jumps and therefore a higher volatility. Second, if doctors with different specializations fail to aggregate information, it might take them longer to learn about an adequate medical treatment or they might even never learn about it (Eyster and Rabin 2014). Third, different agents may not adopt new products or technologies such as fertilizer (Rosenblat and Möbius 2004) for a longer time when they do not learn about all the available signals. This may also affect marketing strategies and distort incentives for innovation.

The remainder of the paper is organized as follows. Section 4.2 relates the paper to the literature. Section 4.3 introduces the model. Section 4.4 derives and discusses the main results. Section 4.5 presents some extensions, including the case of agents with correlation neglect. Section 4.6 concludes. All proofs can be found in the appendix.

4.2 Related Literature

This chapter contributes to five strands of literature. First, this chapter extends a new and still small literature which combines learning with endogenous network formation.[10] The most closely-related paper from a technical perspective is Acemoglu et al. (2014). In their paper, which focuses on the importance of so-called information hubs (some highly connected agents) for social learning, agents start with free links to other agents in exogenously given groups and linking to other agents is possible at a constant cost. By contrast, in this chapter, there are no exogenous groups, and linking costs have more than just two different levels. As a consequence, the arising network structures are different. Song (2016) looks at a sequential decision process in which agents can decide to observe some of their predecessors' actions and finds that agents' influence must be infinitesimally influential for maximal learning to occur. In this chapter agents can observe signals rather than actions and all agents decide simultaneously. Tan (2015) also combines opinion formation with endogenous network formation and, furthermore, a preference of people to talk to people with the same beliefs. She finds that there is a steady state with either consensus or extreme polarization. In contrast to this chapter, however, she looks at a very different network formation process, in which in every period some randomly selected agent can delete or form a certain link. Moreover, there is no learning but opinions are chosen. None of the above papers looks at a gradual change in inter-regional linking cost, which is the focus of this chapter. Further, none of them analyzes a framework in which agents have different types and live in different regions.

Second, this chapter adds to the literature which investigates the consequences of progress in inter-regional communication technologies. Rosenblat and Möbius (2004) find that such technological progress always decreases the separation between two individuals when the number of agents is large (which is equivalent to increasing disagreement in this chapter), while Van Alstyne and Brynjolfsson (2005) find that technological progress always leads to more separation (unless agents have a preference for diversity). By contrast, this chapter finds a U-shaped relationship between technological progress and the degree of separation; while Rosenblat and Möbius (2004) do not find this result because of the combination of

[10] There are several papers that look at either just learning in networks or just network formation. For learning in networks, see e.g., DeGroot (1974), Bala and Goyal (1998), DeMarzo, Vayanos and Zwiebel (2003), Gale and Kariv (2003), Golub and Jackson (2010), Mueller-Frank (2011), Acemoglu, Dahleh, Lobel and Ozdaglar (2011), Golub and Jackson (2012), Jadbabaie, Molavi, Sandroni and Tahbaz-Salehi (2012), Jadbabaie, Molavi and Tahbaz-Salehi (2013), and Mossel, Sly and Tamuz (2015). For network formation, see e.g., Jackson and Wolinsky (1996), Bala and Goyal (2000), Watts (2001), Jackson and Watts (2002), Jackson and Rogers (2005), Hojman and Szeidl (2008), and Galeotti and Goyal (2010).

randomness and only two types in their model, Van Alstyne and Brynjolfsson (2005) do not look at long-run effects.

Third, this chapter provides another explanation of how agents can become more polarized, hold inaccurate beliefs, or agree to disagree (Aumann 1976). The explanations in the literature include different priors and bimodal preferences (Dixit and Weibull 2007), different priors and uncertain signal interpretations (Glaeser and Sunstein 2013, Acemoglu, Chernozhukov and Yildiz 2016), multi-dimensional beliefs and private information (Andreoni and Mylovanov 2012, Glaeser and Sunstein 2013), confirmatory bias (Rabin and Schrag 1999), correlation or selection neglect (Glaeser and Sunstein 2013), wishful thinking (Bénabou 2013, Le Yaouang 2016), learning with a misspecified model of the world (Bohren 2016, Heidhues, Kőszegi and Strack 2016a), not taking into account strategic behavior of others (Eyster and Rabin 2010, Glaeser and Sunstein 2013, Gagnon-Bartsch and Rabin 2016), and reputation-concerned journalists (Shapiro 2015). While all of those explanations might be factors that play a role for forming more extreme beliefs, this chapter focuses on rational updating and updating with correlation neglect in a setting of common priors and unidimensional beliefs. It investigates the change in polarization due to technological progress.

Fourth, this work speaks to an empirical literature that looks at the polarization of voters. So far, the evidence for increased polarization beyond sorting is mixed at best. (See for example Fiorina, Abrams and Pope (2005), Abramowitz and Saunders (2008), Hetherington and Weiler (2009), Gentzkow and Shapiro (2011), Falck, Gold and Heblich (2014), Lelkes, Sood and Iyengar (forthcoming), Fiorina (2016), and Poy and Schüller (2016).) This chapter offers an explanation for why it might take a while before significanty increased polarization can be observed.

Fifth, this chapter contributes to the literature of modeling correlation neglect and selection neglect and examining the consequences and implications.[11] Implicitly, all papers using models that build on DeGroot (1974) and DeMarzo et al. (2003) incorporate correlation neglect. More explicitly, correlation neglect is modeled by a series of papers by Levy and Razin.[12] The most closely related one of those papers in terms of application is Levy and Razin (2015b). It finds that correlation neglect may help to aggregate information in elections, because voters are more likely to vote according to their information rather than their political preference. Ortoleva and Snowberg (2015) show that correlation neglect leads

[11] For experimental evidence on correlation neglect and selection neglect, see Eyster and Weizsäcker (2011), Enke and Zimmermann (2015), and Enke (2015). For evidence that people may avoid double-counting in other situations, see Möbius, Phan and Szeidl (2015).

[12] Levy and Razin (2015a, 2015b, 2015c, 2016a, 2016b).

to overconfidence, more extreme beliefs, and a higher voter turnout. In this chapter, correlation neglect also leads to overconfidence but is never beneficial. Similar to those papers, correlation neglect in this chapter is modeled in such a way that the belief of agents with correlation neglect coincides with the rational belief when correlation is absent. In this chapter correlation neglect is modeled in an even simpler way and therefore more portable.

4.3 Model

This section presents the basic model with rational agents and explains how agents form links, what signals they observe, and how they make their decisions.

There are $n \in \mathbb{N}$ regions with $m \in \mathbb{N}$ different agents each. Let agent θk be an agent of type $\theta \in M = \{1, ..., m\}$ in region $k \in N = \{1, ..., n\}$. I assume that both the number of types m and the number of regions n are large enough so that agents can always form another link to an agent in the same region or to an agent with the same type. There are $\bar{t} + 1$ periods: $t = 0, 1, ..., \bar{t}$, where $\bar{t} \geq 2$.

Decision problem. There is an unknown true state $r \in \{0, 1\}$, where $Prob(r = 1) = Prob(r = 0) = 0.5$. At the end of period $t = 1$, each agent decides between $d_{\theta k} = 1$ and $d_{\theta k} = 0$ and tries to match the state. She receives a gross utility of $\overline{u}_{1\theta}$ if $d_{\theta k} = r = 1$, $\underline{u}_{1\theta}$ if $d_{\theta k} = 1 \neq r$, $\overline{u}_{0\theta}$ if $d_{\theta k} = r = 0$, and $\underline{u}_{0\theta}$ if $d_{\theta k} = 0 \neq r$, where $\overline{u}_{i\theta} > \underline{u}_{j\theta}$ for all $i \neq j \in \{0; 1\}$. For example, an agent optimizes her links for decision problems in which her utility depends on whether something is good or popular. More concretely, she might wonder whether to watch a movie. If the movie is going to win an Academy Award ($r = 1$), an agent would like to watch the movie ($d_{\theta k} = 1$), while she would rather not watch the movie ($d_{\theta k} = 0$) if it is not going to win an Academy Award ($r = 0$).

Link formation. At the beginning of $t = 0$ each agent can form directed links to other agents.[13] These links are long-run relationships and cannot be changed later. It is more costly for an agent to form a link to an agent who does not live in the same region and it is the costlier the further away that agent's type (e.g., ideology or income) is. If agent θk wants to form a link to agent ηl, this costs her $c_T|\theta - \eta| + c_R \mathbb{1}_{l \neq k}$, where $c_T > 0$ is the type-specific linking cost and $c_R > 0$ is the inter-regional linking cost. I assume that an agent rather forms

[13] A link is *directed* if θk could have no link to ηl, even if ηl had a link to θk. One interpretation would be that ηl pays attention to θk, but θk does not pay attention to ηl.

an inter-regional link to a close-by region than to a region further away, potentially because of cultural or language differences.[14] This is a conservative assumption, because a link to a region further away would allow information to spread even faster.

The network consisting of agents and links can be represented by an *adjacency matrix* A, of which entry $a_{\theta k,\eta l} = 1$ if there is an (unweighted) link from agent θk to agent ηl and 0 otherwise. The total cost of link formation for agent θk is the sum of the cost of all her links, $C(\mathbf{a_{\theta k}}) := \sum_{\eta \mathbf{l} \in \mathbf{M} \times \mathbf{N} : \mathbf{a_{\theta k},_{\eta l}} = \mathbf{1}} (\mathbf{c_T} |\theta - \eta| + \mathbf{c_R} \mathbb{1}_{\mathbf{l} \neq \mathbf{k}})$, where $\mathbf{a_{\theta k}}$ is row θk of the adjacency matrix $A \in \{0;1\}^{m \cdot n \times m \cdot n}$. I interpret the increasing importance and availability of social media in people's lives as a decrease in the inter-regional cost.

Signals and beliefs. In $t = 0$, after having formed links, each agent observes a set of signals, $S^1_{\theta k}$. This set consists of up to three different types of signals: $\overline{s}_U \geq 1$ signals uniquely obtained by one agent, $\overline{s}_T \geq 0$ type-specific signals, and $\overline{s}_R \geq 0$ region-specific signals. Each $s^i_{\theta k} \in S^1_{\theta k}$, $i \in \{1, ..., \overline{s}_U + \overline{s}_T + \overline{s}_R\}$, is either 0 or 1 and equal to r with probability $\alpha \in (0.5, 1)$. Sometimes I will refer to the signals as *high* and *low*. Unless noted otherwise, signals are (conditional on the true state) independent and identically distributed. Signals uniquely obtained by one agent are always (conditional on the true state) independent and identically distributed. Every agent obtains some information no one else obtains by herself. Type-specific signals are the same for all agents of the same type, i.e., $s^i_{\theta k} = s^i_{\theta l}$ for all $i \in \{\overline{s}_U + 1, ..., \overline{s}_U + \overline{s}_T\}$. For example, agents of the same type might watch the same TV programs. Similarly, region-specific signals are the same for all agents in the same region, i.e., $s^i_{\theta k} = s^i_{\eta k}$ for all $i \in \{\overline{s}_U + \overline{s}_T + 1, ..., \overline{s}_U + \overline{s}_T + \overline{s}_R\}$. For example, agents in the same region might read the same local newspaper.

In $t = 1$, an agent also observes the signals of the agents she has formed links to. For example, if an agent has formed a link to an agent with the same type but in a different region, she will observe $\overline{s}_T$ independent type-specific signals (because the other agent obtains the same type-specific signals), $2\overline{s}_R$ region-specific signals (because the other agent obtains different region-specific signals), and $2\overline{s}_U$ uniquely obtained signals. These are in total $\overline{s}_T + 2\overline{s}_R + 2\overline{s}_U$ (conditional on the true state) independent and identically distributed signals.

Opinions. An agent also uses the links she formed to make an optimal decision in $t = 1$ to form an opinion about other questions that barely affect her utility. Exactly as in her decision problem above, there is a true state $r \in \{0, 1\}$, where $Prob(r = 1) = Prob(r = 0) = 0.5$, each

[14] One could model this as an extra cost of $\varepsilon|k - l|$ per link, where $\varepsilon > 0$ is very small.

agent observes $\overline{s}_U + \overline{s}_T + \overline{s}_R$ that can be uniquely obtained by one agent, type-, or region-specific in $t = 0$, and in $t = 1$, she also observes the signals of the agents she has formed links to. In addition to that, in period $t > 1$, an agent also observes the signals the agents she has formed links to have observed in the previous period $t-1$, i.e., the signals of her contacts at a geodesic distance of t.[15] She then forms her beliefs in a Bayesian way. Consequently, opinion formation differs from the belief formation regarding her decision problem in the following two ways. First, opinion formation is an ongoing process that can take place for a longer time and also takes into account the information of indirect contacts. Second, her links are not formed in order to hold accurate opinions, e.g., because the utility differences between different opinions are far less important than the utility differences that are possible for her decision problems or because she does not take into account the indirect effects of links when forming links. I prefer the first interpretation because it allows for the agents' behavior to be considered rational. This means that I will investigate opinion formation as a byproduct of links that were formed for information acquisition.[16]

Let $S^t_{\theta k}$ be the set of signals agent θk has observed up until and including period t. For example, $S^1_{\theta k} = \{s^i_{\theta k} : i \in \{1, ..., \overline{s}_U+\overline{s}_T+\overline{s}_R\}\} \cup \{s^i_{\eta l} : a_{\theta k,\eta l} = 1 \text{ and } i \in \{1, ..., \overline{s}_U+\overline{s}_T+\overline{s}_R\}\} = \bigcup_{\eta l : a_{\theta k,\eta l}=1} S^0_{\eta l} \cup S^1_{\theta k}$. Furthermore, let $\underline{S}^t_{\theta k}$ be the set of conditionally independent signals agent θk has observed up until and including period t.

Timing. In $t = 0$, agents form links and then receive a signal. In $t = 1$, agents observe the signals of the agents they have formed links to and then make a decision. In each $t > 1$, agents observe the signals of agents at a geodesic distance of t, which allows us to look at how agents' opinions develop.

Formal maximization problem. In $t = 1$ each agent maximizes her expected utility from the decision given the signals she observed:

$$\mathbb{E}[u_\theta(r, d_{\theta k}(S^1_{\theta k}))] := \tag{4.1}$$

[15] A *path* from agent θk to agent ηl is a sequence of links that connects θk to ηl, i.e., such that $a_{\theta k,\sigma(11)} = a_{\sigma(11),\sigma(12)} = a_{\sigma(12),\sigma(13)} = ... = a_{\sigma(\zeta i),\eta l} = 1$ for some permutation σ on the set of agents in the network, where ζi is an agent in the network. The *shortest path* is the path with the minimum number of such links. The *geodesic distance* is the length of the shortest path.

[16] Similarly, Allen and Gale (2000) look at contagion in a network that was formed for risk-sharing.

$$\begin{aligned}\max_{d_{\theta k}(S^1_{\theta k})\in[0,1]} d_{\theta k} \quad & \left[Prob(r=1|S^1_{\theta k})\overline{u}_{1\theta} + (1-Prob(r=1|S^1_{\theta k}))\underline{u}_{1\theta}\right] \\ -(d_{\theta k}-1) \quad & \left[Prob(r=1|S^1_{\theta k})\underline{u}_{0\theta} + (1-Prob(r=1|S^1_{\theta k}))\overline{u}_{0\theta}\right].\end{aligned}$$

Obviously, if the expression following $d_{\theta k}$ is strictly greater than the expression following $-(d_{\theta k}-1)$, it is optimal for the agent to choose $d_{\theta k}=1$, and $d_{\theta k}=0$ if it is strictly smaller. In case the two expressions are equal, i.e., the agent is indifferent, I assume that she chooses each alternative with probability one half. For the sake of brevity, I will not state that the agent mixes with equal probability for each alternative in lemmas, propositions, nor corollaries. Her (unconditional) *expected utility from the decision* is equal to $\mathbb{E}[\mathbb{E}[u_\theta(r, d_{\theta k}(S^1_{\theta k}))]]$. In $t=0$, each agent forms links that maximize her expected utility which consists of the expected utility from the decision minus the cost for the links:

$$\max_{\mathbf{a}_{\theta k}\in\{0,1\}^{mn}:a_{\theta k\theta k}=0} \mathbb{E}[U_{\theta k}] := \max_{\mathbf{a}_{\theta k}\in\{0,1\}^{mn}:a_{\theta k\theta k}=0} \mathbb{E}[\mathbb{E}[u_\theta(r, d_{\theta k}(S^1_{\theta k}))]] - C(\mathbf{a}_{\theta k}). \tag{4.2}$$

4.4 Analysis

This section starts by calculating how agents decide in $t=1$ and then what links they form in $t=0$. After that it will look at the consequences for the resulting decisions in $t=1$. Finally, this section will show what opinions will be formed.

4.4.1 Belief Formation and Decision

When an agent observed i high and j low conditionally independent signals, her belief that $r=1$ is equal to $\frac{\binom{i+j}{i}\alpha^i(1-\alpha)^j}{\binom{i+j}{i}\alpha^i(1-\alpha)^j+\binom{i+j}{i}(1-\alpha)^i\alpha^j} = \frac{\alpha^{i-j}}{\alpha^{i-j}+(1-\alpha)^{i-j}}$. Because this expression only depends on the difference between high and low conditionally independent signals for a given α, this difference is a sufficient statistic for the agent's belief. Let $z^t_{\theta k}$ be the difference between high and low conditionally independent signals agent θk observed up until period t.

Lemma 4.1. *(i) $z^t_{\theta k}$ is a sufficient statistic for agent θk's belief in period t.*
(ii) In period $t=1$ the agent chooses $d_{\theta k}=1$ if and only if $z^1_{\theta k} > z^*_\theta := \frac{\log\left(\frac{\overline{u}_{1\theta}-\underline{u}_{0\theta}}{\overline{u}_{0\theta}-\underline{u}_{1\theta}}\right)}{\log\left(\frac{1}{\alpha}-1\right)}$.

How agents decide at the time they have to make a decision depends on their beliefs about the state of the world and what they can gain by matching the state (or lose by not matching it), depending on each state. This can be captured by a threshold of the difference between high and low observed conditionally independent signals. This threshold also has (among others) the following expected properties: First, when the utilities in the different states are symmetric, the agent chooses $d_{\theta k} = 1$ if and only if she received more high than low conditionally independent signals. Second, when the utility in any state increases after $d_{\theta k} = 1$ (or the utility in any state decreases after $d_{\theta k} = 0$), the necessary threshold for the difference between high and low conditionally independent signals z_θ^* decreases. Third, when the precision of the signal increases, i.e., $\alpha \to 1$, the threshold decreases, too.

The model captures the simple intuition that agents benefit from more independent signals. An exception is the case $|\underline{S}_{\theta k}^t| < |z_\theta^*|$. Then, signals are not of any value to an agent, because her decision will be independent of the signals as those will never be enough to exceed the threshold. Another exception may occur when the threshold of the difference between high and low signals, z_θ^*, is an integer. Then agents may only benefit from every two additional signals. For example, imagine a situation in which the different utility levels are exactly symmetric with respect to the states, so $z_\theta^* = 0$. If an agent has an odd number of signals, one additional signal will not help her, because it either will not influence her decision (if the threshold is either exceeded or not with and without the additional signal) or make the agent indifferent and not provide a strictly higher utility (if the threshold is met exactly with the additional signal). Consequently, she does not benefit from just one additional signal.[17]

Lemma 4.2. *(i) The expected utility from the decision, $\mathbb{E}[\mathbb{E}[u_\theta(r, d_{\theta k}(S_{\theta k}^1))]]$, is increasing in $|\underline{S}_{\theta k}^t|$. For $|\underline{S}_{\theta k}^t| \geq |z_\theta^*|$, it is strictly increasing in every additional conditionally independent signal if $z_\theta^* \notin \mathbb{Z}$ and strictly increasing in every two additional conditionally independent signals if $z_\theta^* \in \mathbb{Z}$.*
(ii) The additional expected utility from the decision approaches zero for $|\underline{S}_{\theta k}^t| \to \infty$.

Lemma 4.2 states that usually agents value additional information, but the value-added is decreasing. Therefore, the willingness to pay for two additional links is positive and decreasing and the number of links agents want to form is limited for a given cost per link. This helps us to avoid corner solutions if link formation is not prohibitively costly (so that no agents

[17]For brevity and clarity, I will ignore the non-generic cases in which $z_\theta^* \in \mathbb{Z}$ and assume $z_\theta^* \notin \mathbb{Z}$ for the rest of this chapter after Lemma 4.2.

would want to form links), link formation is not free (so that agents would want to link to everybody), and the number of regions, n, is large (so that agents have a choice whom to link to).

4.4.2 Link Formation

This section investigates how agents form links. A first observation is captured by the following lemma.

Lemma 4.3. *Agents form inter-regional links only with agents of their own type.*

Forming links to other types is cheaper within an agent's own region and forming links to an agent in another region is cheapest if this agent has the same type. Recall that we assumed that the number of regions is large, and so the agent can always link to another agent of her own type. Note that Lemma 4.3 also holds if agents receive type- or region-specific signals, despite an agent with a different type in a different region providing potentially more conditionally independent signals at a first glance. However, an agent rather forms links to that type in her own region and to an agent of her own type in a different region.

Lemma 4.3 simplifies the analysis a lot, because it allows us to focus on the trade-off between the most expensive domestic link versus a link to an agent of the same type in another region.

Moreover, Lemma 4.3 implies that there will be sorting by types when inter-regional links are formed. If an agent forms a link to agents of the same type, she will observe more signals (and their realizations) of agents of her type. Thus, her beliefs and also her decisions will become more similar to those of other agents of her type. As an example, if types are ideologies, people who form inter-regional links will make decisions more similar to the decisions of people of the same ideology and become more consistent with one ideology. This is in line with the finding of the Pew Research Center (2014) that the share of consistent liberals rose from 5% in 1994 to 23% in 2014, and the share of consistent conservatives rose from 13% to 20%.

The following lemma captures the intuition that if inter-regional links are cheaper, agents will generally form more links, but will also substitute domestic links with inter-regional links.

Lemma 4.4. *(i) Suppose* $\overline{s}_T \geq \overline{s}_R$*. Agents form (weakly) more links when the inter-regional linking cost,* c_R*, decreases.*
(ii) Conditional on forming a positive number of links, agents form (weakly) fewer domestic links when the inter-regional linking cost, c_R*, decreases.*

There are two subtleties. First, the condition that the number of type-specific signals should not be lower than the number of region-specific signals, $\overline{s}_T \geq \overline{s}_R$, ensures that replacing domestic links with inter-regional links does not make fewer links optimal. If the condition did not hold, so $\overline{s}_T < \overline{s}_R$, it could happen that two domestic links are replaced with just one inter-regional link, because an inter-regional link provides access to more information. A second inter-regional link might not be worth the cost due to decreasing value-added of additional signals. Second, the number of domestic links might increase with a decreasing cost of inter-regional links if an agent did not form links at all at a higher cost of inter-regional links, as the next proposition will show.

Let me first define what it means for a type to be less central.

Definition 4.1. *Type* θ *is called less central than* η*, if it is further away from the median type* $\frac{m+1}{2}$*, i.e., if* $|\theta - \frac{m+1}{2}| > |\eta - \frac{m+1}{2}|$.

I may also refer to very low and very high types as *extreme* types. For example, these could be people with radical left or right views. They would probably find it hard to communicate with each other, but also people with more moderate views would not share a lot of common ground with them and therefore might find it hard to form a relationship with them. It is not important for the model that the types can be ordered on a line. The intuition for the results would continue to hold as long as there are some more central types with whom most other types share some common ground for a relationship and some less central types that other less central and not close types would find it quite difficult to talk to.

Proposition 4.1. *(i) Conditional on forming a positive number of links, the less central an agent's type, the (weakly) higher the fraction of her inter-regional links.*
(ii) For all types θ*, there exist* $c_T, c_R > 0$ *and utility parameters* $\overline{u}_{\theta i} = \overline{u}_{\eta i}$ *and* $\underline{u}_{\theta i} = \underline{u}_{\eta i}$ *for* $i \in \{0, 1\}$ *such that only types more central than* θ *form links to other agents and some types less central form links when* c_R *is reduced.*

The first part of the proposition means that agents who are not at the center of the society tend to have a higher fraction of inter-regional links. For them, it is more expensive to form links within a region, because they cannot necessarily link to higher and lower types in equal numbers as more central types can do. Therefore they find forming inter-regional links relatively more attractive.

The second part of the proposition shows that people with extreme types may not form links to other people, but may be isolated. As in part (i), it is more expensive for them to form links within a region. So it might not make sense to form links at all, if a lower number of signals would never let them influence their decision anyway. When inter-regional communication becomes cheaper, it also becomes cheaper to acquire this minimum number of signals that could make a difference for the decision, and so it might make sense to form several links. The present model predicts that some of these links should be inter-regional and some domestic at first (when c_R is just low enough), which would mean that social media might help to reintegrate people also in their local communities. However, it might also be possible that the differences between types are not uniform as in the present model, but that some extreme types may feel more distant and therefore face even higher costs when they want to form links to other types. Then it may be the case that these extreme types just form links to agents with the same type online.

Not only more extreme types, but also agents with higher stakes tend to form a higher fraction of inter-regional links.

Proposition 4.2. *Let $h > 0$ be a factor that multiplies an agent's gross utility in every state and every decision.*
(i) The fraction of inter-regional links this agent forms (weakly) increases in h.
(ii) For all c_R, there exists an $\overline{h} > 0$ such that the agent with $h > \overline{h}$ forms inter-regional links.

Proposition 4.2 says that agents with higher stakes form inter-regional links at a higher inter-regional linking cost than agents with lower stakes would and for the same cost, they form a higher fraction of such links. If the inter-regional linking cost decreases because of technological progress, agents with high stakes start forming inter-regional links earlier. This also implies that the sorting of politicians and other professionals with a high demand for information should start earlier and also be more extensive (at least before there is perfect sorting for everyone). The intuition is that agents with high stakes want to gather more information for their decision and therefore want to form more links. Intra-regional links, however, become increasingly costly as the types become less similar, so these agents rather form inter-regional links to agents with the same type at a high cost than domestic links to agents with very different types at an even higher cost.

4.4.3 Learning in endogenous networks

Now we look at the outcomes of learning in endogenous networks. This section does so for decisions in $t = 1$.

Remark 4.1. *Agents are (weakly) better off with a decrease in the inter-regional linking cost.*

Obviously, utility-maximizing rational agents become weakly better off as the cost of inter-regional links decreases and they become strictly better off when they form inter-regional links and the cost of inter-regional links decreases even further. However, the quality of their information, i.e., their expected utility before the link formation cost, does not necessarily increase monotonically. For example, if there are more type-specific signals than region-specific signals, i.e., $\overline{s}_T > \overline{s}_R$, it may happen that a domestic link is replaced by a less informative inter-regional link when the latter is sufficiently cheaper, and the agents form no additional links. Then the number of conditionally independent signals she observes is lower. Section 4.5.1 shows that agents with correlation neglect may be worse off.

4.4.4 Opinions in Endogenous Networks

Now we look at beliefs in $t > 1$. We say that there is a *consensus* if all agents hold the same belief. When there is no consensus, people disagree, and disagreement is stronger when the sum of differences between agents' beliefs is higher or it takes longer to reach a consensus (for a given number of agents). A *component* is a collection of agents in the network and all their links in which there is a (undirected) path between every pair of agents in this collection of agents and no path to agents not in this collection of agents. In a *strongly connected* component or graph, there is a directed path from every agent to every other agent.

Proposition 4.3. *(i) For all $c_T > 0$, there exist $\overline{c_R} > \underline{c_R} > 0$ such that for all $c_R > \overline{c_R}$ and all $c_R < \underline{c_R}$ the society consists of several components that are not connected with each other. Consensus occurs only by chance, even if $\bar{t} \to \infty$.*
(ii) For all sufficiently small $c_T > 0$, there exist $\overline{c_R} \geq \overline{c_R}^* > \underline{c_R}^* \geq \underline{c_R} > 0$ such that for all $c_R \in (\underline{c_R}^*, \overline{c_R}^*)$ the society consists of one strongly connected component. There is consensus if $\bar{t} \to \infty$.

Part (i) says that when the cost of inter-regional links is too high, agents will only form links within their region. Then there will be no connection between different regions. A consensus between two different regions, i.e., that all agents in these regions hold the same belief, therefore only occurs when the agents received the same combination of signals by chance. If the cost of inter-regional links is very low, agents will only link to agents of the same type. Then there will be no connection between different types. Again, consensus only occurs by chance. Part (ii) says that for an intermediate cost range of inter-regional linking, there is a directed path from every agent to every other agent in the society, so eventually all agents will have observed all signals and hold the same belief. (They might still decide differently, though, because their utility is type-specific.) To ensure that agents want to form some domestic links, c_T has to be sufficiently low. As we know from Proposition 4.1, agents of different types may differ in their link formation behavior. Therefore, the cost thresholds in part (i) do not necessarily have to coincide with the cost thresholds in part (ii). It may happen that an agent with a central type forms links both domestically and inter-regionally, while another agent does not form links at all or only domestically. Figure 4.1 illustrates Proposition 4.3 with five regions and five types.

Even if the society eventually reaches a consensus, this may happen quite fast or the society might stay polarized for a long time before reaching a consensus. How fast beliefs converge, depends on the length of shortest paths in the society. The minimal number of periods after which a consensus is reached for sure is equal to the maximal shortest path length plus one, because then all agents have observed all signals.

Lemma 4.5. *Suppose that the society is strongly connected, $\bar{s}_R = \bar{s}_T$, and $\underline{u}_{i\theta} = \underline{u}_{i\eta}, \overline{u}_{i\theta} = \overline{u}_{i\eta}$ for all $i \in \{0, 1\}, \theta, \eta \in M$. Then there exists a threshold $c_R^* > 0$ such that the maximal shortest path length weakly decreases in c_R for $c_R < c_R^*$ and weakly increases in c_R for $c_R > c_R^*$.*

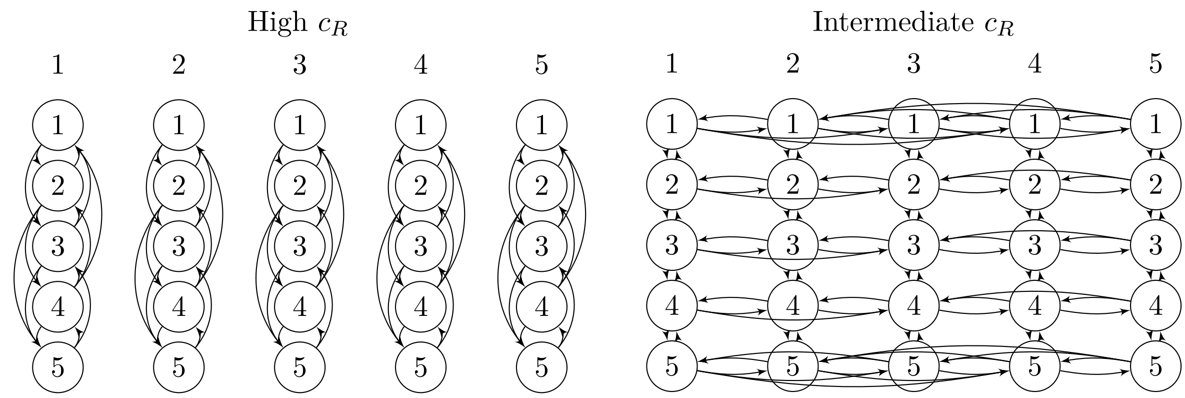

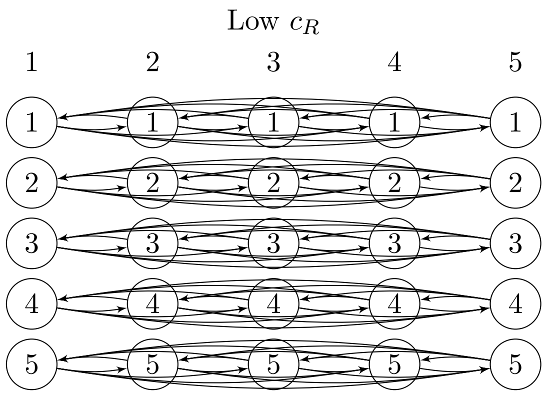

The numbers in the circles are the different types, the numbers in the top row are the different regions. If c_R *is high, agents only link within their region, and so agents will disagree across regions, even if* $t \to \infty$*. If* c_R *is intermediate, there will be a consensus from* $t = 6$ *on. If* c_R *is low, agents only link agents of the same type, and so agents will disagree across types, even if* $t \to \infty$*.*

Figure 4.1: Example with five regions and five types.

This means that the amount of time the society needs to reach a consensus first decreases when the cost of inter-regional links decreases, reaches a minimum, and then increases again. The reason is as follows. When the cost of inter-regional links decreases, more inter-regional links and fewer domestic links are formed. This speeds up the inter-regional exchange of information and slows down the exchange of information between types. When the cost of inter-regional links is high, there are only few inter-regional links, so the effect of faster inter-regional exchange of information dominates. When the cost of inter-regional links is low, the inter-regional exchange of information is already fast and the speed of exchange of information between types becomes the limiting factor and is reduced by replacing domestic links with inter-regional links. While there are many exceptions and special cases to consider if the gross utilities are type-dependent or the number of type-specific signals is not equal to the number of region-specific signals, the basic intuition remains the same. Together with Proposition 4.3, Lemma 4.5 shows that polarization follows a U-shape when the cost of inter-regional links decreases.

Proposition 4.4. *Suppose $\overline{s}_R = \overline{s}_T$ and $\underline{u}_{i\theta} = \underline{u}_{i\eta}, \overline{u}_{i\theta} = \overline{u}_{i\eta}$ for all $i \in \{0,1\}, \theta, \eta \in M$. For all sufficiently small $c_T > 0$, there exists a threshold $c_R^* > 0$ such that the maximal shortest path length weakly decreases in c_R for $c_R < c_R^*$ and weakly increases in c_R for $c_R > c_R^*$.*

4.5 Extensions and Discussion

4.5.1 Correlation Neglect

Experiments (Eyster and Weizsäcker 2011, Enke and Zimmermann 2015) have found that people have difficulties taking correlation of signals into account when they aggregate information and may double-count signals. This section investigates the case in which agents have correlation neglect, i.e., in which agents treat (conditionally) correlated signals as (conditionally) independent signals. It shows that most results continue to hold with a slight modification that lies in the nature of correlation neglect: Agents form their beliefs and base their decisions on signals they perceive as (conditionally) independent, but those do not have to be actually (conditionally) independent. However, unlike rational agents, agents with correlation neglect may be worse off when the inter-regional linking cost decreases.

An agent with correlation neglect, whom I may also call *naive*, treats all signals she observes as conditionally independent signals. Let $\hat{\underline{S}}_{\theta k}^t = S_{\theta k}^t$ be the set of *perceived* conditionally independent signals agent θk has observed up until and including period t. Partial naivete, i.e., $\underline{S}_{\theta k}^t < \hat{\underline{S}}_{\theta k}^t < S_{\theta k}^t$ can also be modeled easily, but I will focus on complete naivete here. I conjecture partial naive agents to behave and form beliefs in between (completely) naive and rational agents.[18] Note that naive agents (whether completely or partially) cannot be told apart from rational agents if all signals indeed are conditionally independent.

In $t = 1$ each naive agent maximizes her *perceived* expected utility from the decision given the signals she observed:

$$\hat{\mathbb{E}}[u_\theta(r, d_{\theta k}(S_{\theta k}^1))] := \tag{4.3}$$

$$\max_{d_{\theta k}(S_{\theta k}^1) \in [0,1]} d_{\theta k} \quad \left[\widehat{Prob}(r = 1|S_{\theta k}^1)\overline{u}_{1\theta} + (1 - \widehat{Prob}(r = 1|S_{\theta k}^1))\underline{u}_{1\theta}\right]$$
$$-(d_{\theta k} - 1) \quad \left[\widehat{Prob}(r = 1|S_{\theta k}^1)\underline{u}_{0\theta} + (1 - \widehat{Prob}(r = 1|S_{\theta k}^1))\overline{u}_{0\theta}\right],$$

[18]There is no reason to expect a discontinuity as often found with partially naive present-biased agents (**?**, Eliaz and Spiegler 2006, Heidhues and Kőszegi 2010, Englmaier, Fahn and Schwarz 2016), because for partial correlation neglect, the equations determining the beliefs and actions are influenced gradually.

where $\hat{\mathbb{E}}$ and $\widehat{Prob}$ use the perceived probabilities as if all signals were conditionally independent.

In $t = 0$, each naive agent forms links that maximize her perceived expected utility which consists of the perceived expected utility from the decision minus the cost for the links:

$$\max_{\mathbf{a}_{\theta k} \in \{0,1\}^{mn}: a_{\theta k \theta k}=0} \hat{\mathbb{E}}[U_{\theta k}] := \max_{\mathbf{a}_{\theta k} \in \{0,1\}^{mn}: a_{\theta k \theta k}=0} \hat{\mathbb{E}}[\hat{\mathbb{E}}[u_\theta(r, d_{\theta k}(S^1_{\theta k}))]] - C(\mathbf{a}_{\theta k}). \tag{4.4}$$

Let furthermore $\hat{z}^t_{\theta k}$ be the difference between high and low perceived conditionally independent signals agent θk observed up until period t, which is equal to the difference between high and low signals. This difference is a sufficient statistic for a naive agent's belief and she forms links and decides just like a rational agent, but on grounds of $\hat{z}^t_{\theta k}$ instead of $z^t_{\theta k}$.

Remark 4.2. *All previous lemmas and propositions for rational agents also hold for naive agents when one replaces $z^t_{\theta k}$ with $\hat{z}^t_{\theta k}$, $\underline{S}^t_{\theta k}$ with $\hat{\underline{S}}^t_{\theta k}$, "expected" with "perceived expected", and "conditionally independent" with "perceived conditionally independent."*

The reason is that agents' decisions and links depend on their expectations. Consequently, that the patterns of link and opinion formation do not change compared to the case of rational agents, if all agents are naive. This does not mean, however, that the thresholds for naive agents are the same as for rational agents. In general, the thresholds could be higher or lower. To see that, assume first that both $\overline{s}_T$ and $\overline{s}_R$ are large, i.e., there are many type-specific and region-specific signals. Then agents with correlation neglect think that they already observe many independent signals when they only link to a few agents within their region. By contrast, rational agents are aware that they do not have that many independent signals and are willing to link inter-regionally at a higher cost than agents with correlation neglect. Assume then that $\overline{s}_T$ is very large and $\overline{s}_R$ is small. A rational agent prefers to link domestically as she is aware that many of the signals would be correlated if she linked to agents with the same type. Agents with correlation neglect, however, are not aware of this, and will link inter-regionally at higher inter-regional linking costs than rational agents will.

In societies with both rational and naive agents, it could theoretically happen that rational agents only form domestic links and naive agents only form inter-regional links or the other way around. Consequently, societies with both rational agents and agents with correlation neglect may have more difficulties reaching a consensus for two reasons: Firstly and obviously, rational agents and agents with correlation neglect agree only by chance. Secondly, if rational agents and agents with correlation neglect are not distributed uniformly, the society, i.e., the graph that includes all agents, is less likely to be strongly connected.

Remark 4.2 implies that naive agents think they are better off with a lower inter-regional linking cost, but what about their actual utility? The following two lemmas state how their actual utility is affected by additional correlated signals.

Lemma 4.6. *Suppose an agent with correlation neglect observes a set of conditionally independent signals S_0 and a set of perfectly correlated signals S_1, where all signals in S_1 are conditionally independent from all signals in S_0. The expected utility from the decision is decreasing in $|S_1|$. It is strictly decreasing for $|S_1| < |S_0| - |z_\theta^*|$.*
Furthermore, the expected utility from the decision decreases when both $|S_1|$ and $|S_0|$ increase by the same amount. It strictly decreases for $|S_1| < |S_0| - |z_\theta^*|$. The reduction in expected utility converges to zero for $|S_0| \to \infty$.

Lemma 4.6 says that an agent with correlation neglect is in general hurt by correlated signals. The decrease in expected utility even exceeds the gain by an additional independent signal. To see the intuition, imagine the agent observes an equal number of independent and perfectly correlated signals and chooses 1 if the difference between high and low signals is strictly positive. When the correlated signals are high, even though the state of the world is 0, only if all independent signals are low, the agent makes the (ex-post) right choice. This is similar when the agent observes more correlated and equally more independent signals. However, the probability that all independent signals are low decreases with the number of signals.
However, there are some situations in which an additional correlated signal does not hurt or even benefits the agent. First, if the condition $|S_1| < |S_0| - |z_\theta^*|$ is not met, the decision is based on only the correlated signals anyway, so one additional correlated signal does not make it worse. Second, naive agents can potentially benefit from more correlated signals when there is another set of perfectly correlated signals with a higher cardinality, as Lemma 4.7 below shows. Then, if these other correlated signals are wrong, the first set of correlated signals is more effective in making up for it. At the same time, however, the independent signals less often make a difference and which effect dominates depends on $|S_0|, |S_1|, |S_2|, \alpha$, and the different utility levels. An example of a situation where the agent benefits from more correlated signals is when $|S_1|$ and $|S_2|$ are already very high compared to $|S_0|$ so that the

independent signals do not play a role for the decision when the correlated signals are all high or all low. Then one additional correlated signal does not hurt the agent when the correlated signals are all high or all low, but makes the independent signals more decisive when the signals of one group of correlated signals are high and the signals of the other group are low. Furthermore, Lemma 4.7 below says that an agent with correlation neglect benefits from replacing a signal that is perfectly correlated with one set of signals with a signal that is perfectly correlated with another set of signals if the cardinality of the latter set is lower than the cardinality of the first set plus one, i.e., if one group of signals becomes less dominant for the decision.

Lemma 4.7. *Suppose an agent with correlation neglect observes a set of conditionally independent signals S_0, a set of perfectly correlated signals S_1, and another set of perfectly correlated signals S_2, where all signals in S_i are conditionally independent from all signals in S_0 and S_j, $i \neq j \in \{1, 2\}$.*
(i) The expected utility from the decision can strictly increase in $|S_1|$ if $|S_1| < |S_2|$ and $|S_2| < |S_1| + |S_0| - |z^*_\theta|$. If $|S_1| > |S_2|$, and $|S_1| < |S_2| + |S_0| - |z^*_\theta|$, it strictly decreases in $|S_1|$.
(ii) The expected utility from the decision strictly increases if $|S_1|$ is decreased by one and $|S_2|$ is increased by one if and only if $|S_1| > |S_2|+1$. It strictly decreases if and only if $|S_1| \leq |S_2|$.

Lemma 4.6 implies that the U-shape of disagreement with respect to the inter-regional linking cost is amplified if agents are naive: If agents only link domestically, they are exposed to many region-specific signals of that region and therefore likely to be overly influenced by these signals, because they treat them as conditionally independent. Similarly, if agents only link inter-regionally, they are exposed to many type-specific signals of their type and therefore are likely to be overly influenced by these signals. Because of this overweighting of certain signals, the distribution of beliefs has a greater variance compared to rational agents' beliefs, i.e., disagreement in these cases is stronger. The following corollary assumes $s^i_{\theta k} = s^j_{\theta k} \forall i, j = \overline{s}_U + 1, ..., \overline{s}_T$ and $s^i_{\theta k} = s^j_{\theta k} \forall i, j = \overline{s}_U + \overline{s}_T + 1, ..., \overline{s}_U + \overline{s}_T + \overline{s}_R$, i.e., all type-specific signals of each type are perfectly correlated and all region-specific signals are perfectly correlated in each region.

Corollary 4.1. *Suppose all type-specific signals of each type are perfectly correlated and all region-specific signals are perfectly correlated in each region. Further suppose all agents only link inter-regionally or all agents only link domestically. Disagreement is stronger if all agents are naive than if all agents are rational.*

We are now ready to say how agents with correlation neglect are affected when inter-regional link formation becomes cheaper. It is difficult to make general statements as for example replacing the last domestic link can make those agents better off when the number of region-specific signals is high enough. Similarly, already the first inter-regional link or forming an additional inter-regional link could make those agents worse off when the number of type-specific signals is high enough. Part (ii) of Proposition 4.5 shows that naive agents are likely to suffer from worse decisions due to a decreasing inter-regional linking cost if there are many type-specific signals.

Proposition 4.5. *Suppose all type-specific signals of each type are perfectly correlated and all region-specific signals are perfectly correlated in each region.*
(i) Then the expected utility from the decision of a naive agent increases when c_R decreases if the number of inter-regional links does not exceed the number of domestic links, the number of total links is unchanged, and $\overline{s}_T \leq \overline{s}_R$. The expected utility from the decision of a naive agent decreases when c_R decreases if the number of domestic links does not exceed the number of inter-regional links, the number of total links is unchanged, and $\overline{s}_T \geq \overline{s}_R$.
(ii) If $\overline{s}_T \geq \frac{\overline{s}_U+\overline{s}_T+\overline{s}_R}{2}$, the expected utility from the decision of a naive agent decreases when c_R decreases if the number of domestic links does not exceed the number of inter-regional links.

Agents with correlation neglect tend to benefit from a lower inter-regional link formation cost when they replace the first domestic links with inter-regional ones. However, when the number of inter-regional links is already high, those agents' decisions become worse. This is true in particular when there are at least as many type-specific signals as other signals. Also note that agents who make worse decisions when they observe more signals have more 'extreme' beliefs in general. They make ex-post wrong decisions because some signal realizations do not coincide with the state of the world and are over-weighted; similarly, these signal realizations would have been over-weighted if they had coincided with the state of the world and would have led to a more confident belief.

Corollary 4.2. *A decrease in c_R can make agents with correlation neglect worse off, while it always makes rational agents better off.*

This means that a decrease in inter-regional link formation cost may increase the difference in utilities between rational agents and agents with correlation neglect. Consequently, if limiting the difference in utilities is among a policy maker's goals, the rise of social media provides an additional incentive to protect or educate naive citizens.

Proposition 4.6. *Suppose $\overline{s}_R = 0$ and $\overline{s}_T \geq \overline{s}_U$ and all type-specific signals are perfectly correlated for each type. Then the expected utility from the decision of a naive agent decreases when c_R decreases.*
Suppose $\overline{s}_T = 0$ and $\overline{s}_R \geq 0$. Then the expected utility from the decision of a naive agent increases when c_R decreases.

Proposition 4.6 shows that when type-specific correlation is strong, a decrease in inter-regional link formation cost leads to worse decisions of agents with correlation neglect and at the same time therefore also to more extreme beliefs and more differences across types. For example, when within a country, people with the same hobby or sympathizing with the same party also have a similar media consumption behavior, they will be more likely to buy the same products or share the same opinions when the inter-regional link formation cost decreases, whereas people with a different hobby or sympathizing with a different party will be more likely to buy completely different products or have very different opinions.
When region-specific correlation is strong, a decrease in inter-regional link formation cost leads to better decisions and smaller differences across regions. For example, people in different countries watch different news programs and read different newspapers. Via social media people now also get more exposure to news and perspectives in other countries. This leads to better decisions and more similar beliefs across countries.

4.5.2 Social Algorithms and Vicious Circles

Social algorithms, which filter content on some social media platforms, allow an agent to more easily observe signals of the same type on a social media platform, but at the same time make it more difficult to observe the signals of a different type. I model this as a decrease in inter-regional linking cost combined with an increase in the type-specific linking cost. As these two changes do not interact directly and the paper already analyzed the decrease in inter-regional linking cost, it remains to be shown what an increase in the type-specific linking cost entails.

Lemma 4.8. *If the type-specific linking cost, c_T, increases, agents form (weakly) fewer domestic links. Agents also form (weakly) more inter-regional links, if they form a positive number of links after the increase in c_T.*

Lemma 4.8 states that an increase in the type-specific linking cost usually leads to an increase in inter-regional links and a decrease in domestic links. This is very similar to the effect of a decrease in the inter-regional linking cost. An exception occurs, if the increase in the type-specific linking cost makes the links so expensive that an agent does not want to form links at all anymore. In combination with a decrease in the inter-regional linking cost, this seems unlikely, though. Consequently, in general social algorithms accelerate and amplify the effects of an increasing importance of social media.

Lemma 4.8 also implies that, if increasing disagreement across types leads to increasing costs to link to different types, this leads to fewer links to different types, which in turn leads to more disagreement across types. A vicious circle may emerge.

4.5.3 Strategic Communication

I have assumed that every agent can observe the signals of the agents she has linked to, which is equivalent to other agents telling the truth. This is plausible if an individual's utility is not affected by the choices of others, if individuals have a strong preference for truthful communication towards their contacts, or if individuals do not take potential consequences of strategic communication into account. In another interpretation of the model, the assumption that every agent can observe the signals of the agents she has linked to is a quite innocent one: If signals are hard evidence and agents know how many signals other agents observe, concealing a signal would not be effective, because the realization of that signal could be inferred. Consequently, the results are robust to strategic behavior in some interpretations of the model. However, because situations in which signals are hard evidence may be rare and correlation neglect seems implausible in such situations, there is some room for future research.

4.5.4 Decisions with Several Alternatives

In many situations, people do not only have to pick one of two alternatives, but face several alternatives. These can either be ordered or not. As an example of ordered alternatives, an agent could face the question of how much to invest in a certain project. Her decision will depend on how confident she is that the project will be successful. In this case, the different alternatives can easily be represented by different thresholds. As an example of alternatives that cannot be ordered, an agent might have to choose between a yellow, blue, or red car. In this case, the problem can be split up into several questions between two alternatives, in which the states of the world might be correlated. In both of these cases, the link formation problem becomes more involved, but the intuition for the results does not change.

4.6 Conclusion

In this chapter I combine an endogenous network formation model with learning to investigate the question of how social media impact belief formation. In the model, agents live in different regions and have different types that can be ordered on a line. I find that agents with less central types and agents with higher stakes form a higher fraction of inter-regional links. These are also the agents who are willing to form inter-regional links at a higher cost than other agents. Moreover, I show that if these links and indirect links also influence agents' opinions, disagreement in a society is U-shaped with respect to the increasing importance of social media, which I model as a decrease in inter-regional linking cost. As an extension, I propose a portable model of correlation neglect and demonstrate that correlation neglect may amplify the U-shape and that agents with correlation neglect might be worse off with a lower inter-regional linking cost due to increasing importance of social media. Furthermore, I identify situations in which agents with correlation neglect may benefit from perfectly correlated signals.

It is hard to empirically identify whether the increasing importance of social media is still helping us to reduce disagreement in a society or is already increasing it, but the mechanism in this chapter shows that the disagreement is likely to go up in the future. While the policy maker will probably be unable and unwilling to stop technological progress, easing domestic communication by subsidizing local events or clubs might be helpful to counteract this increasing disagreement. I also demonstrate that social media increases the need to integrate all people in society, because with social media they are more likely to radicalize themselves online. Moreover, the model suggests that educating citizens with correlation neglect becomes more important with growing influence of social media for two reasons. First, increasing disagreement in a society is stronger when citizens have correlation neglect. Second, people with correlation neglect might be worse off with social media, while rational people are better off.

In general, networks can be difficult to analyze. As this chapter demonstrates, imposing a simple and plausible structure on link formation costs allows one to derive a variety of results and predictions, and thus this seems promising for future research. For example, this might allow to investigate how a social network changes when a new technology or product is introduced or how people could adjust their interactions with other people if they do not want to be infected by a spreading disease. The advantages of an approach similar to the one in this chapter are that the network changes endogenously due to strategic considerations (in contrast to models with a random network formation process) and that the structure is less simplified than in other models of strategic network formation but still tractable.

This chapter assumed that agents are uniformly distributed across regions. While this seems to be a reasonable simplification to analyze a model of network formation and learning and to derive some new mechanisms and insights, this is certainly not a realistic assumption and further research is needed. It might be interesting to see the interaction between social media and the incentives for people to move to a different region. On the one hand, the increasing number of possibilities provided by social media makes moving less desirable, because inter-regional communication becomes easier. On the other hand, the growing influence of social media may promote alienation (as discussed in Section 4.5.2) and thus make moving more desirable. Furthermore, the model in this chapter, but with a different distribution of agents across regions, could serve as a basic framework to analyze in which regions investments in communication technologies have the highest impact.

Appendix A

Appendix for Chapter 1.

A.1 Proofs and Calculations

Proof of Lemma 1.1. First, we show that the parties agree to trade $x^*(\theta)$. In contradiction, let $(\hat{x}, \hat{p})$ with $\hat{x} \neq x^*(\theta)$ be the outcome of renegotiation. The resulting generalized Nash product is

$$GNP(\hat{x}, \hat{p}) = [\hat{p} - c(\hat{x}, \theta) - d_C]^{\alpha} \times [v(\hat{x}) - \hat{p} - d_P]^{1-\alpha}. \tag{A.1}$$

Consider the alternative contract with design $x^*(\theta)$ and price $p^* = \hat{p} + v(x^*(\theta)) - v(\hat{x})$. By construction, the procurer is indifferent between the two contracts. The contractor's net payoff under the alternative contract is

$$p^* - c(x^*(\theta), \theta) - d_C. \tag{A.2}$$

Hence, the contractor prefers the alternative contract if and only if

$$\begin{aligned} \hat{p} + v(x^*(\theta)) - c(x^*(\theta), \theta) - v(\hat{x}) - d_C &\geq \hat{p} - c(\hat{x}, \theta) - d_C \\ \iff v(x^*(\theta)) - c(x^*(\theta), \theta) &\geq v(\hat{x}) - c(\hat{x}, \theta), \end{aligned} \tag{A.3}$$

which holds by the definition of $x^*(\theta)$ and the fact that $\overline{x}$ does not maximize the social surplus. Thus, $GNP(x^*(\theta), p^*) > GNP(\hat{x}, \hat{p})$ a contradiction to the assumption that $(\hat{x}, \hat{p})$ is the outcome of renegotiation.

Taking the partial derivative of the generalized Nash product with respect to p yields

$$\frac{\partial GNP}{\partial p} = \alpha \left[\frac{v(x^*(\theta)) - p - v(\overline{x}) + \overline{p}}{p - c(x^*(\theta), \theta) - \overline{p} + c(\overline{x}, \theta)}\right]^{\alpha} - (1-\alpha)\left[\frac{p - c(x^*(\theta), \theta) - \overline{p} + c(\overline{x}, \theta)}{v(x^*(\theta)) - p - v(\overline{x}) + \overline{p}}\right]^{1-\alpha}. \quad \text{(A.4)}$$

We set the partial derivative equal to zero and solve for the renegotiation price

$$\hat{p} = \overline{p} + \alpha[v(x^*(\theta)) - v(\overline{x})] + (1-\alpha)[c(x^*(\theta), \theta) - c(\overline{x}, \theta)]. \quad \text{(A.5)}$$

□

Proof of Lemma 1.2. It is a well-known result that in a second price auction it is a (weakly) dominant strategy for each bidder to bid his type. Placing a bid equal to the type, corresponds to placing a price bid so that the profit equals zero in our setup. Placing a higher bid reduces the probability of winning the auction without affecting the price $\overline{p}$. A lower bid is not optimal because in the additional cases where the supplier now wins the auction, he makes losses.

It remains to be shown that $\theta_1 < \theta_2$ implies $\overline{p}(\theta_1) < \overline{p}(\theta_2)$. This property of the bidding function follows immediately from Assumption 1.2. Note that $S(\theta) \equiv \max_x\{v(x) - c(x, \theta)\}$ and thus $S(\theta_1) > S(\theta_2)$ by Assumption 1.2. □

Proof of Proposition 1.2. It is readily established that the bidding strategy $\overline{p}(\theta)$ for types $\theta \leq r$ constitutes a Nash equilibrium of the (reduced) auction game. Thus, the procurer maximizes the following expression via the design $\overline{x}$ and the threshold type r:

$$\begin{aligned}
&\int_{\underline{\theta}}^{r}\int_{\underline{\theta}}^{\ddot{\theta}} \Big[(1-\alpha)S(\dot{\theta}) + \alpha S(\ddot{\theta}) - (1-\alpha)[c(\overline{x}, \ddot{\theta}) - c(\overline{x}, \dot{\theta})]\Big] \\
&n(1-F(\dot{\theta}))^{n-1} f(\dot{\theta}) n(n-1)(1-F(\ddot{\theta}))^{n-2} F(\ddot{\theta}) f(\ddot{\theta}) d\dot{\theta} d\ddot{\theta} \\
&+ \int_{r}^{\overline{\theta}}\int_{\underline{\theta}}^{r} \Big[(1-\alpha)S(\dot{\theta}) + \alpha S(r) - (1-\alpha)[c(\overline{x}, r) - c(\overline{x}, \dot{\theta})]\Big] \\
&n(1-F(\dot{\theta}))^{n-1} f(\dot{\theta}) n(n-1)(1-F(\ddot{\theta}))^{n-2} F(\ddot{\theta}) f(\ddot{\theta}) d\dot{\theta} d\ddot{\theta}.
\end{aligned}$$

The first summand describes the situation when the second-lowest bid is below the maximum bid, while the second summand describes the situation when the maximum bid binds for determining the price. Since this expression is decreasing in $[c(\overline{x}, \ddot{\theta}) - c(\overline{x}, \dot{\theta})]$ and $[c(\overline{x}, r) - c(\overline{x}, \dot{\theta})]$, an optimal solution requires $\overline{x} = x_1$ by Assumption 1.3. □

Proof of Lemma 1.3. Each supplier has an incentive to place the bid (x^b, p^b) that maximizes the score $G(x,p)$ subject to the supplier's break-even constraint.

Bidding a lower score reduces the probability of winning without affecting the concluded contract in case the supplier wins the auction. As in a second price auction, the concluded contract is independent of the bid placed by the winner.

Bidding a higher score increases the probability of winning. In the additional cases where the supplier now wins, he has to match a score at which he makes losses.

Hence, the optimal bid solves:

$$\begin{aligned} &\max_{x,p} \quad g(x) - p \\ &\text{s.t.} \quad p + \alpha S(\theta) - \alpha v(x) - (1-\alpha)c(x,\theta) \geq 0. \end{aligned}$$

The solution is $x^B(\theta)$ and $p^b(\theta)$, which concludes the proof; see also Che (1993). $\square$

Proof of Lemma 1.4.

$$\begin{aligned} G(\theta_2) = &g(x^b(\theta_2)) - \alpha v(x^b(\theta_2)) - (1-\alpha)c(x^b(\theta_2),\theta_2) + \alpha S(\theta_2) \\ &< g(x^b(\theta_2)) - \alpha v(x^b(\theta_2)) - (1-\alpha)c(x^b(\theta_2),\theta_1) + \alpha S(\theta_1) \\ &\leq g(x^b(\theta_1)) - \alpha v(x^b(\theta_1)) - (1-\alpha)c(x^b(\theta_1),\theta_1) + \alpha S(\theta_1) \\ &= G(\theta_1). \end{aligned}$$

The first inequality follows from Assumption 1.2 and the second inequality holds by the definition of $x^b(\cdot)$. $\square$

Proof of Corollary 1.1. Follows immediately from Proposition 1.3. $\square$

Proof of Proposition 1.4. If renegotiation takes place, then the outcome is characterized by Lemma 1.1. The expected ex post utility of the contractor from contract $(\overline{x}, \overline{p})$ is

$$\pi(\overline{x},\overline{p},\theta) = (1-b)[\hat{p}(\overline{p},\theta) - c(x^*(\theta),\theta) + b[\overline{p} - c(\overline{x},\theta)] \tag{A.6}$$

$$= \overline{p} + (1-b)\alpha[S(\theta) - v(\overline{x})] - [1-\alpha(1-b)]c(\overline{x},\theta). \tag{A.7}$$

From the above expression the next result is readily obtained.

Lemma A.1. *The symmetric equilibrium bidding strategy is*

$$\overline{p}(\theta) = (1-b)\alpha[v(\overline{x}) - S(\theta)] + [1-\alpha(1-b)]c(\overline{x}, \theta).$$

The above lemma can be proven by the usual steps (as in the proof of Lemma 1.2).

The most efficient type $\dot{\theta}$ wins the auction and the initial price is determined by the second most efficient type $\ddot{\theta}$, which is given by

$$\overline{p}(\ddot{\theta}) = (1-b)\alpha[v(\overline{x}) - S(\ddot{\theta})] + [1-\alpha(1-b)]c(\overline{x}, \ddot{\theta}). \tag{A.8}$$

The final price is given by

$$\begin{aligned}\hat{p}(\dot{\theta}, \ddot{\theta}) = \alpha(1-b)v(\overline{x}) - \alpha(1-b)S(\ddot{\theta}) + [1-\alpha(1-b)]c(\overline{x}, \ddot{\theta}) \\ + \alpha[v(x^*(\dot{\theta})) - v(\overline{x})] + (1-\alpha)[c(x^*(\dot{\theta}), \dot{\theta}) - c(\overline{x}, \dot{\theta})].\end{aligned} \tag{A.9}$$

The procurer's ex post utility for given realizations of $\dot{\theta}$ and $\ddot{\theta}$ is

$$u(\overline{x}, \dot{\theta}, \ddot{\theta}) = b[v(\overline{x}) - \overline{p}(\ddot{\theta})] + (1-b)[v(x^*(\dot{\theta}) - \hat{p}(\dot{\theta}, \ddot{\theta})]. \tag{A.10}$$

Inserting the expressions for $\hat{p}$ and $\overline{p}$ in the procurer's utility and rearranging yields

$$\begin{aligned}u(\overline{x}, \dot{\theta}, \ddot{\theta}) = (1-b)[(1-\alpha)S(\dot{\theta}) + \alpha S(\ddot{\theta})] \\ - (1-b)(1-\alpha)[c(\overline{x}, \ddot{\theta}) - c(\overline{x}, \dot{\theta})] + b[v(\overline{x}) - c(\overline{x}, \ddot{\theta})].\end{aligned} \tag{A.11}$$

Noting that the procurer maximizes $\mathbb{E}[u(\overline{x}, \dot{\theta}, \ddot{\theta})]$ by choosing $\overline{x}$ completes the proof. □

Supplementary Calculations to Example 1.1. The first-order condition of the procurer's maximization problem is

$$\mathbb{E}_{\dot{\theta},\ddot{\theta}}\Big[b[1-\ddot{\theta}\overline{x}] - (1-\alpha)(1-b)\overline{x}(\ddot{\theta} - \dot{\theta})\Big] = 0. \tag{A.12}$$

When an interior solution exists, i.e $\overline{x} > 1/2$ is optimal, then this solution satisfies the first-order condition. By using the distribution of the lowest and the second lowest type realization (order statistics), the first-order condition can be written as

$$\int_1^2 [b(1-\ddot{\theta}\overline{x}) - (1-\alpha)(1-b)\ddot{\theta}\overline{x}]n(n-1)(\ddot{\theta}-1)(2-\ddot{\theta})^{n-2}d\ddot{\theta} + \int_1^2 (1-\alpha)(1-b)\dot{\theta}\overline{x}n(2-\dot{\theta})^{n-1}d\dot{\theta} = 0. \quad \text{(A.13)}$$

Solving the above expression for $\overline{x}$ yields

$$\overline{x} = \frac{b(1+n)}{1-\alpha+b(2+\alpha+n)} < 1. \quad \text{(A.14)}$$

A corner solution, $\overline{x} = x_1 = 1/2$, is optimal if

$$1 - \alpha(1-b) - bn \geq 0,$$

which is equivalent to

$$b \leq \frac{1-\alpha}{n-\alpha} < 1. \quad \text{(A.15)}$$

Notice that for $\alpha < 1$ the critical b is strictly positive.

An important observation is that $\overline{x}$ is increasing in b. Thus, the highest $\overline{x}$ is selected for $b \to 1$. For $b \to 1$ (or $\alpha \to 1$), we obtain

$$\overline{x} = \frac{1+n}{3+n} > \frac{1}{2}, \quad \text{(A.16)}$$

which is the socially optimal design for the expected second lowest type. If the probability of bargaining breakdown is sufficiently high, specifying the simplest design ex ante is not optimal. Now, renegotiation may lead to a downward as well as to an upward adjustment of x ex post. Upward renegotiation takes place if $x^*(\dot{\theta}) > \overline{x}$, which is equivalent to ($b = 1$)

$$\dot{\theta} < \frac{3+n}{1+n}. \quad \text{(A.17)}$$

Hence, the probability with which a cost overrun occurs is $\rho^U = Prob(\dot{\theta} < (3+n)/(1+n)$. Due to the continuous design, the probability of a downward adjustment ex post is $\rho^D = 1 - \rho^U$. Hence, a cost overrun occurs ex post with probability

$$\rho^U(n) = F_{(1)}\left(\frac{3+n}{1+n}\right) = 1 - \left(2 - \frac{3+n}{1+n}\right)^n = 1 - \left(\frac{n-1}{n+1}\right)^n, \quad \text{(A.18)}$$

where $F_{(1)}(\cdot)$ denotes the c.d.f. of the lowest type realization. Importantly, $\rho^U(n)$ is decreasing in n and

$$\lim_{n\to\infty} \rho^U(n) = 1 - \frac{1}{\mathrm{e}^2} \approx 0.865, \tag{A.19}$$

where e is Euler's number.

Appendix B

Appendix for Chapter 2.

B.1 Proofs

Proof of Lemma 2.1.

In what follows, we analyze a slightly more general case in which $k_1 = k_2 = k$. We divide the analysis into two cases.

First, suppose that the firm sells the service only to naive consumers. In this case, the firm sets $p^a = \frac{1}{\beta}k$. Naive consumers do not switch, whereas time-consistent consumers switch either in period 1 or in period 2. The firm's profits are $\pi = \frac{\alpha}{\beta}k$ and the consumers' long-run utilities are $u^N = \bar{u} - \frac{1}{\beta}k + 2a$ and $u^{TC} = \bar{u} - k + 2a$.

Second, suppose that the firm sells the service to both naive and time-consistent consumers. In this case, the firm sets $p^a = k$. The firm's profits are k and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u} - k + 2a$.

By comparing the above two cases, we obtain the result. $\square$

Proof of Lemma 2.2.

(i) Note that time-consistent consumers do not switch in period 1 if and only if $p^a \leq \underline{k}$. Naive consumers do not switch in period 1 because $-\underline{k} \leq -\beta\overline{k}$.

First, suppose that the firm sells the service only to naive consumers. In this case, the firm sets $p^a = \frac{1}{\beta}\overline{k}$. Naive consumers do not switch, whereas time-consistent consumers switch in period 1. The firm's profits are $\pi = \frac{\alpha}{\beta}\overline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - \frac{1}{\beta}\overline{k} + 2a$ and $u^{TC} = \bar{u} - \underline{k} + 2a$.

Second, suppose that the firm sells the service to both naive and time-consistent consumers. In this case, the firm sets $p^a = \underline{k}$. The firm's profits are $\pi = \underline{k}$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u} - \underline{k} + 2a$.

By comparing the above two cases, we obtain the result.

(ii) Note that time-consistent consumers do not switch in period 1 if and only if $p^a \leq \underline{k}$. Naive consumers do not switch in period 1 if and only if $p^a \leq \frac{1}{\beta}\underline{k}$.

First, suppose that the firm sells the service only to naive consumers. In this case, the firm sets $p^a = \frac{1}{\beta}\underline{k}$. Naive consumers do not switch, whereas time-consistent consumers switch in period 1. The firm's profits are $\pi = \frac{\alpha}{\beta}\underline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - \frac{1}{\beta}\underline{k} + 2a$ and $u^{TC} = \bar{u} - \underline{k} + 2a$.

Second, suppose that the firm sells the service to both naive and time-consistent consumers. In this case, the firm sets $p^a = \underline{k}$. The firm's profits are $\pi = \underline{k}$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u} - \underline{k} + 2a$.

By comparing the above two cases, we obtain the result. $\square$

Proof of Proposition 2.1.

The conditions in which the policy in $t = 1$ increases the equilibrium price and decreases social welfare, $1 \geq \frac{\alpha}{\beta} > \Delta_k \geq \beta$, are immediate from Lemma 2.1 and Lemma 2.2.

Given $1 \geq \frac{\alpha}{\beta} > \Delta_k \geq \beta$, the total consumer surplus under no policy is $\bar{u} + 2a - \bar{k}$, whereas under the policy in $t = 1$ it is $\bar{u} + 2a - \alpha\frac{1}{\beta}\overline{k} - (1 - \alpha)\underline{k}$. Comparing these two cases, we get the condition in which the policy in $t = 1$ decreases consumer welfare if and only if $\frac{\alpha}{\beta} + (1 - \alpha)\Delta_k > 1$. $\square$

Proof of Lemma 2.3.

We divide the analysis into two cases.

First, suppose that the firm sells the service only to naive consumers. In this case, the firm sets $p^a = \frac{1}{\beta}\underline{k}$. Naive consumers do not switch, whereas time-consistent consumers switch in period 2. The firm's profits are $\pi = \frac{\alpha}{\beta}\underline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - \frac{1}{\beta}\underline{k} + 2a$ and $u^{TC} = \bar{u} - \underline{k} + 2a$.

Second, suppose that the firm sells the service to both naive and time-consistent consumers. In this case, the firm sets $p^a = \underline{k}$. The firm's profits are $\pi = \underline{k}$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u} - \underline{k} + 2a$.

By comparing the above two cases, we obtain the result. □

Proof of Proposition 2.2.

Immediate from Lemma 2.1 and Lemma 2.3. □

Proof of Proposition 2.3.

Immediate from Lemma 2.2 and Lemma 2.3. □

Proof of Corollary 2.1. Note that consumers do not have an incentive to change their behavior, if not only k_2 but also k_1 is reduced to $\underline{k}$ for the equilibrium prices. Hence, the firm does not have an incentive to change its pricing strategy. If reducing a period's switching cost is costly, social welfare is lower when the switching cost is reduced in two periods compared to when the switching cost is reduced in only in $t = 2$. □

Proof of Proposition 2.4.

Before the proof, we characterize the consumer switching behavior. For notational simplicity, let β^i be consumer i's degree of present bias where time-consistent consumers have $\beta^{TC} = 1$ and (partially) naive present-biased consumers have $\beta^N = \beta < 1$. Similarly, let $\hat{\beta}^i$ be consumer i's belief about her degree of present bias where $\hat{\beta}^{TC} = 1$ and $\hat{\beta}^N = \hat{\beta} \in (\beta, 1]$.

Note that consumers do not take any action in $t = T$. We first analyze the switching decision in $t = T - 1$. Suppose that consumers bought the base product and kept using the firm's add-on. Then, consumers do not switch to the competitive fringe if and only if $-k_{T-1} + \beta^i a \leq \beta^i (a - p_T^a)$ or equivalently $p_T^a \leq \frac{k_{T-1}}{\beta^i}$. We next analyze consumer behavior in period $\tau < T - 1$. Consumers think that they will not switch in any future period if and only if $\hat{\beta}^i \sum_{i=t+1}^{T} p_i^a \leq k_t$ for all $t > \tau$. Given this belief, consumers' switching behavior in period τ can be divided into the following two cases. First, if $\hat{\beta}^i \sum_{i=t+1}^{T} p_i^a \leq k_t$ for all $t > \tau$, consumers do not switch in period τ if and only if $\beta^i \sum_{i=t+1}^{T} p_i^a \leq k_\tau$ because they think that they will never switch in any future period $t > \tau$. Second, if there exists a period $t > \tau$ such that $\hat{\beta}^i \sum_{i=t+1}^{T} p_i^a > k_t$, by backward induction consumers form a belief about whether they will switch or not in each future period, and as a result, they think they will switch in period $\hat{t} > \tau$. Given $\hat{t}$, they do not switch in period τ if and only if $k_\tau > \beta^i \left(k_{\hat{t}} + \sum_{i=\tau+1}^{\hat{t}} p_i^a \right)$.

Given these, each consumer buys the base product in $t = 0$ if and only if her perceived utility is equal to or greater than the outside option. We here explicitly describe the consumer behavior on the purchase of the base product in $t = 0$. Given the switching decisions regarding the add-on, each consumer takes up the base product in $t = 0$ if and only if one (or both) of the following two conditions is satisfied; (i) the total perceived utility of buying the base product and the add-on from the monopoly firm exceeds the outside option: $\beta^i[v + (T-1)a - p^v - \sum_{t=3}^{T} p_t^a] \geq \beta^i \bar{u}$, (ii) the total perceived utility of buying the base product and switching in period $\hat{t}$ exceeds the outside option for some $\hat{t} \in \{2, \cdots, T-1\}$: $\beta^i \left[v + (T-1)a - p^v - k_{\hat{t}} - \sum_{t=3}^{\hat{t}} p_t^a\right] \geq \beta^i \bar{u}$. Note that (i) is equivalent to $p^v \leq V_T - \sum_{t=3}^{T} p_t^a$, where $V_T := v + (T-1)a - \bar{u}$ denotes the total net consumption value of the product.

It is easy to show that the firm sells its add-on to some consumers in every period: $p_t^a \leq \frac{k_t}{\beta^i}$. It is also easy to show that if time-consistent consumers pay p_t^a, then naive consumers also pay p_t^a. From the above two participation constraints in $t = 0$, we can divide the firm's maximization problem into two cases: $p^v \leq V_T - \sum_{t=3}^{T} p_t^a$ and $V_T - \sum_{t=3}^{T} p_t^a < p^v \leq V_T - \min_s[k_s + \sum_{t=3}^{s} p_t^a]$. In the former case, it is optimal for the firm to sell the add-on to both naive and time-consistent consumers. In the latter case, the firm sells the add-on only to naive consumers from period $\hat{t}$ on.

We now analyze the optimal pricing of the firm. Note that if no consumer had an option to opt out of the add-on, the firm would set its total price equal to its overall consumption value minus the consumers' outside option, i.e., $p^v + \sum_{t=3}^{T} p_t^a = V_T$.

To complete the proof, we show two lemmas. We first investigate the situation in which switching costs are high in all periods, i.e., $k_t = \overline{k}$ for all $t \in \{1, \cdots, T-1\}$. This is the case when the policymaker does not employ any policy. The firm faces a trade-off between exploiting naive consumers with a high add-on price and selling its add-on to all consumers with a moderate add-on price. Note that the add-on prices can be different between periods. Lemma B.1 summarizes the result of the case:

Lemma B.1. *Suppose $k_t = \overline{k}$ for all $t \in \{1, \cdots, T-1\}$.*

If $\alpha + (T-3)(1-\beta)\alpha > \beta$, the firm sets $p^v = V_T - \overline{k}$, $p_t^a = \frac{1-\beta}{\beta}\overline{k}$ in $t \in \{3, \cdots, T-1\}$, and $p_T^a = \frac{1}{\beta}\overline{k}$. Time-consistent consumers switch before paying any add-on price, whereas naive consumers never switch. The firm's profits are $\pi = V_T - c^v + (\frac{\alpha}{\beta} - 1)\overline{k} + (T-3)\frac{\alpha}{\beta}(1-\beta)\overline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - (T-2)\frac{1-\beta}{\beta}\overline{k}$ and $u^{TC} = \bar{u}$.

If $\alpha + (T-3)(1-\beta)\alpha \leq \beta$, the firm sets $p^v + \sum_{t=3}^{T} p_t^a = V_T$ with $\sum_{t=3}^{T} p_t^a \leq \overline{k}$. No consumer switches. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

Proof. In what follows, we analyze a slightly more general case in which $k_t = k$ for all t. We divide the analysis into two cases.

First, suppose that the firm sells the add-on only to naive consumers. In this case, the maximal add-on price the firm can charge to naive consumers is $p_T^a = \frac{1}{\beta}k$ as we showed in Lemma 2.1. Given this, in period $t = T-2$ naive consumers prefer switching in the next period $t+1$ to switching in the current period t if and only if $\beta(-p_t^a - k) \geq -k$ or equivalently $p_t^a \leq \frac{1-\beta}{\beta}k$. Also, note that if the firm sets $p_t^a = \frac{1-\beta}{\beta}k$, then for any $\hat{\beta} > \beta$ naive consumers (wrongly) believe that they will switch in t. By recursively applying this argument for $t < T-2$, the firm sets $p^v = V_T - k$, $p_T^a = \frac{1}{\beta}k$, and $p_t^a = \frac{1-\beta}{\beta}k$ for all $t \in \{3, \cdots, T-1\}$. Naive consumers do not switch, whereas time-consistent consumers switch before paying any add-on price. The firm's profits are $\pi = V_T - c^v + (\frac{\alpha}{\beta} - 1)k + (T-3)\frac{\alpha}{\beta}(1-\beta)k$ and the consumers' long-run utilities are $u^N = \bar{u} - (T-2)\frac{1-\beta}{\beta}k$ and $u^{TC} = \bar{u}$.

Second, suppose that the firm sells the add-on to both naive and time-consistent consumers. In this case, the firm sets $p^v + \sum_{t=3}^T p_t^a = V_T$ with $\sum_{t=3}^T p_t^a \leq k$.[1] The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

By comparing the above two cases, we obtain the result. $\square$

We next analyze the situation in which the switching cost is lower only in the first period, i.e., $k_1 = \underline{k}$, $k_t = \overline{k}$ for all $t \in \{2, \cdots, T-1\}$. This is the case if the policymaker employs the policy when consumers are enrolled in the add-on service. Lemma B.2 summarizes the result of this case:

Lemma B.2. *Suppose $k_1 = \underline{k}, k_t = \overline{k}$ for all $t \in \{2, \cdots, T-1\}$.*

(i) Suppose $\Delta_k \geq \beta$. If $\alpha + (T-3)(1-\beta)\alpha > \beta\Delta_k$, the firm sets $p^v = V_T - \underline{k}$, $p_t^a = \frac{1-\beta}{\beta}\overline{k}$ in $t \in \{3, \cdots, T-1\}$, and $p_T^a = \frac{1}{\beta}\overline{k}$. Time-consistent consumers switch in period 1, whereas naive consumers never switch. The firm's profits are $\pi = V_T - c^v - \underline{k} + \frac{\alpha}{\beta}[1 + (T-3)(1-\beta)]\overline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} + \underline{k} - \frac{1}{\beta}[1 + (T-3)(1-\beta)]\overline{k}$ and $u^{TC} = \bar{u}$.

If $\alpha + (T-3)(1-\beta)\alpha \leq \beta\Delta_k$, the firm sets $p^v + \sum_{t=3}^T p_t^a = V_T$ with $\sum_{t=3}^T p_t^a \leq \underline{k}$. No consumer switches. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

[1] In addition to ($p^v = V_T - k$, $\sum_{t=3}^T p_t^a = k$), there are multiple equilibria for charging a higher p^v and a lower $\sum_{t=3}^T p_t^a$. We can pin down the equilibrium base-product price to $p^v = V_T - k$ by assuming that a tiny proportion of consumers exits the market at the end of $t = 1$ and cannot use the add-on. The same argument can be applied to the subsequent lemmas.

(ii) Suppose $\Delta_k < \beta$. If $\alpha > \beta$, the firm sets $p^v = V_T - \underline{k}$, $p_t^a = \frac{1-\beta}{\beta}\underline{k}$ in $t \in \{3, \cdots, T-1\}$, and $p_T^a = \frac{1}{\beta}\underline{k}$. Time-consistent consumers switch in period 1, whereas naive consumers never switch. The firm's profits are $\pi = V_T - c^v + (\frac{\alpha}{\beta} - 1)\underline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - \frac{1-\beta}{\beta}\underline{k}$ and $u^{TC} = \bar{u}$.

If $\alpha \leq \beta$, the firm sets $p^v + \sum_{t=3}^T p_t^a = V_T$ with $\sum_{t=3}^T p_t^a \leq \underline{k}$. No consumer switches. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

Proof. (i) Notice that time-consistent consumers do not switch in period 1 if and only if $\sum_{t=3}^T p_t^a \leq \underline{k}$. Because $-\underline{k} < -\beta\overline{k}$, naive consumers do not switch in period 1.

First, suppose that the firm sells the add-on only to naive consumers. In this case, the firm sets $p^v = V_T - \underline{k}$, $p_T^a = \frac{1}{\beta}\overline{k}$, and $p_t^a = \frac{1-\beta}{\beta}\overline{k}$ for all $t \in \{3, \cdots, T-1\}$ as in Lemma B.1. Naive consumers do not switch, whereas time-consistent consumers switch in period 1. The firm's profits are $\pi = V_T - c^v - \underline{k} + \frac{\alpha}{\beta}[1 + (T-3)(1-\beta)]\overline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} + \underline{k} - \frac{1}{\beta}[1 + (T-3)(1-\beta)]\overline{k}$ and $u^{TC} = \bar{u}$.

Second, suppose that the firm sells the add-on to both naive and time-consistent consumers. In this case, the firm sets $p^v + \sum_{t=3}^T p_t^a = V_T$ with $\sum_{t=3}^T p_t^a \leq \underline{k}$. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

By comparing the above two cases, we obtain the result.

(ii) Notice that time-consistent consumers do not switch in period 1 if and only if $\sum_{t=3}^T p_t^a \leq \underline{k}$. Naive consumers do not switch in period 1 if and only if $\beta \sum_{t=3}^T p_t^a \leq \underline{k}$, because given $-\underline{k} \geq -\beta\overline{k}$ naive consumers always prefer switching in period 1 to switching in any subsequent period.

First, suppose that the firm sells the add-on only to naive consumers. In this case, the firm sets $p^v = V_T - \underline{k}$ and $\sum_{t=3}^T p_t^a = \frac{1}{\beta}\underline{k}$. Naive consumers do not switch, whereas time-consistent consumers switch in period 1. The firm's profits are $\pi = V_T - c^v + (\frac{\alpha}{\beta} - 1)\underline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - \frac{1-\beta}{\beta}\underline{k}$ and $u^{TC} = \bar{u}$.

Second, suppose that the firm sells the add-on to both naive and time-consistent consumers. In this case, the firm sets $p^v + \sum_{t=3}^T p_t^a = V_T$ with $\sum_{t=3}^T p_t^a \leq \underline{k}$. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

By comparing the above two cases, we obtain the result. □

Comparing Lemmas B.1 and B.2 completes the proof of Proposition 2.4. □

Proof of Proposition 2.5.

Lemma B.3 summarizes the result of the alternative policy:

Lemma B.3. *Suppose $k_t = \underline{k}$ for any t which satisfies $p^a_{t+1} > p^a_t$ with $p^a_2 = 0$.*

If $\alpha + (T-3)(1-\beta)\alpha > \beta$, the firm sets $p^v = V_T - \underline{k}$, $p^a_t = \frac{1-\beta}{\beta}\underline{k}$ in $t \in \{3, \cdots, T-1\}$, and $p^a_T = \frac{1}{\beta}\underline{k}$. Time-consistent consumers switch before paying any add-on price, whereas naive consumers never switch. The firm's profits are $\pi = V_T - c^v + (\frac{\alpha}{\beta} - 1)\underline{k} + (T-3)\frac{\alpha}{\beta}(1-\beta)\underline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - (T-2)\frac{1-\beta}{\beta}\underline{k}$ and $u^{TC} = \bar{u}$.

If $\alpha + (T-3)(1-\beta)\alpha \leq \beta$, the firm sets $p^v + \sum_{t=3}^T p^a_t = V_T$ with $\sum_{t=3}^T p^a_t \leq \underline{k}$. No consumer switches. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

Proof. Note that consumer behavior in each case is described in the proof of Proposition 2.4. We divide the analysis into two cases.

We first analyze the case in which the firm sells the add-on to both naive and time-consistent consumers. In this case, the firm sets $p^v + \sum_{t=3}^T p^a_t = V_T$ with $\sum_{t=3}^T p^a_t \leq \underline{k}$. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

Second, suppose that the firm sells the add-on only to naive consumers. In this case, the firm sets $p^v = V_T - \underline{k}$, $p^a_T = \frac{1}{\beta}\underline{k}$, and $p^a_t = \frac{1-\beta}{\beta}\underline{k}$ for all $t \in \{3, \cdots, T-1\}$. The firm voluntarily reduces the switching cost to $\underline{k}$ in any period after the firm is forced to do so by the policy. To show this, suppose that $k_t = \underline{k}$. On the one hand, decreasing k_{t+1} makes naive consumers more likely to believe that they will switch in future, and hence makes them more likely to procrastinate their switching decision by relaxing the constraint of not switching in period t: $\underline{k} \geq \beta(p^a_{t+1} + k_{t+1})$ or equivalently $p^a_{t+1} \leq \frac{1}{\beta}k_t - k_{t+1}$. On the other hand, it tightens the constraint of not switching in period $t+1$: $k_{t+1} \geq \beta(p^a_{t+2} + k_{t+2})$. However, the latter constraint is not binding because under the policy the firm has to decrease its switching cost whenever charging a higher price. To see this, suppose that k_t was reduced due to the policy but the firm did not decrease k_{t+1} voluntarily (i.e., $k_t = \underline{k}, k_{t+1} = \overline{k}$). In such a case, the firm can charge at most $p^a_{t+1} \leq \frac{1}{\beta}\underline{k} - \overline{k}$, which is less than $\frac{1-\beta}{\beta}\underline{k}$. To make use of the relaxed constraint $k_{t+1} \geq \beta(p^a_{t+2} + k_{t+2})$, the firm would have to increase the price p^a_{t+2}. By doing so, however, the firm would also have to reduce k_{t+1} by the policy and can charge at most $p^a_{t+2} \leq \frac{1}{\beta}\underline{k} - k_{t+2} \leq \frac{1-\beta}{\beta}\underline{k}$. Hence, compared to the situation in which the firm sets the switching cost to $\underline{k}$ in any period after the policy is implemented, the firm cannot increase its profits by setting a higher switching cost under the policy. Given that, the firm charges a positive add-on price in $t = 3$, and then keeps the add-on prices constant with setting a low k. Naive consumers do not switch, whereas time-consistent consumers switch before paying any add-on price. The firm's profits are $\pi = V_T - c^v + (\frac{\alpha}{\beta} - 1)\underline{k} + (T-3)\frac{\alpha}{\beta}(1-\beta)\underline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - (T-2)\frac{1-\beta}{\beta}\underline{k}$ and $u^{TC} = \bar{u}$.

By comparing the above two cases, we obtain the result. □

Comparing Lemma B.1 and Lemma B.3 completes the proof of Proposition 2.5. □

Proof of Proposition 2.6.

Immediate from Lemma B.1 with $k_t = \underline{k}$ for all t and Lemma B.3. □

Proof of Proposition 2.7.

Suppose that a symmetric equilibrium exists in which firms earn positive profits. Then, each firm can profitably deviate by offering the same add-on price and a slightly lower base-product price, because the deviating firm can attract all consumers and each consumer's behavior regarding the add-on purchase does not change—a contradiction.

As firms make zero profits in equilibrium, the base-product price equals the production cost minus the total profits from the add-on. Similar to the analysis in Section 2.4.1, the outcomes are summarized as follows:

First, suppose that $k_1 \geq k_2$. If $\alpha > \beta$, there exists an equilibrium in which $p^v = c^v - \frac{\alpha}{\beta}k_2$ and $p_3^a = \frac{1}{\beta}k_2$. If $\alpha \leq \beta$, there exists an equilibrium in which $p^v = c^v - k_2$ and $p_3^a = k_2$.

Second, suppose that $k_1 = \underline{k} < \overline{k} = k_2$ and $\Delta_k \geq \beta$. If $\alpha > \beta\Delta_k$, there exists an equilibrium in which $p^v = c^v - \frac{\alpha}{\beta}\overline{k}$ and $p_3^a = \frac{1}{\beta}\overline{k}$. If $\alpha \leq \beta\Delta_k$, there exists an equilibrium in which $p^v = c^v - \underline{k}$ and $p_3^a = \underline{k}$.

Third, suppose that $k_1 = \underline{k} < \overline{k} = k_2$ and $\Delta_k < \beta$. If $\alpha > \beta$, there exists an equilibrium in which $p^v = c^v - \frac{\alpha}{\beta}\underline{k}$ and $p_3^a = \frac{1}{\beta}\underline{k}$. If $\alpha \leq \beta$, there exists an equilibrium in which $p^v = c^v - \underline{k}$ and $p_3^a = \underline{k}$.

By comparing the above three cases, we obtain the result. □

Proof of Proposition 2.8.

The argument in the proof of Proposition 2.7 implies that in any positive-profit equilibrium firms set $p^v = 0$. Also, if all consumers pay the add-on price, then the standard Bertrand-type price competition argument leads to $p^v + p_3^a = c^v$. However, if only naive consumers pay the add-on price, then the firms may be able to earn positive profits because of the constraint $p^v \geq 0$. To see this, consider a candidate equilibrium $p^v = 0$ and $p_3^a = \frac{1}{\beta}k_2$. If a firm deviates from the candidate equilibrium and charges the add-on price to both naive and time-consistent consumers, the deviating firm can charge a total payment of at most $p^v + p_3^a = \min\{k_1, k_2\}$ in order to attract these consumers. The analysis of the case in which $p^v = 0$ and $p_3^a = \frac{1}{\beta}k_1$ is a candidate equilibrium is the same. Similar to the previous analysis, the outcomes are summarized as follows:

First, suppose that the policy is enacted in $t = 2$ or when the policy is enacted in $t = 1$ and $\Delta_k < \beta$. If $\frac{1}{N}(\frac{\alpha}{\beta}\underline{k} - c^v) > \max\{\underline{k} - c^v, 0\}$, there exists a positive-profit equilibrium in which $p^v = 0$ and $p_3^a = \frac{1}{\beta}\underline{k}$. If $\frac{1}{N}(\frac{\alpha}{\beta}\underline{k} - c^v) \leq \max\{\underline{k} - c^v, 0\}$, there exists a zero-profit equilibrium in which $p^v + p_3^a = c^v$.

Second, suppose that the policy is enacted in $t = 1$ and $\Delta_k \geq \beta$. If $\frac{1}{N}(\frac{\alpha}{\beta}\overline{k} - c^v) > \max\{\underline{k} - c^v, 0\}$, there exists a positive-profit equilibrium in which $p^v = 0$ and $p_3^a = \frac{1}{\beta}\overline{k}$. If $\frac{1}{N}(\frac{\alpha}{\beta}\overline{k} - c^v) \leq \max\{\underline{k} - c^v, 0\}$, there exists a zero-profit equilibrium in which $p^v + p_3^a = c^v$.

By comparing the above two cases, we obtain the result. $\square$

B.2 Supplementary Material

In this supplementary material, we investigate further extensions of the model introduced in Section 2.4.1. First, we incorporate fully-sophisticated present-biased consumers into the model. Second, we analyze the model incorporating heterogeneous base-product demand or heterogeneous add-on demand. Finally, we investigate the effects and caveats of imposing deadlines.

B.2.1 Incorporating Sophisticated Consumers

We here analyze the model in which some present-biased consumers are perfectly aware of their self-control problems. Assume that a proportion $\alpha^s > 0$ of consumers is present-biased but is perfectly sophisticated about own present biasness: $\hat{\beta} = \beta < 1$. A proportion $\alpha^n > 0$ of consumers is naive present-biased and the remaining proportion $1-\alpha^s-\alpha^n$ is time consistent. For simplicity we focus on the case in which $T = 3$. The result is summarized as follows:

Proposition B.1. *Suppose that a proportion $\alpha^s > 0$ of consumers is sophisticated present-biased, a proportion $\alpha^n > 0$ of consumers is naive present-biased, and the remaining proportion $1 - \alpha^s - \alpha^n$ is time consistent.*

Then, a policy in $t = 1$ strictly decreases social welfare compared to no policy when $\frac{\alpha^n}{\beta} > \Delta_k \geq \beta$ and $\alpha^s(V_3 - c^v) \geq [\frac{\alpha^n}{\beta} - (1 - \alpha^s)]\overline{k}$. Under any parameters, a policy in $t = 2$ strictly increases consumer welfare and weakly increases social welfare compared to no policy.

Proposition B.1 shows that the firm may have a stronger incentive to charge a relatively low add-on price under the presence of fully-sophisticated present-biased consumers. Intuitively, if these consumers want to plan but cannot commit to switch, they foresee that they will fail to do so and therefore either do not buy the base product or cancel the contract earlier. Our implications regarding the timing of policies, however, remain unchanged.

Proof of Proposition B.1.

Note that actual consumer behavior in $t = 2$ does not change because $\hat{\beta}$ is not relevant to the consumer's actual decision in $t = 2$. In $t = 1$, sophisticated consumers think that they will not switch in $t = 2$ if and only if $p_3^a \leq k_2/\beta$. Conditional on this belief, these consumers' switching behavior in $t = 1$ can be divided into the following two cases. First, if $p_3^a \leq k_2/\beta$, consumers think they will not switch in $t = 2$; given this, they do not switch in $t = 1$ if and only if $p_3^a \leq k_1/\beta$. Second, if $p_3^a > k_2/\beta$, consumers think they will switch in $t = 2$; given this, they do not switch in $t = 1$ if and only if $k_1 > \beta k_2$. Note that the firm never sets $p_3^a > k_2/\beta$, because otherwise all consumers opt out of the firm's add-on.

Given the take up of the base product, sophisticated consumers opt out in $t = 1$ if $k_2 < p_3^a \leq k_2/\beta$ and $p_3^a < k_1/\beta$; otherwise, naive and sophisticated consumers receive the same ex-post utility. Selling the add-on only to naive and sophisticated consumers is not optimal, because sophisticated consumers would not buy the base product in $t = 0$ if their utility is less than $\bar{u}$ and hence, selling the add-on also to time-consistent consumers is more profitable. Note that when the firm sells the add-on to time-consistent consumers, then the firm can earn at most $\pi = V_3 - c^v$. Note also that whenever the firm sells only to naive consumers and sets $p_3^a = k_2/\beta$ in Section 2.4.1, naive consumers' long-run utility is less than $\bar{u}$. Since sophisticated consumers anticipate this, they do not buy the base product.

Under no policy, when the firm sells the add-on to all consumers, the profits are $\pi = V_3 - c^v$. When the firm sells the add-on with $p_3^a = \frac{1}{\beta}\overline{k}$, sophisticated consumers do not buy the base product in $t = 0$. The profits are $\pi = (1 - \alpha^s)(V_3 - c^v - \overline{k}) + \frac{\alpha^n}{\beta}\overline{k}$. Hence, the firm sells the add-on only to naive consumers if and only if $[\frac{\alpha^n}{\beta} - (1 - \alpha^s)]\overline{k} > \alpha^s(V_3 - c^v)$, and otherwise it sells the add-on to all consumers.

Suppose that there is a policy in $t = 1$ and $\Delta_k \geq \beta$. When the firm sells the add-on to all consumers, the maximal profits are $\pi = V_3 - c^v$. When the firm sells the add-on only to naive consumers, sophisticated consumers buy the base product and switch in $t = 1$. Hence the profits are at most $\pi = V_3 - c^v - \underline{k} + \frac{\alpha^n}{\beta}\overline{k}$. The firm sells the add-on only to naive consumers if and only if $\alpha^n > \beta\Delta_k$.

Suppose that there is a policy in $t = 1$ and $\Delta_k < \beta$. When the firm sells the add-on to all consumers, the maximal profits are $\pi = V_3 - c^v$. When the firm sells the add-on with $p_3^a = \frac{1}{\beta}\underline{k}$, sophisticated consumers do not buy the base product in $t = 0$. The profits are $\pi = (1 - \alpha^s)(V_3 - c^v - \underline{k}) + \frac{\alpha^n}{\beta}\underline{k}$. Hence, the firm sells the add-on only to naive consumers if and only if $[\frac{\alpha^n}{\beta} - (1 - \alpha^s)]\underline{k} > \alpha^s(V_3 - c^v)$, and otherwise it sells the add-on to all consumers.

Suppose that there is a policy in $t = 2$. When the firm sells the add-on to all consumers, the profits are $\pi = V_3 - c^v$. When the firm sells the add-on with $p_3^a = \frac{1}{\beta}\underline{k}$, sophisticated consumers do not buy the base product in $t = 0$. The profits are $\pi = (1 - \alpha^s)(V_3 - c^v - \underline{k}) + \frac{\alpha^n}{\beta}\underline{k}$. Hence, the firm sells the add-on only to naive consumers if and only if $[\frac{\alpha^n}{\beta} - (1 - \alpha^s)]\underline{k} > \alpha^s(V_3 - c^v)$, and otherwise it sells the add-on to all consumers. $\square$

B.2.2 Continuous Distributions

Here we analyze the case in which the valuation for the base product or the valuation for the add-on is heterogeneous. To illustrate the point in a simple manner, consider the case in which $T = 3$ and either v or a is uniformly distributed. With slightly abbreviating the notation, let $\tilde{v} = v + 2a - \bar{u}$ be uniformly distributed on $[\underline{v}, \overline{v}]$ (independent of whether consumers are time consistent or not). We denote the price for the base product when a policy is employed in period t by p^{vt} and the price for the base product when no policy is employed by p^{vn}.

The following proposition shows that imposing a policy may cause an increase in the base-product price. Similar to the logic of Grubb (2015), under downward-sloping base-product demand, even a policy in $t = 2$ may increase the equilibrium base-product price. Further, Proposition B.2 shows that when all consumers sufficiently value both the base product and the add-on, the results of incorporating heterogeneities are qualitatively the same:

Proposition B.2. *Suppose that the valuation for the base product or the valuation for the add-on is uniformly distributed and $a \geq \underline{k}$ and $\tilde{v} \geq c_v$ for all consumers.*

Then, the prices for the add-on are the same as in Section 2.4.1. Under any parameters, $p^{vn} \leq p^{v1}, p^{v2}$. If $1 \geq \frac{\alpha}{\beta} > \Delta_k \geq \beta$, then $p^{v1} < p^{v2}$; otherwise, $p^{v1} = p^{v2}$.

The intuition is as follows. As in a simple monopoly problem, a firm faces the trade-off between charging a high price for the base product (but only serving few consumers) and serving many consumers (but only making small profits per consumer). In addition to the profits with the base product, the firm makes a constant average profit per consumer from the add-on. If a policy reduces the average profit per consumer from the add-on, a higher number of consumers is less profitable for the firm, so the policy increases the price for the base product.

Note that an additional inefficiency can arise when $\tilde{v} + 2a < c_v$ for some naive consumers, but the firm sells the base product to these consumers below cost in order to enroll them in the add-on.

Also, if $a < \underline{k}$ for some consumers, these consumers would not take up the add-on of the competitive fringe after canceling the add-on from the firm (when we interpret $\underline{k}$ is the re-registration cost). As these consumers do not take up the add-on of the competitive fringe, under the policy in $t = 2$ they would not switch the firm's add-on if $p_3^a \leq a(< \underline{k})$. Hence, in this case the firm might want to reduce the price for the add-on.

Proof of Proposition B.2.

Let $V(\tilde{v})$ be the cumulative distribution function of $\tilde{v}$, which is differentiable for the relevant values of $\tilde{v}$ (for other values, the firm would either not sell at all or sell to all consumers, in which case we would be back to Section 2.4.1). Because the distribution of $\tilde{v}$ is independent of whether consumers are time consistent or not, it can be shown that the optimal prices for the add-on are the same as in Section 2.4.1. Note that consumers buy the base product if and only if their perceived utility in $t = 0$ is greater than or equal to $\bar{u}$. The firm maximizes the number of consumers who buy the product times the profit per consumer which is given in Section 2.4.1.

If no policy is implemented and $\alpha \leq \beta$, consumers buy the base product if and only if $\tilde{v} \geq p^{vn} + p_3^a$. The firm solves $\max_{p^{vn}}[1 - V(p^{vn} + p_3^a)](p^{vn} + p_3^a - c^v) \Rightarrow (p^{vn} + p_3^a - c^v)(-V'(p^{vn} + p_3^a)) + [1 - V(p^{vn} + p_3^a)] = 0 \Rightarrow p^{vn} + p_3^a = \frac{\overline{v}+c^v}{2}$. If $\alpha > \beta$, $p^{vn} = \frac{\overline{v}+c^v-\left(1+\frac{\alpha}{\beta}\right)\overline{k}}{2}$ and $p_3^a = \frac{1}{\beta}\overline{k}$.

If a policy in $t = 1$ is implemented, $\Delta_k \geq \beta$, and $\alpha > \beta\Delta_k$, then $p^{v1} = \frac{\overline{v}+c^v-\left(1+\frac{\alpha}{\beta\Delta_k}\right)\underline{k}}{2}$ and $p_3^a = \frac{1}{\beta}\overline{k}$. If $\Delta_k \geq \beta$ and $\alpha \leq \beta\Delta_k$, then $p^{v1} + p_3^a = \frac{\overline{v}+c^v}{2}$. If $\Delta_k < \beta$ and $\alpha > \beta$, then $p^{v1} = \frac{\overline{v}+c^v-\left(1+\frac{\alpha}{\beta}\right)\underline{k}}{2}$ and $p_3^a = \frac{1}{\beta}\overline{k}$. If $\Delta_k < \beta$ and $\alpha \leq \beta$, then $p^{v1} + p_3^a = \frac{\underline{v}+c^v}{2}$.

If a policy in $t = 2$ is implemented and $\alpha > \beta$, $p^{v2} = \frac{\overline{v}+c^v-\left(1+\frac{\alpha}{\beta}\right)\underline{k}}{2}$ and $p_3^a = \frac{1}{\beta}\underline{k}$. If $\alpha \leq \beta$, $p^{v2} + p_3^a = \frac{\overline{v}+c^v}{2}$.

By comparing the above cases, we obtain the result. $\square$

B.2.3 Deadlines

So far, we have shown the consequences of policies which decrease consumers' switching costs in certain periods. Now we examine an alternative policy intervention. Specifically, in this subsection we analyze an extended model in which the policymaker can sufficiently increase switching costs in certain periods so that consumers cannot cancel their contract in those periods. By doing so, the policymaker can impose a deadline of switching decisions to consumers. In our illustrative model, it is optimal for the policymaker to prevent consumers from switching after the first two periods:

Proposition B.3. *Assume that the policymaker can prohibit consumers to switch in certain periods: the policymaker chooses* $\mathfrak{T} \subseteq \{2, \cdots, T-1\}$ *such that* $k_t = \infty$ *for all* $t \in \mathfrak{T}$*. Then, choosing* $\mathfrak{T}^* = \{3, \cdots, T-1\}$ *maximizes consumer and social welfare. If* $\alpha + (T-3)(1-\beta)\alpha > \beta$*, consumer welfare is strictly larger than choosing* $\mathfrak{T} = \emptyset$*. If in addition* $\beta > \alpha$*, social welfare is also strictly larger than choosing* $\mathfrak{T} = \emptyset$*.*

The intuition of Proposition B.3 is as follows: since consumers are not able to cancel after the second period under the policy, naive consumers cannot falsely believe that they will switch in a future period.[2] Given this, they will not procrastinate their switching decision if the total future payment for the add-on is high. This finding is in line with the theoretical literature which analyzes the effects of imposing deadlines (O'Donoghue and Rabin 1999b, Herweg and Müller 2011).

Proposition B.3 implies that $(T+1)$-period models can be reduced to the four-period model when the policymaker can impose an optimal deadline. In such a case, p_3^a is interpreted as the sum of all payments that have to be made after the second period. If the policymaker cannot regulate the prices directly but can change the switching costs, decreasing the switching cost in the second period and imposing the deadline in the second period maximize social welfare in our basic model. Namely, when the policymaker cannot regulate the prices directly, imposing $k_2 = \underline{k}$ and $\mathfrak{T}^* = \{3, \cdots, T-1\}$ becomes the optimal policy.[3]

Unlike a choice-enhancing or an active-choice policy which decreases the switching costs, however, one should be very cautious about imposing such a strict deadline in practice. For example, imposing a deadline may decrease welfare if add-on values or switching costs are changing over time. Also, as we discussed in Section 2.5, imposing a deadline can be welfare harmful when consumers also have other psychological biases.

Furthermore, imposing a deadline might not be feasible as the firm might be able to circumvent the deadline by (pretendedly) changing the product features of the add-on such that consumers receive extraordinary termination rights. Corollary B.1 states that the firm indeed has an incentive to do so:

[2] Note that imposing a deadline in $t = 1$ is also optimal, although the deadline in $t = 1$ does not seem to be legal in practice because consumers who get a free trial do not have an option to cancel the service later.

[3] Note that when the policymaker can regulate the prices directly, simply imposing $p_t^a = 0$ for all t maximizes the social welfare.

Corollary B.1. *Assume that the policymaker can prohibit consumers to switch in certain periods: she can choose* $\mathfrak{T} \subseteq \{2, \cdots, T-1\}$ *such that* $k_t = \infty$ *for all* $t \in \mathfrak{T}$*. If* $\alpha + (T-3)(1-\beta)\alpha > \beta$*, then profits of the firm under* $\mathfrak{T} = \emptyset$ *are strictly higher than under* $\mathfrak{T} = \{3, \cdots, T-1\}$*.*

Consequently, if the firm can credibly commit to pretendedly change its terms and conditions of the add-on, the deadline policy in Proposition B.3 may not be effective and the policymaker can only force the firm to lower the switching cost in certain periods.

Proof of Proposition B.3.

Note first that time-consistent consumers' utility is not affected by the policy and is always $u^{TC} = \bar{u}$.

Let $\underline{t}$ and $\bar{t}$ be the first and the last period of a sequence of periods such that $k_t = \infty$ for all $t \in \{\underline{t}, \cdots, \bar{t}\}$. Then, naive consumers do not switch in period $\underline{t} - 1$ if and only if $k_{\underline{t}-1} \geq \beta \left(\sum_{t=\underline{t}}^{\tau} p_t^a + k_\tau \right)$ for some $\tau \in \{\bar{t}+1, \cdots, T-1\}$ or $k_{\underline{t}-1} \geq \beta \left(\sum_{t=\underline{t}}^{T} p_t^a \right)$. So the maximum total payment the firm can charge (weakly) decreases when the number of periods in which consumers cannot switch increases. Charging lower prices potentially benefits naive consumers and potentially increases social welfare when time-consistent consumers do not switch anymore.

Analogous to Lemma B.1 with $T = 3$, if $t = 2$ is the last period in which a consumer can cancel the contract and if $\alpha > \beta$, then the firm sets $p^v = V_T - \overline{k}$ and $\sum_{t=3}^{T} p_t^a = \frac{1}{\beta}\overline{k}$. Naive consumers do not switch, whereas time-consistent consumers switch either in period 1 or period 2. The firm's profits are $\pi = V_T - c^v + (\frac{\alpha}{\beta} - 1)\overline{k}$ and the consumers' long-run utilities are $u^N = \bar{u} - \frac{1-\beta}{\beta}\overline{k}$ and $u^{TC} = \bar{u}$. If $t = 2$ is the last period in which a consumer can cancel the contract and if $\alpha \leq \beta$, then the firm sets $p^v + \sum_{t=3}^{T} p_t^a = V_T$. The firm's profits are $\pi = V_T - c^v$ and the consumers' long-run utilities are $u^N = u^{TC} = \bar{u}$.

Comparing this to Lemma B.1 delivers the result. □

Appendix C

Appendix for Chapter 3.

C.1 Preliminaries: Simplifying the Problem

First, we derive a number of preliminary results that help us to prove our main results.

Lemma C.1. *Without loss of generality, the following holds for all histories:*

- $U^v\left(h^t \cup \{d_t = 0\}\right) = \frac{\overline{u}}{1-\delta}$
- *If the agent selects effort* $\tilde{e}_t \neq e_t^v$ *in the virtual contract, then* $b^v(h^t, \tilde{e}_t) = 0$ *and* $U^v\left(h^t \cup \{\tilde{e}_t\}\right) = \frac{\overline{u}}{1-\delta}$
- *If the agent selects* $\tilde{e}_t \neq e_t^r$ *in the real contract, then* $b^r(h^t, \tilde{e}_t) = 0$ *and* $U^v\left(h^t \cup \{\tilde{e}_t\}\right) = \frac{\overline{u}}{1-\delta}$.

Proof. The optimality of $b^v(h^t, \tilde{e}_t) = 0$ for $\tilde{e}_t \neq e_t^v$ and $b^r(h^t, \tilde{e}_t) = 0$ for $\tilde{e}_t \neq e_t^r$ is straightforward because lower values of $b^r(h^t, \tilde{e}_t)$ and $b^v(h^t, \tilde{e}_t)$ relax (IC) constraints.

Concerning continuation values after the agent either chose $d_t = 0$ or did not exert the equilibrium effort level, assume that, for example, there is a history h^t with $U^v\left(h^t \cup \{d_t = 0\}\right) > \frac{\overline{u}}{1-\delta}$. Replace the contract following the history $h^t \cup \{d_t = 0\}$ with the following stationary contract for all subsequent histories: $e^r = e^v = e^{FB}$, $w^r = w^v = \overline{u}$, $b^r = b^v = c(e^{FB})$. This relaxes (IR) and (IC) constraints for history h^t, and keeps all off-equilibrium constraints following the history $h^t \cup \{d_t = 0\}$ satisfied. The same can be done for all other cases. $\square$

The next Lemma also simplifies the analysis.

Lemma C.2. *Without loss of generality, we can set* $w^r(h^t) = w^v(h^t) = 0$ *for all histories* h^t.

Proof. Assume $w^v(h^t) > 0$. Reducing $w^v(h^t)$ by ε and increasing $b^v(h^t)$ by $\varepsilon > 0$ does not tighten any constraint, but relaxes (vIC). Assume $w^r(h^t) > 0$. Reducing $w^r(h^t)$ by $\varepsilon > 0$ and increasing $b^r(h^t)$ by ε does not tighten any constraint but relaxes (rIC). $\square$

Hence, the constraints (rIRA) and (rIC), and (vIRA) and (vIC), respectively, are identical, allowing us to omit (rIC) and (vIC).

Now, the remaining constraints are

$$b^r(h^t, e^r_t) - c(e^r_t) + \beta\delta U^v\left(h^t \cup \{d_t = 1\}\right) \geq \overline{u} + \beta\delta\frac{\overline{u}}{1-\delta}. \tag{rIRA}$$

$$U^v\left(h^t \cup \{d_t = 1\}\right) \geq \frac{\overline{u}}{1-\delta} \tag{vIRA}$$

$$\Pi^{rv}(h^t) \geq \frac{\overline{\pi}}{1-\delta} \tag{vrIRP}$$

$$\Pi^{v}(h^t) \geq \frac{\overline{\pi}}{1-\delta} \tag{vIRP}$$

$$b^r(h^t, e^r_t) - c(e^r_t) + \beta\delta U^v\left(h^t \cup \left\{\hat{\imath}_t = r\right\}\right) \geq b^v(h^t, e^v_t) - c(e^v) + \beta\delta U^v\left(h^t \cup \left\{\hat{\imath}_t = v\right\}\right) \tag{rC}$$

$$b^v(h^t, e^v_t) - c(e^v) + \delta U^v\left(h^t \cup \left\{\hat{\imath}_t = v\right\}\right) \geq b^r(h^t, e^r_t) - c(e^r) + \delta U^v\left(h^t \cup \left\{\hat{\imath}_t = r\right\}\right) \tag{vC}$$

In the next two Lemmas, we derive the structure of $\boldsymbol{C}$, hence prove 3.1 and show that the real contract is stationary, as well as the virtual contract with the exception of the first period where it is chosen.

Lemma C.3. *$b^r(h^t, e^r) - c(e^r_t)$ is the same for all histories h^t.*

Proof. Assume there are two histories $h^{\overline{t}}$ and $h^{\underline{t}}$ such that the agent's real net payoff differs for both histories, and without loss of generality assume that $b^r(h^{\overline{t}}, e^r_{\overline{t}}) - c(e^r_{\overline{t}}) > b^r(h^{\underline{t}}, e^r_{\underline{t}}) - c(e^r_{\underline{t}})$. Then, replace $b^r(h^{\overline{t}}, e^r_{\overline{t}}) - c(e^r_{\overline{t}})$ by $b^r(h^{\underline{t}}, e^r_{\underline{t}}) - c(e^r_{\underline{t}})$, as well as the virtual contract following the history $h^{\overline{t}}$ by the virtual contract following the history $h^{\underline{t}}$. This increases the principal's real profits, without affecting any constraint for any other history – also not in earlier periods, because there the agent does not expect to choose the real contract in any future period. Hence, the agent's incentives in earlier periods are not affected by what happens if he chooses the real contract in a future periods. $\square$

Therefore, the real contract is history-independent, and we can omit dependence on time and histories when describing the elements of C^r.

In a next step, we show that the virtual contract can be independent of calendar time in a sense that without loss of generality, its components are only contingent on the number of subsequent previous periods in which the virtual contract has been chosen. Denote C^v_τ as the τ's subsequent period where the virtual contract has been chosen (hence, $\tau \geq 1$), independent of the remaining components of the history of the game.

Lemma C.4. *Without loss of generality, the virtual contract $C^v(h^t)$ is of the form C^v_τ for all histories of the game. Furthermore, it is stationary for all $\tau \geq 2$.*

Proof. First, we show that the agent's expected continuation utility after chosing the real contract in any period, $U^v\left(h^t \cup \left\{\hat{i}_t = r\right\}\right)$ or, if using the form C^v_τ, $U^v_1(h^t)$, is the same for all histories. To the contrary, assume there are two equilibrium histories $h^{\bar{t}}$ and $h^{\underline{t}}$ where the agent's respective future virtual utility differs, i.e., assume $U^v_1(h^{\bar{t}}) \neq U^v_1(h^{\underline{t}})$ and without loss of generality $U^v_1(h^{\bar{t}}) > U^v_1(h^{\underline{t}})$. Replace all components of the history following $h^{\underline{t}}$ with the components of the history following $h^{\bar{t}}$. Then, all constraints following $h^{\bar{t}}$ remain satisfied, and $b^r - c(e^r)$ can be reduced (before, rIRA was slack at the history $h^{\bar{t}}$).

Hence, $U^v_1(h^t)$ is the same for all histories h^t, i.e., we can write U^v_1. To prove that the virtual contract can be stationary for all $\tau > 2$, note that the respective elements of C^v_τ are only relevant for the constraints (vIRA), (vIRP) and (vC). Now take the history $h^{\tilde{t}}$ where the value $U^v_2(\cdot)$ assumes its highest value. If there are histories with a lower U^v_2, replace the respective elements of the virtual contract with ones determining $U^v_2(h^{\tilde{t}})$, implying that also $U^v_2(\cdot)$ can be independent of calendar time. Finally, take the per-period value of U^v_2, $(1-\delta)U^v_2$, and set all per-period utilities for the virtual contract for $\tau \geq 2$ equal to $(1-\delta)U^v_2$. This is clearly feasible and violates no constraint. □

These two Lemmas also prove Lemma 3.1. Hence, in the following, we omit time- and history-dependence when describing the elements of C^r. Regarding C^v_τ, we also omit dependence on τ for all $\tau \geq 2$. For $\tau = 1$, we keep the subscript "1". Therefore, the remaining constraints are

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b^v_1 - c(e^v_1) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right] \geq 0 \tag{rIRA}$$

$$(b^v_1 - c(e^v_1) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u}) \geq 0 \tag{v1IRA}$$

$$(b^v - c(e^v) - \overline{u}) \geq 0 \tag{vIRA}$$

$$e^r\theta - b^r + \delta\left[(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v)\right] \geq \frac{\overline{\pi}}{1-\delta} \tag{vrIRP}$$

$$(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v) \geq \frac{\overline{\pi}}{1-\delta} \tag{v1IRP}$$

$$e^v\theta - b^v \geq \overline{\pi} \tag{vIRP}$$

$$\begin{aligned} -c(e^r) + b^r - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right] \\ \geq (b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta\beta}{1-\delta}(b^v - c(e^v) - \overline{u}) \end{aligned} \tag{rC}$$

$$\begin{aligned} (b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}(b^v - c(e^v) - \overline{u}) \\ \geq (b^r - c(e^r) - \overline{u}) + \delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}(b^v - c(e^v) - \overline{u})\right] \end{aligned} \tag{vC}$$

Note that we added $\overline{u}\left(1 + \frac{\beta\delta}{1-\delta}\right)$ on both sides of (rC) and $\frac{\overline{u}}{1-\delta}$ on both sides of (vC).

In a next step, we prove Proposition 3.1, which can be rephrased as

Proposition C.1. *If the agent is naive and has $\beta \in (0,1)$, then in the profit-maximizing menu of contracts, $-c(e^r) + b^r < \overline{u}$.*

Proof. First, assume $-c(e^r) + b^r > \overline{u}$. Then, (vC) and (v1IRA) imply that $(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{(1-\delta)}(b^v - c(e^v) - \overline{u})$ must be strictly positive as well. Change $\boldsymbol{C}$ in the following way: Set $b^r - c(e^r) - \overline{u} = b_1^v - c(e_1^v) - \overline{u} = (b^v - c(e^v) - \overline{u}) = 0$, which leaves all constraints satisfied and increases the principal's profits. Note that these considerations already allow us to omit (vrIRP), (v1IRA) and (vC).

Now, assume that $-c(e^r)+b^r = \overline{u}$. Change $\boldsymbol{C}$ in the following way: First, set $(b_1^v - c(e_1^v) - \overline{u}) = (b^v - c(e^v) - \overline{u}) = 0$ for given effort levels, which satisfies all constraints. Then, reduce b_1^v by ε and increase b^v by $\varepsilon\frac{1-\delta}{\delta\beta}$. This increases $b_1^v + \frac{\delta}{1-\delta}b^v$ and hence relaxes (rIRA), (v1IRA), (vIRA) and (rC) (and does not violate limited liability as well as (v1IRP) and (vIRP) constraints for ε sufficiently small), and therefore allows the principal to reduce b^r. $\square$

$-c(e^r)+b^r < \overline{u}$ immediately implies that, given (rIRA), the (v1IRA) constraint automatically holds and can be omitted. The same is true for (vC). Furthermore, the (vrIRP) constraint can be omitted: $e^r\theta - b^r - \overline{\pi}$ will be strictly positive in a profit-maximizing equilibrium, and (v1IRP) yields $(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v) \geq \frac{\overline{\pi}}{1-\delta}$.

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right] \geq 0 \tag{rIRA}$$

$$(b^v - c(e^v) - \overline{u}) \geq 0 \tag{vIRA}$$

$$(e_1^v\theta - b_1^v) + \frac{\delta}{1-\delta}(e^v\theta - b^v) \geq \frac{\overline{\pi}}{1-\delta} \tag{v1IRP}$$

$$e^v\theta - b^v \geq \overline{\pi} \tag{vIRP}$$

$$\begin{aligned}-c(e^r) + b^r - \overline{u} + \beta\delta\left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right]\\ \geq (b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta\beta}{1-\delta}(b^v - c(e^v) - \overline{u})\end{aligned} \tag{rC}$$

Lemma C.5. *(vIRA) and (v1IRP) constraints are slack and can hence be omitted in profit-maximizing equilibrium.*

Proof. First, assume that (vIRA) binds. Increasing b^v by ε and reducing $e_1^v\theta - b_1^v$ by $\frac{\delta}{1-\delta}\varepsilon$ keeps all constraints unaffected with the exception of (vIRA) and (rC) which are relaxed.

Concerning (v1IRP), first note that if it binds, the same has to be true for (vIRP). Otherwise, we could reduce b_1^v by ε (if (vIRP) is slack but (v1IRP) binds, $b_1^v > 0$ for sure) and increase b^v by $\varepsilon\frac{1-\delta}{\delta}$. This would not affect (v1IRP) and (rIR), but relax (rC).

Now, assume that (v1IRP) and (vIRI) bind. This implies that $b_1^v > 0$, that the agent gets the whole virtual surplus, and constraints are:

$$b^r - c(e^r) - \overline{u} + \beta\delta\left[(e_1^v\theta - c(e_1^v) - \overline{\pi} - \overline{u}) + \frac{\delta}{1-\delta}(e^v\theta - c(e^v) - \overline{\pi} - \overline{u})\right] \geq 0 \tag{rIR}$$

$$-c(e^r) + b^r - \overline{u} + \beta\delta \left[(e_1^v\theta - c(e_1^v) - \overline{\pi} - \overline{u}) + \frac{\delta}{1-\delta}(e^v\theta - c(e^v) - \overline{\pi} - \overline{u}) \right]$$
$$\geq (e_1^v\theta - c(e_1^v) - \overline{\pi} - \overline{u}) + \frac{\delta\beta}{1-\delta}(e^v\theta - c(e^v) - \overline{\pi} - \overline{u}) \qquad \text{(rC)}$$

There, the right hand side of (rC) is positive, hence that (rIR) is slack. Therefore, a slight reduction of b_1^v, accompanied with a reduction of b^r or an increase of e^r (to keep (rC) unaffected) would keep (rIR) satisfied and increase the principal's profits. □

C.2 Proofs of Lemmas and Propositions from the Main Part

Proofs of Lemma 3.2 and Proposition 3.2.

Given the remaining constraints, the principal's maximization problem gives rise to the following Lagrange function:

$$\begin{aligned}
L =& \frac{e^r\theta - b^r}{1-\delta} + \lambda_{rIR} \left[b^r - c(e^r) - \overline{u} + \beta\delta \left[(b_1^v - c(e_1^v) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u}) \right] \right] \\
&+ \lambda_{rC} \left[-c(e^r) + b^r - \overline{u} - (b_1^v - c(e_1^v) - \overline{u})(1-\beta\delta) - \delta\beta (b^v - c(e^v) - \overline{u}) \right] \\
&+ \lambda_{vIRP} (e^v\theta - b^v - \overline{\pi}),
\end{aligned}$$

with first-order conditions

$$\begin{aligned}
\frac{\partial L}{\partial b^r} &= \frac{-1}{1-\delta} + \lambda_{rIR} + \lambda_{rC} = 0 \\
\frac{\partial L}{\partial e^r} &= \frac{\theta}{1-\delta} - c(e^r)'(\lambda_{rIR} + \lambda_{rC}) = 0 \\
\frac{\partial L}{\partial (b_1^v - c(e_1^v))} &= \lambda_{rIR}\beta\delta - \lambda_{rC}(1-\beta\delta) = 0 \\
\frac{\partial L}{\partial b^v} &= \lambda_{rIR}\beta\delta\frac{\delta}{1-\delta} - \lambda_{rC}\delta\beta - \lambda_{vIRP} = 0 \\
\frac{\partial L}{\partial e^v} &= -\lambda_{rIR}\beta\delta\frac{\delta}{1-\delta}c(e^v)' + \lambda_{rC}\delta\beta c(e^v)' + \lambda_{vIRP}\theta = 0.
\end{aligned}$$

Hence, $\lambda_{rIR} = \frac{1}{1-\delta} - \lambda_{rC}$ and $\theta - c(e^r)' = 0$, i.e., $e^r = e^{FB}$. Rearranging these conditions further yields $\lambda_{rC} = \lambda_{rIR}\frac{\beta\delta}{(1-\beta\delta)}$, $\lambda_{vIRP} = \lambda_{rIR}\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}$, and

$\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}\left(\theta - c(e^v)'\right) = 0$, i.e., $e^v = e^{FB}$.

Hence, (rC), (rIR) and (vIRP) can only bind simultaneously, which also implies that all of them must bind. To the contrary, assume that (rC) and (rIR) do not bind. Then, b^r can be further reduced until one of them binds, further increasing the principal's profits. Using these results gives the values for b^r, u^r and π^r. □

Proof of Proposition 3.3.

We solve this by backward induction: In the last period, there are no future periods left, so the agent is just compensated for first-best effort: $e^r_T = e^{FB}$, $b^r_T = c(e^{FB})$. In the second to last period, the agent cannot be fooled regarding the last period as he will end up choosing the contract giving him the highest utility. So again, the agent is just compensated for first-best effort: $e^r_{T-1} = e^{FB}$, $b^r_{T-1} = c(e^{FB})$. In $T-2$, the agent can be fooled by offering him a virtual contract in which he earns the whole production in the last period if he has earned less (or has worked harder) in the second to last period: $e^v_T = e^v_{T-1} = e^{FB}$, $b^v_{T-1} = c(e^{FB}) + \overline{u} - \beta\delta\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$, $b^v_T = e^{FB}\theta - \overline{\pi}$, where b^v_{T-1} makes the $T-1$ agent just indifferent between the virtual and the real contract. If the virtual contract above is promised to him, he is willing to work below his outside option in $T-2$: $e^r_{T-2} = e^{FB}$, $b^r_{T-2} = c(e^{FB}) + \overline{u} - \beta(1-\beta)\delta^2\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. Analogously, in earlier periods the promised virtual contract will always promise the whole surplus to the agent from the period after the entry period on, so the entry period bonus that makes the agent indifferent is given by $b^v_t = c(e^{FB}) + \overline{u} - \beta\delta\sum_{j=0}^{T-t-1}\delta^j\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. This in turn allows the principal to exploit the agent by setting $b^r_t = c(e^r_t) + \overline{u} - \beta(1-\beta)\delta^2\sum_{j=0}^{T-t-2}\delta^j\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. □

Proof of Proposition 3.4.

Note that none of the steps we performed to simplify the original problem is affected by the presence of a limited liability constraint. Hence, relevant constraints are the same, and the Lagrange function becomes:

$$\begin{aligned} L =& \frac{e^r\theta - b^r}{1-\delta} + \lambda_{rIR}\left[b^r - c(e^r) - \overline{u} + \beta\delta\left[(b^v_1 - c(e^v_1) - \overline{u}) + \frac{\delta}{1-\delta}(b^v - c(e^v) - \overline{u})\right]\right] \\ &+ \lambda_{rC}\left[-c(e^r) + b^r - \overline{u} - (b^v_1 - c(e^v_1) - \overline{u})(1-\beta\delta) - \delta\beta\left(b^v - c(e^v) - \overline{u}\right)\right] + \lambda_{b^r}b^r \\ &+ \lambda_{vIRP}\left(e^v\theta - b^v - \overline{\pi}\right), \end{aligned}$$

with first-order conditions:

$$\begin{aligned}
\frac{\partial L}{\partial b^r} &= \frac{-1}{1-\delta} + \lambda_{rIR} + \lambda_{rC} + \lambda_{b^r} = 0 \\
\frac{\partial L}{\partial e^r} &= \frac{\theta}{1-\delta} - c(e^r)'\left(\lambda_{rIR} + \lambda_{rC}\right) = 0 \\
\frac{\partial L}{\partial\left(b_1^v - c(e_1^v)\right)} &= \lambda_{rIR}\beta\delta - \lambda_{rC}\left(1-\beta\delta\right) = 0 \\
\frac{\partial L}{\partial b^v} &= \lambda_{rIR}\beta\delta\frac{\delta}{1-\delta} - \lambda_{rC}\delta\beta - \lambda_{vIRP} = 0 \\
\frac{\partial L}{\partial e^v} &= -\lambda_{rIR}\beta\delta\frac{\delta}{1-\delta}c(e^v)' + \lambda_{rC}\delta\beta c(e^v)' + \lambda_{vIRP}\theta = 0
\end{aligned}$$

Hence, $\lambda_{rIR} = \frac{1}{1-\delta} - \lambda_{rC} - \lambda_{b^r}$ and $\frac{\theta - c(e^r)'}{1-\delta} + c(e^r)'\lambda_{b^r} = 0$. Rearranging further yields $\lambda_{rC} = \lambda_{rIR}\frac{\beta\delta}{(1-\beta\delta)}$, $\lambda_{vIRP} = \lambda_{rIR}\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}$, and $\frac{\beta\delta^2(1-\beta)}{(1-\delta)(1-\beta\delta)}\left(\theta - c(e^v)'\right) = 0$. Furthermore, (rC), (rIR) and (vIRP) all bind simultaneously, which follows from the same arguments as in the prove to Proposition 3.2.

Hence, if $b^r = 0$, e^r is above the level given by $\theta - c(e^r)' = 0$. Plugging binding constraints into utilities gives the desired values. □

Proof of Proposition 3.5.

We show that (vC) does not bind.

Ignore the non-negativity constraints. (rIRP) binds. If not, one could increase b^v by ε and decrease b_1^v by $\frac{\delta\hat{\beta}}{1-\delta\hat{\beta}}\epsilon$. Then one could decrease b^r slightly without violating any constraint and thereby increasing the principal's profits.

(rC) binds: Note that (vC) can only bind simultaneously if $\beta = \hat{\beta}$. Assume $\beta < \hat{\beta}$ and (rC) does not bind. Then one could decrease b^r by ε (or increase e^r accordingly) and increase b_1^v by $\frac{\varepsilon}{\beta\delta}$ without violating any constraint, but increasing the principal's profits.

(rIR) binds: Otherwise one could decrease both b^r and b_1^v without violating any constraint, but increasing the principal's profits. □

Proof of Proposition 3.6.

We will first approach the solution to the principal's screening problem under the assumption that offering one menu of contracts for each agent is optimal. Then we will show that offering one menu of contracts for each agent is indeed optimal.

Separation by Menu Choice Assume the principal wants the agents to choose different contracts. To develop an idea about the structure of these contracts, take $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$, the profit-maximizing contracts derived in the main part, with the slight modification that the virtual contract is only offered from $t = 1$ on. Hence, agents can only choose the real contract in $t = 0$. We will discuss how to optimally modify these contracts below.

Assume the principal offered $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$. Then agent 2's expected utility level when choosing $\boldsymbol{C}_1$ was

$\tilde{U}_2^r = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1) > 0$. This is positive because of $\beta_2 > \beta_1$, hence agent 2 puts more weight on future utilities than agent 1 does, and since both agent did not expect to get a rent before. Furthermore, because the respective (rC) constraints have been binding before, agent 2 would actually go for the virtual contract.

When choosing $\boldsymbol{C}_2$ (and expecting to select the virtual contract in the future), agent 1 gets:

$\tilde{U}_1^r = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_1-\beta_2) < 0$. Hence, agent 1 would stick to $\boldsymbol{C}_1$.

Separation by menu choice involves giving at least one agent an expected rent (which will also materialize in a real rent compared to the case with symmetric information). The principal could either adjust $\boldsymbol{C}_2$ in a way that it becomes more attractive for agent 2 (without making it too attractive for agent 1), or adjust $\boldsymbol{C}_1$ in a way that it becomes less attractive for agent 2.

First, note that the principal is restricted in increasing 2's virtual surplus - simply because this is already made as attractive as feasible for agent 2, with the exception of the first period where it is expected to be chosen. Hence, the principal has the following opportunities to make $\boldsymbol{C}_2$ more attractive: She can include an additional payment in period $t = 0$ (which only consists of the real contract), which we denote X_2. Alternatively, she can increase the payments in the (first period of the) virtual contract in $t = 1$. Since she still wants agent 2 to actually choose the real contract in $t = 1$ and since the (rC) constraint has been binding before, she must increase the payments in the real contract in period $t = 1$ by the same amount. This amount is denoted by Y_2. Offering agent 2 a contract where she chooses C_2 but then takes the virtual contract is dominated by this contract. Then, the agent would still receive Y_2, but additionally capture the rents from later periods in the virtual contract.

Finally, the principal could reduce agent 1's payoff from the virtual contract and instead increase the real payoff he receives in period $t = 0$. We denote this amount by Z_1. More precisely, an increase of 1's real contract by Z_1 goes hand in hand with a reduction of his virtual payments by an amount Z_{1v} (in order to keep the rIR constraint for agent 1 binding) and potentially with an increase of his real payments in later periods (in order to keep the (rC) constraint for agent 1 binding). Note that a decrease of the virtual surplus in $t = 1$ does not only affect the contract in $t = 1$, but also limits what contracts the principal can offer in any later period. The principal cannot always simply decrease the payment of the virtual contract starting in the next period. If Z_{1v} is large and only the payment of the virtual contract starting in the next period was reduced, the agent would plan to choose the real contract in the next period and to choose the virtual contract only in the period after that (which would violate (vC)). Reducing the payment of the real contract is not an option, because the agent would eventually rather quit than choose the real contract. So the principal has to decrease the payment of the later virtual contract and even increase the next period's real contract in order to make the agent choose the real contract in the next period. If the necessary reduction of the payment in the later virtual contract is large, even later contracts might have to be changed by the same logic. The larger Z_{1v}, the more later periods are affected. When we talk about Z_1, we mean the full set of these adjustments. However, we first assume that any costs of these additional adjustments after $t = 0$ are zero, solve the simplified problem, and take the actual costs into account thereafter.

Now, expected payoffs when choosing the intended contracts and when all other components remain unchanged are

$U_1^r = Z_1 - \beta_1 \delta Z_{1v} = 0$ (hence $Z_{1v} = \frac{Z_1}{\beta_1 \delta}$) and $U_2^r = X_2 + \beta_2 \delta Y_2$.

When deviating and selecting the other agent's menu, an agent's expected payoffs (and expecting to pursue the virtual contracts there) are

$\tilde{U}_1^r = \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u} \right) (1-\beta_2)(\beta_1 - \beta_2) + X_2 + \beta_1 \delta Y_2$ and

$\tilde{U}_2^r = \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u} \right) (1-\beta_1)(\beta_2 - \beta_1) + Z_1 - \frac{\beta_2}{\beta_1} Z_1$.

If separation by menu-choice is intended, each agent must have an incentive to choose his intended contract, i.e., the no-deviation (ND) constraints $U_i^r \geq \tilde{U}_i^r$ must hold. Plugging in the respective values, we get

$$X_2 + \beta_1 \delta Y_2 \leq \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u} \right) (1-\beta_2)(\beta_2 - \beta_1). \qquad \text{(ND1)}$$

for agent 1 and

$$X_2 + \beta_2 \delta Y_2 + Z_1 \left(\frac{\beta_2 - \beta_1}{\beta_1} \right) \geq \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u} \right) (1 - \beta_1)(\beta_2 - \beta_1) \quad \text{(ND2)}$$

for agent 2.

Compared to the situation where the principal can observe each agent's β_i, in our simplified problem, she has to make additional real expected payments of

$$K = s_1 Z_1 + (1 - s_1)(X_2 + \delta Y_2).$$

The profit-maximizing set of menus of contracts that induces a separation by menu choice now minimizes these costs, subject to (ND1) and (ND2).

First of all, note that (ND2) must bind. Otherwise, any of the payments could be reduced, thereby also relaxing (ND1) and reducing the principal's real costs. Furthermore, note that when comparing X_2 and Y_2, the principal would ceteris paribus always prefer to use X_2, i.e., using X_2 is cheaper than using Y_2: A reduction of Y_2 by ε requires increasing X_2 by $\delta\beta_2\varepsilon$ in order to keep (ND2) satisfied. This adjustment would lead to a cost change of $\delta\beta_2\varepsilon - \delta\varepsilon < 0$.

The following lemma provides a lower bound of the total effective costs of using Z_1, X_2 and Y_2 as a function of s_1. We make use of the fact that (ND2) binds and that costs are linear in payments.

Lemma C.6. *The following use of Z_1, X_2 and Y_2 minimizes the cost of separation by menu choice if there were no costs due to Z_1 later than in $t = 0$:*

- $s_1 \leq \frac{(\beta_2 - \beta_1)}{\beta_2}$

  $X_2 = Y_2 = 0$ *and* $Z_1 = \beta_1 (1 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u} \right)$*; costs are*

$$K = s_1 \beta_1 (1 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u} \right)$$

- $\frac{(\beta_2 - \beta_1)}{\beta_2} < s_1 \leq 1 - \beta_1$

  $Y_2 = 0$, $X_2 = (1 - \beta_2)(\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u} \right)$ *and*

$Z_1 = \beta_1 (\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$; *costs are*

$$K = [s_1\beta_1 + (1-s_1)(1-\beta_2)] (\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$$

- $s_1 > 1 - \beta_1$

  $Z_1 = 0,\ X_2 = (1 - \beta_1 - \beta_2)(\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$ *and*

  $\delta Y_2 = (\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$;

  costs are

$$K = (1-s_1)(2 - \beta_1 - \beta_2)(\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right).$$

Proof. By Lagrange optimization, we minimize costs, subject to (ND) as well as non-negativity constraints. There, note that Z_1 and Y_2 cannot be negative: If Z_1 was negative, player 1's (rIR) constraint would not hold (and it will not be optimal to increase agent 1's virtual surplus, since this would further tighten the (ND2) constraint). If Y_2 were negative, player 2's (rIR) constraint would not hold in period $t = 1$. X_2 can be negative, but only if Y_2 is increased accordingly. Hence, the constraint $X_2 + \delta\beta_2 Y_2 \geq 0$ must hold as well.

This gives the Lagrange function

$L = -s_1 Z_1 - (1-s_1)(X_2 + \delta Y_2)$

$+\lambda_{ND1} \left[\frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_2-\beta_1) - X_2 - \beta_1\delta Y_2\right]$

$+\lambda_{ND2} \left[X_2 + \beta_2\delta Y_2 + Z_1 \left(\frac{\beta_2-\beta_1}{\beta_1}\right) - \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1)\right]$

$+\mu_X (X_2 + \delta\beta_2 Y_2) + \mu_Y Y_2 + \mu_Z Z_1$

and first-order conditions

$\frac{\partial L}{\partial Z_1} = -s_1 + \lambda_{ND2} \left(\frac{\beta_2-\beta_1}{\beta_1}\right) + \mu_Z = 0,$

$\frac{\partial L}{\partial X_2} = -(1-s_1) - \lambda_{ND1} + \lambda_{ND2} + \mu_X = 0,$

$\frac{\partial L}{\partial Y_2} = -(1-s_1)\delta - \beta_1\delta\lambda_{ND1} + \beta_2\delta\lambda_{ND2} + \delta\beta_2\mu_X + \mu_Y = 0.$

We know that $\lambda_{ND2} > 0$, furthermore rearranging and substituting gives the three conditions

$\frac{\partial L}{\partial Z_1} : \lambda_{ND2} = \frac{\beta_1(s_1-\mu_Z)}{\beta_2-\beta_1}$, (I)

$\frac{\partial L}{\partial X_2} : \lambda_{ND1} = \frac{-(\beta_2-\beta_1)+s_1\beta_2}{\beta_2-\beta_1} - \frac{\beta_1\mu_Z}{\beta_2-\beta_1} + \mu_X$, (II)

$\frac{\partial L}{\partial Y_2}$: $s_1 - (1-\beta_1) - \beta_1\mu_Z + \frac{\mu_Y}{\delta} + \mu_X(\beta_2-\beta_1) = 0$. (III)

Combining (II) and (III) implies that $\lambda_{ND1} = \frac{(1-s_1)(1-\beta_2)-\frac{\mu_Y}{\delta}}{\beta_2-\beta_1}$.

In the following, we just go through all potential cases and analyze whether they are feasible and if yes under which conditions.

1. $s_1 - (1-\beta_1) > 0$. Then, (III) implies that $\mu_Z > 0$, giving the following potential cases:

   (a) $\mu_Y > 0$: This is not feasible, since for $Y_2 = Z_1 = 0$, obtaining

   $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1)$ from binding (ND2) and plugging it into (ND1) gives $\beta_1 \geq \beta_2$, which is ruled out by assumption.

   (b) $\mu_Y = 0$: Then, (ND1) binds as well. Obtaining

   $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1) - \beta_2\delta Y_2$ from binding (ND2) and plugging it into (ND1) gives

   $Y_2 = \frac{\delta}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(\beta_2-\beta_1)$, implying that

   $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1-\beta_2)(\beta_2-\beta_1)$. Then,

   $X_2 + \delta\beta_2 Y_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_1)(\beta_2-\beta_1) > 0$.

2. $s_1 - (1-\beta_1) < 0$. Then, (III) implies that either $\mu_X > 0$ or $\mu_Y > 0$, or both, giving the following potential cases:

   (a) $\mu_X > 0$, $\mu_Y > 0$: Since $X_2 = Y_2 = 0$, $\mu_Z = 0$, and a binding (ND2) constraint gives $Z_1 = \beta_1(1-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. Furthermore, (II) implies that this is only feasible for $-(\beta_2-\beta_1) + s_1\beta_2 < 0$.

   (b) $\mu_X > 0$, $\mu_Y = 0$: Hence, (ND1) binds as well. Plugging $X_2 = -\delta\beta_2 Y_2$ into (ND2) gives $Z_1 = \beta_1(1-\beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. Plugging $X_2 = -\delta\beta_2 Y_2$ into (ND1) gives

   $\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_2-\beta_1) + \delta Y_2(\beta_2-\beta_1) = 0$, which is not feasible.

   (c) $\mu_X = 0$, $\mu_Y > 0$:

      i. $\mu_Z = 0$: (III) gives $\mu_Y = -\delta[s_1 - (1-\beta_1)]$, hence $\lambda_{ND1} = \frac{s_1\beta_2}{\beta_2-\beta_1} - 1 > 0$ if $-(\beta_2-\beta_1) + s_1\beta_2 > 0$. Then, binding (ND1) gives
      $X_2 = \frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1-\beta_2)(\beta_2-\beta_1)$, and plugging this into

(ND2) gives: $Z_1 = \beta_1 (\beta_2 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$. (II) implies that this is only feasible for $-(\beta_2 - \beta_1) + s_1\beta_2 \geq 0$.

ii. $\mu_Z > 0$: Only using X_2 is not feasible (see 1.(a)).

□

By the above lemma, we obtain restrictions on $\underline{s}$ and the threshold $1 - \beta_1$. Note that the menu offered for low s_1 is $\boldsymbol{C}_2$ and the first-best contract (intended to be chosen by agent 1).[1] It is easy to see that this is optimal even for the actual costs of Z_1[2] if s_1 is low enough: The costs of Z_1 occur for the fraction of agents of type 1, s_1, whereas X_2 and Y_2 have to be paid to the fraction of agents of type 2, $1 - s_1$. As Z_1 can be substituted by X_2 and Y_2 in a linear way (to separate the types), there exists an $\underline{s}$ such that $X_2 = Y_2 = 0$ is optimal for all $s_1 \leq \underline{s}$. Hence, when there are only few agents of type 1, the principal does not exploit them, while fully exploiting the agents of type 2.
To see that it is optimal to alter both $\boldsymbol{C}_1$ and $\boldsymbol{C}_2$ compared to the case without screening for intermediate values of s_1, note that the lower bound for the cost of using Z_1 is equal to the actual costs if Z_1 is low enough and does not require to alter any contracts but the real contract at the begining of the game and the virtual contract starting in the following period. Hence, the principal should choose a positive Z_1 for $s_1 < 1 - \beta_1$. At the same time, the principal should not only use Z_1 as the actual costs for using Z_1 would be strictly larger than the lower bound in this case.

Note that also (vIRA), (vIC), and (vC) are fulfilled, as they are not affected by the changes.[3] For (vrIRP) and (vIRP) observe that the compensation in the real contract in the first period and in the virtual contract's first period are not higher than the principal's surplus in these periods: $c(e^{FB}) + \overline{u} - \beta_2(1 - \beta_2)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$
$+ (1 - \beta_2)(\beta_2 - \beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) \leq e^{FB}\theta - \overline{\pi}$ and

[1] To see this, note that $Z_1 = \beta_1 (1 - \beta_1) \frac{\delta^2}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$ comes with $Z_{1v} = (1 - \beta_1) \frac{\delta}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$ which means that agent 1's perceived net utility from $t = 1$ on is reduced to zero. At the same time, making the agent accept a real contract that puts his utility below his outside option requires a benefit in the future virtual contract such that the perceived combined utility of accepting the real contract and choosing the virtual contract in the future is at least (and optimally exactly) zero. As this must hold even if the virtual contract is discounted with $\beta\delta$ from the time when the agent accepts the real contract, this means that taking the future real contract and the virtual contract after that must have a positive net utility from agent's perspective in $t = 0$. This is a contradiction to it being reduced to zero.

[2] In fact, the costs of increasing Z_1 are $s_1 \left(Z_1 + \sum_{i=1}^{\infty} \max\{0, \frac{Z_1}{\delta^2(1-\beta_1)} - \frac{\beta_1}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)(1 - \delta^i(1 - \beta_1))\}\right)$. To simplify the analysis, we use a cost function that does not take into account potential payments in later real contracts (due to a high Z_1) as a lower bound.

[3] (vC) is fulfilled by construction if Z_1 is large.

$c(e^{FB}) + \overline{u} - \beta_2 \frac{\delta}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) + (\beta_2 - \beta_1) \frac{\delta}{1-\delta} \left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right) \leq e^{FB}\theta - \overline{\pi}$. For the real contract in the second period the compensation does not exceed the principal's surplus when $(\beta_2 - \beta_1 - \beta_2(1-\beta_2)\delta - \beta_2\delta)\frac{\delta}{1-\delta} \leq 1$, which is true.

Separation by Action Now we calculate the costs for the principal when she offers only one menu of contracts which is supposed to be chosen by both agents and show that these costs are higher than making agents choose different menus of contracts.

Instead of offering a menu of contracts for each type of agent, the principal could just let agent 2 choose 1's contract, taking into account that 2 would then go for the virtual contract and make a career. In this case, it is without loss of generality to assume that only the profit-maximizing menu for agent 1, $\boldsymbol{C}_1$, is offered. Note that it cannot be optimal to induce separation by menu choice and then let agent 2 actually choose the virtual contract (unless $\beta_2 = 1$). Such a setting would give agent 2 a higher real rent than the one derived above because (ND) constraints would still have to hold.

If agent 2 is offered $\boldsymbol{C}_1$, his expected as well as real utility is

$$U_2^r = \tilde{U}_2^r = (1-\beta_1)(\beta_2 - \beta_1)\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right),$$

whereas agent 1 expects to get nothing.

Compared to the situation with symmetric information, though, the principal also foregoes the benefits from exploitation – because agent 2 not only select $\boldsymbol{C}_1$, but also goes for the virtual contract. Recall that under symmetric information, the net per-period profits the principal generates are $\pi^r - \overline{\pi} = \left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)\left(1 + \beta(1-\beta)\frac{\delta^2}{1-\delta}\right)$. Hence, under symmetric information, the principal's total profits when dealing with agent 2 would be $\frac{\left(e^{FB}\theta - c(e^{FB}) - \overline{u} - \overline{\pi}\right)}{1-\delta}\left(1 + \beta_2(1-\beta_2)\frac{\delta^2}{1-\delta}\right)$. If letting agent 2 choose $\boldsymbol{C}_1$, the principal's profits dealing with agent 2 are $(1+\delta)\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$
$+ (2-\beta_1)\beta_1\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right)$.

Therefore, the principal's costs when letting agent 2 choose $\boldsymbol{C}_1$, taking into account that he then actually goes for the virtual contract, compared to the case of symmetric information (which also served as our benchmark above), are

$$\tilde{K} = (1-s_1)\left[1 + \frac{\beta_2(1-\beta_2)}{1-\delta} - (2-\beta_1)\beta_1\right]\frac{\delta^2}{1-\delta}\left(e^{FB}\theta - c(e^{FB}) - \overline{\pi} - \overline{u}\right).$$

Comparing these costs to the costs of separation by menu choice for large s_1 (which serve as an upper bound), the condition equivalent to separation by action being cheaper,

$$\left[1 + \frac{\beta_2(1-\beta_2)}{1-\delta} - (2-\beta_1)\,\beta_1\right] \leq (2-\beta_1-\beta_2)(\beta_2-\beta_1),$$

shows that separation by action can only be optimal if $\beta_2 = 1$.

For $\beta_2 = 1$, the cost expressions for separation by actions and for separation by menu choice are the same for $s_1 > 1-\beta_1$, so separation by action is (weakly) optimal. $\square$

Appendix D

Appendix for Chapter 4.

D.1 Proofs

Proof of Lemma 4.1.

This proof is slightly more general. Let the probability that $r = 1$ be p and the probability that $r = 0$ be $1 - p$. (i) When an agent observed i high and j low conditionally independent signals, her belief that $r = 1$ is $\frac{p\binom{i+j}{i}\alpha^i(1-\alpha)^j}{p\binom{i+j}{i}\alpha^i(1-\alpha)^j+(1-p)\binom{i+j}{i}(1-\alpha)^i\alpha^j} = 1 - \frac{(1-p)\binom{i+j}{j}\alpha^j(1-\alpha)^i}{(1-p)\binom{i+j}{j}\alpha^j(1-\alpha)^i+p\binom{i+j}{j}(1-\alpha)^j\alpha^i} = \frac{p\alpha^{i-j}}{p\alpha^{i-j}+(1-p)(1-\alpha)^{i-j}}$. Because this expression only depends on the difference between high and low conditionally independent signals for a given α, this difference is a sufficient statistic for the agent's belief.

(ii) It follows from equation 4.3 that an agent picks $d_{\theta k} = 1$ if and only if $Prob(r = 1|S^1_{\theta k})\overline{u}_{1\theta} + (1 - Prob(r = 1|S^1_{\theta k}))\underline{u}_{1\theta} > Prob(r = 1|S^1_{\theta k})\underline{u}_{0\theta} + (1 - Prob(r = 1|S^1_{\theta k}))\overline{u}_{0\theta} \Leftrightarrow Prob(r = 1|S^1_{\theta k}) > \frac{\overline{u}_{0\theta}-\underline{u}_{1\theta}}{\overline{u}_{1\theta}-\underline{u}_{1\theta}+\overline{u}_{0\theta}-\underline{u}_{0\theta}}$. This is equivalent to $z^1_{\theta k} > \frac{\log\left(\frac{p}{1-p}\frac{\overline{u}_{1\theta}-\underline{u}_{0\theta}}{\overline{u}_{0\theta}-\underline{u}_{1\theta}}\right)}{\log\left(\frac{1}{\alpha}-1\right)}$. □

Proof of Lemma 4.2.

First, observe that

$$\begin{aligned}
\max_{d_{\theta k}(S^1_{\theta k})} d_{\theta k} \quad & [Prob(r = 1|S^1_{\theta k})\overline{u}_{1\theta} + (1 - Prob(r = 1|S^1_{\theta k}))\underline{u}_{1\theta}] \\
-(d_{\theta k} - 1) \quad & [Prob(r = 1|S^1_{\theta k})\underline{u}_{0\theta} + (1 - Prob(r = 1|S^1_{\theta k}))\overline{u}_{0\theta}] \\
= \frac{1}{2}[\sum_{i=\frac{|S^t_{\theta k}|+z^*_\theta}{2}+1}^{|S^t_{\theta k}|} \quad & \binom{|S^t_{\theta k}|}{i}(\alpha^i(1-\alpha)^{|S^t_{\theta k}|-i}\overline{u}_{1\theta} + \alpha^{|S^t_{\theta k}|-i}(1-\alpha)^i\underline{u}_{1\theta}) \\
+ \sum_{i=0}^{\frac{|S^t_{\theta k}|+z^*_\theta}{2}} \quad & \binom{|S^t_{\theta k}|}{i}(\alpha^i(1-\alpha)^{|S^t_{\theta k}|-i}\underline{u}_{0\theta} + \alpha^{|S^t_{\theta k}|-i}(1-\alpha)^i\overline{u}_{0\theta})],
\end{aligned}$$

where the first sum adds up the probability-weighted terms for $d_{\theta k} = 1$ and the second sum adds up the probability-weighted terms for $d_{\theta k} = 0$. For example, if the agent has observed no high signals and the threshold z^*_θ is not negative, she will choose $d_{\theta k} = 0$. If $r = 1$ (which occurs with probability 0.5), the probability of zero high signals is $(1-\alpha)^{|\underline{S}^t_{\theta k}|}$, which is multiplied by $\underline{u}_{0\theta}$. If $r = 0$ (which also occurs with probability 0.5), the probability of zero high signals is $\alpha^{|\underline{S}^t_{\theta k}|}$, which is multiplied by $\overline{u}_{0\theta}$.

Second, assume for now $z^*_\theta = \frac{\log\left(\frac{\overline{u}_{1\theta}-\underline{u}_{0\theta}}{\overline{u}_{0\theta}-\underline{u}_{1\theta}}\right)}{\log\left(\frac{1}{\alpha}-1\right)} \notin \mathbb{Z}$. We look at the expected utility from the decision, if we increase $|S^1_{\theta k}$ by one: There are two cases: either (i) $\lfloor\frac{|\underline{S}^t_{\theta k}|+1+z^*_\theta}{2}\rfloor = \lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor$ or (ii) $\lfloor\frac{|\underline{S}^t_{\theta k}|+1-z^*_\theta}{2}\rfloor = \lfloor\frac{|\underline{S}^t_{\theta k}|-z^*_\theta}{2}\rfloor + 1$.

Case (i): Because $\binom{n+1}{k} = \binom{n}{k-1} + \binom{n}{k}$ and with an index shift for one term,

$$
\begin{aligned}
& 2[\mathbb{E}[U_{\theta k}||\underline{S}^t_{\theta k}|+1] - \mathbb{E}[U_{\theta k}||\underline{S}^t_{\theta k}|]] \\
= & \sum_{i=\lfloor\frac{|\underline{S}^t_{\theta k}|+1+z^*_\theta}{2}\rfloor+1}^{|\underline{S}^t_{\theta k}|+1} \binom{|\underline{S}^t_{\theta k}|+1}{i} \alpha^i (1-\alpha)^{|\underline{S}^t_{\theta k}|+1-i} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
+ & \sum_{i=0}^{\lfloor\frac{|\underline{S}^t_{\theta k}|+1+z^*_\theta}{2}\rfloor} \binom{|\underline{S}^t_{\theta k}|+1}{i} \alpha^{|\underline{S}^t_{\theta k}|+1-i} (1-\alpha)^i (\overline{u}_{0\theta} - \underline{u}_{1\theta}) \\
- & \sum_{i=\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor+1}^{|\underline{S}^t_{\theta k}|} \binom{|\underline{S}^t_{\theta k}|}{i} \alpha^i (1-\alpha)^{|\underline{S}^t_{\theta k}|-i} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
- & \sum_{i=0}^{\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor} \binom{|\underline{S}^t_{\theta k}|}{i} \alpha^{|\underline{S}^t_{\theta k}|-i} (1-\alpha)^i (\overline{u}_{0\theta} - \underline{u}_{1\theta}) \\
= & \binom{|\underline{S}^t_{\theta k}|}{\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor} (\alpha^{\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor+1} (1-\alpha)^{|\underline{S}^t_{\theta k}|-\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
- & (\alpha^{|\underline{S}^t_{\theta k}|-\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor} (1-\alpha)^{\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor+1} (\overline{u}_{0\theta} - \underline{u}_{1\theta}))
\end{aligned}
$$

This is positive by the definition of $\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor$ in case (i). More generally, the difference in expected utilities from the decision going from $|\underline{S}^t_{\theta k}|+2j$ to $|\underline{S}^t_{\theta k}|+2j+1$ signals, where $j \in \mathbb{N}$ is

$$
\begin{aligned}
\binom{|\underline{S}^t_{\theta k}|+2j}{\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor + j} \alpha^j (1-\alpha)^j (\alpha^{\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor+1} (1-\alpha)^{|\underline{S}^t_{\theta k}|-\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
-(\alpha^{|\underline{S}^t_{\theta k}|-\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor} (1-\alpha)^{\lfloor\frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2}\rfloor+1} (\overline{u}_{0\theta} - \underline{u}_{1\theta})),
\end{aligned}
$$

which goes to zero for $j \to \infty$, because $\alpha(1-a\alpha) < \frac{1}{4}$.

Case (ii): Similarly,

$$
\begin{aligned}
& 2[\mathbb{E}[U_{\theta k}||\underline{S}^t_{\theta k}|+1] - \mathbb{E}[U_{\theta k}||\underline{S}^t_{\theta k}|]] \\
= & \sum_{i=\lfloor \frac{|\underline{S}^t_{\theta k}|+1+z^*_\theta}{2} \rfloor +2}^{|\underline{S}^t_{\theta k}|+1} \binom{|\underline{S}^t_{\theta k}|+1}{i} \alpha^i (1-\alpha)^{|\underline{S}^t_{\theta k}|+1-i} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
+ & \sum_{i=0}^{\lfloor \frac{|\underline{S}^t_{\theta k}|+1+z^*_\theta}{2} \rfloor +1} \binom{|\underline{S}^t_{\theta k}|+1}{i} \alpha^{|\underline{S}^t_{\theta k}|+1-i} (1-\alpha)^i (\overline{u}_{0\theta} - \underline{u}_{1\theta}) \\
- & \sum_{i=\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1}^{|\underline{S}^t_{\theta k}|} \binom{|\underline{S}^t_{\theta k}|}{i} \alpha^i (1-\alpha)^{|\underline{S}^t_{\theta k}|-i} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
- & \sum_{i=0}^{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor} \binom{|\underline{S}^t_{\theta k}|}{i} \alpha^{|\underline{S}^t_{\theta k}|-i} (1-\alpha)^i (\overline{u}_{0\theta} - \underline{u}_{1\theta}) \\
= & -\binom{|\underline{S}^t_{\theta k}|}{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1} (\alpha^{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1} (1-\alpha)^{|\underline{S}^t_{\theta k}| - \lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
- & (\alpha^{|\underline{S}^t_{\theta k}| - \lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor} (1-\alpha)^{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1} (\overline{u}_{0\theta} - \underline{u}_{1\theta}))
\end{aligned}
$$

This is negative by the definition of $\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor$ in case (ii). More generally, the difference in expected utilities from the decision going from $|\underline{S}^t_{\theta k}|+2j$ to $|\underline{S}^t_{\theta k}|+2j+1$ signals, where $j \in \mathbb{N}$, without taking into account link formation cost is

$$
\begin{aligned}
\binom{|\underline{S}^t_{\theta k}|+2j}{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1+j} \alpha^j (1-\alpha)^j (\alpha^{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1} (1-\alpha)^{|\underline{S}^t_{\theta k}| - \lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
-(\alpha^{|\underline{S}^t_{\theta k}| - \lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor} (1-\alpha)^{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1} (\overline{u}_{0\theta} - \underline{u}_{1\theta})),
\end{aligned}
$$

which again goes to zero for $j \to \infty$.

Third, consider $z^*_\theta \in \mathbb{Z}$: If the agent is indifferent, choosing either alternative with probability $\frac{1}{2}$ has the same expected utility as always choosing 0. With this in mind the expressions are the same as for $z^*_\theta \notin \mathbb{Z}$. However, $\alpha^{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1} (1-\alpha)^{|\underline{S}^t_{\theta k}| - \lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) - (\alpha^{|\underline{S}^t_{\theta k}| - \lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor} (1-\alpha)^{\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor +1} (\overline{u}_{0\theta} - \underline{u}_{1\theta})$ is zero in case (ii) as $2\lfloor \frac{|\underline{S}^t_{\theta k}|+z^*_\theta}{2} \rfloor + 1 - |\underline{S}^t_{\theta k}| = z^*_\theta$, so not every additional signal is beneficial, but only every two additional conditionally independent signals are beneficial. □

Proof of Lemma 4.3.

Assume this is not true, but an agent has optimally formed a link to an agent with type

$\eta \neq \theta$ via the Internet. Then an agent would rather connect to an agent of type η in her own region and an agent of her own type at the different region at the same total cost, but obtaining more independent signals. If she has already linked to the agent of type η in her own region, too, then linking to an agent of type θ via the Internet gives her (at least) the same number of independent signals as the link to an agent of type η at a strictly lower cost. This is also possible if n is sufficiently large. This is a contradiction to the optimality of the link to an agent with type $\eta \neq \theta$ via the Internet. □

Proof of Lemma 4.4.

I will prove a slightly more general version of Lemma 4.4, also dealing with $z^*_\theta \in \mathbb{Z}$.

Lemma D.1. *(i) Suppose $\overline{s}_T \geq \overline{s}_R$. Agents form (weakly) more links when the inter-regional linking cost, c_R, decreases.*
(ii) Conditional on forming a positive number of links, if $z^*_\theta \notin \mathbb{Z}$, agents form (weakly) fewer domestic links when the inter-regional linking cost, c_R, decreases. If $z^*_\theta \in \mathbb{Z}$, agents form at most one additional domestic link when c_R decreases.

(i) When the cost of inter-regional linking decreases, an agent might form an additional inter-regional link, make no change, or substitute one or several of her domestic links with one or several inter-regional links. We have to show that in the last case, a number of domestic links will be replaced with at least as many inter-regional links. We know from Lemma 4.3 that we have to compare domestic links to agents of different types with inter-regional links to agents of the same type. The former provide access to $\overline{s}_U + \overline{s}_T$ conditionally independent signals, while the latter provide access to $\overline{s}_U + \overline{s}_R$ conditionally independent signals. Hence, if $\overline{s}_T \geq \overline{s}_R$, inter-regional links provide access to weakly fewer signals, and therefore an agent will substitute a domestic with at least as many inter-regional links.
(ii) If an agent forms links at a certain cost c_R, lowering c_R does not change the cost of domestic links. The benefit of forming an additional domestic link is unchanged if $z^*_\theta \notin \mathbb{Z}$, so it is not optimal to form additional domestic links. If $z^*_\theta \in \mathbb{Z}$, it might make sense to form an additional domestic link together with an inter-regional link when an odd number of additional signals provides a higher benefit of creating a new link. Even then, it will never make sense to form two or more additional domestic links, because then it would have made sense to form these links at a higher c_R, too. □

Proof of Proposition 4.1.

(i) Within their region, agents link to the most similar types, i.e., an agent who forms $2i$ domestic links forms these links with i agents of a lower type and i agents of a higher type, if possible (in the case of an odd number, she would be indifferent about one link). If this is not possible, because e.g., there are only $j < i$ agents of a lower type available, the agent would have to pay more for those $i - j$ domestic links as those agents would have a more distant higher type. Therefore, an agent further away from the central type $\frac{m+1}{2}$ forms weakly fewer domestic links and replaces them with inter-regional links at a weakly higher c_R.

For forming inter-regional links, all agents face the same cost. The benefits of agents further away from the central type are weakly higher if $z^*_\theta \notin \mathbb{Z}$. If $z^*_\theta \in \mathbb{Z}$, it can happen that a more central type forms one domestic link more than a less central type and one inter-regional link, while the less central type does not form any inter-regional link (see proof of of Lemma 4.2).

(ii) Let $c_T > 0$ and $\overline{u}_1, \overline{u}_0, \underline{u}_1, \underline{u}_0$ be such that an agent of type $\eta^* + 1$ it is just optimal to form only $2\eta^*$ domestic links to receive exactly $z^*_{\eta^*+1} = z^*_\theta$ conditionally independent signals and let $c_R > c^*_R = \eta^* c_T$. For all other agents with types $\eta \in \{\eta^* + 1, ..., m - \eta^*\}$, it is optimal to form at least the same number of links, too. For agents with types $\eta \notin \{\eta^* + 1, ..., m - \eta^*\}$, it is optimal to not form any links, because by assumption it is not optimal for them to form the same number of links (because this would be more expensive), and forming a positive but lower number of links is not optimal, because the signals could never change the decision. If $c_R < c^*_R$, types η^* and $m + 1 - \eta^*$ would form links, because they could form $2\eta^* - 1$ domestic links and one inter-regional link at costs below the total linking costs of agents of type $\eta^* + 1$ when c_R was greater than c^*_R (and forming links was optimal for them by assumption). □

Proof of Proposition 4.2.

If h increases, the benefit of additional signals increases, while the costs for additional signals stay the same. Therefore, the optimal number of links weakly increases with h. Suppose an agent with $h = 1$ does not form any inter-regional links. If $h \to \infty$, the benefit of an additional inter-regional link approaches infinity. Therefore there exists an $\overline{h} > 1$ such that an agent with $h = \overline{h}$ forms inter-regional links.

□

Proof of Proposition 4.3.

(i) We know from Lemma 4.2 that conditionally independent signals improve an agent's decision. Therefore, if c_R is sufficiently small, all agents will form at least two links. Call the inter-regional linking cost which ensures this c'_R. The number of conditionally independent signals an agent receives from forming a domestic link is $\bar{s}_U + \bar{s}_T$ and the number of conditionally independent signals an agent receives from forming an inter-regional link (with an agent of the same type) is $\bar{s}_U + \bar{s}_R$. If $c_R < \min\{c_T, c_T \frac{\bar{s}_U+\bar{s}_R}{\bar{s}_U+\bar{s}_T}\} =: c''_R$, agents will only form inter-regional links, because they are cheaper and cheaper per signal than domestic links. Set $\underline{c_R} = \min\{c'_R, c''_R\}$. For $c_R < \underline{c_R}$, all agents form no domestic links and at least links to the agents of the same type in the two closest regions. Hence, there is a path from every agents to all agents of the same type, but no path to an agent of a different type. Consequently, even if $\bar{t} \to \infty$, agents will never of observe the signals of other types. Only by chance have the signals the same realizations and consensus occurs.

If the inter-regional link cost is high enough, agents will not form any inter-regional links. Such a high enough cost exists, because the value-added of signals is decreasing in the number of conditionally independent signals as we know from Lemma 4.2. Call the threshold $\overline{c_R}$. If agents do not form any inter-regional links, even if $\bar{t} \to \infty$, they will never of observe the signals of agents in other regions. Only by chance have the signals the same realizations and consensus occurs.

(ii) It is sufficient to show that there exist according parameters such that agents form at least two inter-regional links and at least two domestic links: Then each agent forms links to agents of the same type in the at least two closest regions and to at least two agents in the same region with the closest types. This means that there exists a path from every agent to every other agent. Consequently, for $\bar{t} \to \infty$, agents will observe all signals and then hold the same beliefs.

Let c_R be very high and c_T be such that each agent optimally forms domestic links to observe at least $6\bar{s}_U + 3\bar{s}_R + 4\bar{s}_T$ conditionally independent signals in addition to her own signals. This is possible because of Lemma 4.2. Then reduce c_R so that agents will form inter-regional links and replace domestic links (see Lemma 4.4). Do so till the point where the first agent is indifferent between giving up her second to last domestic link (this will be agents of type 1 and m). For a neighborhood of slightly higher inter-regional linking costs, all agents have at least two domestic links. Because all agents previously observed $6\bar{s}_U + 3\bar{s}_R + 4\bar{s}_T$ conditionally independent signals in addition to their own signals, all agents also have at least two inter-regional links in this neighborhood of inter-regional linking costs. $\square$

Proof of Lemma 4.5.

That the component is strongly connected implies that all agents have domestic links: Suppose one agent did not have a domestic link. Then inter-regional links must be cheaper than domestic links or the agent could not observe any other agent's signals. If inter-regional links are cheaper than domestic links, no agent would form domestic links. This is a contradiction. $\overline{s}_R = \overline{s}_T$ implies that agents replace each domestic link one by one with a domestic link. Because the gross utilities are the same for all agents, an agent has at most one inter-regional link more and at most one domestic link less then a slightly more central agent. Therefore, the maximal path length is determined by the path from agent 11 to agent mn via agent $1n$. Suppose c_R is so high that when increasing it the component would no longer be strongly connected. Let $i_1(c_R) > 0$ the number of inter-regional links agents of type 1 form and $i_1(c_R) > 0$ be the number of domestic links they form. Then the path length from 11 to $1n$ is equal to $\lceil \frac{2(n-1-i_1)}{i_1} \rceil$, because 11's link will go to agent $1(i_1+1)$ and this agent's link will go to agent $1(i_1 + 1 + 0.5i_1)$ (because she forms links in both directions) and so on. The path length from agent $1n$ to agent mn is between $\lceil \frac{m-1-i_2}{i_2} \rceil$ and $\lceil \frac{2(m-1-i_2)}{i_2-1} \rceil$, depending on how many inter-regional links central agents form. If they form inter-regional links, then the path length is $\lceil \frac{m-1-i_2}{i_2} \rceil$, because then they will form twice as many domestic links as agents of type 1. For now, assume all agents form inter-regional links. Because domestic links are replaced one by one with inter-regional links (as long as the component is strongly connected) when c_R decreases, the number of links $b := i_1 + i_2$ stays constant. The maximal path length is $\lceil \frac{2(n-1-i_1)}{i_1} \rceil + \lceil \frac{m-1-i_1-b}{b-i_1} \rceil$. This function has one minimum on $i_1 \in \{1, b-1\}$ and is decreasing before the minimum and increasing after it. If some relatively central agents do not form any inter-regional links for high c_R, then the maximal path length decreases even faster when decreasing c_R, because they do not change their links before they start forming inter-regional links. Because types start forming inter-regional links one type after another, this will increase the path length from agent $1n$ to agent mn by not more than 2 at a time, and so there is no inflection point for the maximal path length. Consequently, the maximal path length follows a U-shape with respect to the inter-regional linking cost. $\square$

Proof of Proposition 4.4.

This immediately follows from Proposition 4.3 and Lemma 4.5 and the fact the maximal path length is ∞ if the society is not strongly connected. $\square$

Proof of Lemma 4.6.

We will show that the expected utility from the decision decreases when both $|S_0|$ and $|S_1|$ increase by the same amount. This together with Lemma 4.2 implies that the expected utility from the decision decreases in $|S_1|$.

Similar to the proof of Lemma 4.2, assume that $z^*_\theta \notin \mathbb{Z}$. We separate the terms into cases in which the signals in S_1 are correct (which happens with probability α) and those in which they are incorrect.

$$
\begin{aligned}
& 2[\mathbb{E}[U_{\theta k}||S_0|+1,|S_1|+1] - \mathbb{E}[U_{\theta k}||S_0|,|S_1|]] \\
= \; & (\overline{u}_{1\theta} - \underline{u}_{0\theta})(1-\alpha) \\
& \left(\sum_{i=\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor + 2}^{|S_0|+1} \binom{|S_0|+1}{i} \alpha^i (1-\alpha)^{|S_0|+1-i} - \sum_{i=\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor + 1}^{|S_0|} \binom{|S_0|}{i} \alpha^i (1-\alpha)^{|S_0|-i} \right) \\
+ \; & (\overline{u}_{0\theta} - \underline{u}_{1\theta})\alpha \\
& \left(\sum_{i=0}^{\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor + 1} \binom{|S_0|+|S_1|+1}{i} \alpha^{|S_0|+1-i} (1-\alpha)^i - \sum_{i=0}^{\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor} \binom{|S_0|}{i} \alpha^{|S_0|-i} (1-\alpha)^i \right) \\
= \; & \binom{|S_0|}{\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor + 1} (\alpha^{\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor + 1} (1-\alpha)^{|S_0| - \lfloor \frac{|S_0|+z^*_\theta}{2} \rfloor + 1} (\overline{u}_{1\theta} - \underline{u}_{0\theta}) \\
- \; & (\alpha^{|S_0| - \lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor + 1} (1-\alpha)^{\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor + 1} (\overline{u}_{0\theta} - \underline{u}_{1\theta})),
\end{aligned}
$$

which is smaller than zero, because $\frac{\overline{u}_{1\theta} - \underline{u}_{0\theta}}{\overline{u}_{0\theta} - \underline{u}_{1\theta}} > \left(\frac{1-\alpha}{\alpha}\right)^{\lfloor \frac{|S_0|+|S_1|+z^*_\theta}{2} \rfloor - |S_0|}$.

Dealing with $z^*_\theta \in \mathbb{Z}$ and showing convergence to zero for $|S_0|$ is analogous to that part in the proof of Lemma 4.2. □

Proof of Lemma 4.7.

(i) There are two cases to deal with: (a) $|S_0| + |S_1| + |S_2| + z^*_\theta \in \bigcup_{j \in \mathbb{Z}} [2j, 2j+1)$ and (b) $|S_0| + |S_1| + |S_2| + z^*_\theta \notin \bigcup_{j \in \mathbb{Z}} [2j, 2j+1)$.

(a) In this case, $\lfloor\frac{|S_0|-|S_1|-|S_2|+z_\theta^*}{2}\rfloor > \lfloor\frac{|S_0|-|S_1|-|S_2|+z_\theta^*-1}{2}\rfloor$ and $\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor > \lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*-1}{2}\rfloor$. Proceeding as in the proof of Lemma 4.6 yields

$$\begin{aligned}
&\binom{|S_0|}{\lfloor\frac{|S_0|-|S_1|-|S_2|+z_\theta^*}{2}\rfloor}(\alpha^{\lfloor\frac{|S_0|-|S_1|-|S_2|+z_\theta^*}{2}\rfloor+2}(1-\alpha)^{|S_0|-\lfloor\frac{|S_0|-|S_1|-|S_2|+z_\theta^*}{2}\rfloor}(\overline{u}_{1\theta}-\underline{u}_{0\theta})\\
&\qquad-(\alpha^{|S_0|-\lfloor\frac{|S_0|-|S_1|-|S_2|+z_\theta^*}{2}\rfloor}(1-\alpha)^{\lfloor\frac{|S_0|-|S_1|-|S_2|+z_\theta^*}{2}\rfloor+2}(\overline{u}_{0\theta}-\underline{u}_{1\theta}))\\
&+\binom{|S_0|}{\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor}(\alpha^{\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(1-\alpha)^{|S_0|-\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(\overline{u}_{1\theta}-\underline{u}_{0\theta})\\
&\qquad-(\alpha^{|S_0|-\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(1-\alpha)^{\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(\overline{u}_{0\theta}-\underline{u}_{1\theta})).
\end{aligned}$$

The first term is always smaller than zero for $|S_1|, |S_2| > 1$. The second term is smaller than zero if $|S_1| > |S_2|$, but greater than zero if $|S_1| + 1 < |S_2|$. If $|S_0| - |S_1| + z_\theta^*$ is small, the second term dominates.

(b) In this case, $\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor < \lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*+1}{2}\rfloor$ and $\lfloor\frac{|S_0|+|S_1|+|S_2|+z_\theta^*}{2}\rfloor < \lfloor\frac{|S_0|+|S_1|+|S_2|+z_\theta^*+1}{2}\rfloor$. Proceeding as in the proof of Lemma 4.6 yields

$$\begin{aligned}
&-\binom{|S_0|}{\lfloor\frac{|S_0|+|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(\alpha^{\lfloor\frac{|S_0|+|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(1-\alpha)^{|S_0|-\lfloor\frac{|S_0|+|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(\overline{u}_{1\theta}-\underline{u}_{0\theta})\\
&\qquad-(\alpha^{|S_0|-\lfloor\frac{|S_0|+|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(1-\alpha)^{\lfloor\frac{|S_0|+|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(\overline{u}_{0\theta}-\underline{u}_{1\theta}))\\
&-\binom{|S_0|}{\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor+1}(\alpha^{\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor+2}(1-\alpha)^{|S_0|-\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor}(\overline{u}_{1\theta}-\underline{u}_{0\theta})\\
&\qquad-(\alpha^{|S_0|-\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor}(1-\alpha)^{\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor+2}(\overline{u}_{0\theta}-\underline{u}_{1\theta})).
\end{aligned}$$

The first term is always smaller than zero for $|S_1|, |S_2| > 1$. The second term is smaller than zero if $|S_1| > |S_2|$, but greater than zero if $|S_1|+1 < |S_2|$. If $|S_0|-|S_1|+z_\theta^*$ is big, the second term dominates.

(ii) Proceeding as in the proof of Lemma 4.6 yields

$$\begin{aligned}
&-\binom{|S_0|}{\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor+1}(\alpha^{\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor+2}(1-\alpha)^{|S_0|-\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor}(\overline{u}_{1\theta}-\underline{u}_{0\theta})\\
&\qquad-(\alpha^{|S_0|-\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor}(1-\alpha)^{\lfloor\frac{|S_0|-|S_1|+|S_2|+z_\theta^*}{2}\rfloor+2}(\overline{u}_{0\theta}-\underline{u}_{1\theta}))\\
&+\binom{|S_0|}{\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor}(\alpha^{\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor+1}(1-\alpha)^{|S_0|-\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor+1}(\overline{u}_{1\theta}-\underline{u}_{0\theta})\\
&\qquad-(\alpha^{|S_0|-\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor+1}(1-\alpha)^{\lfloor\frac{|S_0|+|S_1|-|S_2|+z_\theta^*}{2}\rfloor+1}(\overline{u}_{0\theta}-\underline{u}_{1\theta})).
\end{aligned}$$

Both terms are smaller than zero for $|S_2| > |S_1|$ and greater than zero for $|S_2| + 1 < |S_1|$. If $|S_2| + 1 = |S_1|$ the sum of both terms is zero. $\square$

Proof of Corollary 4.1.

This follows directly from Lemma 4.6. $\square$

Proof of Proposition 4.5.

(i) This follows directly from Lemma 4.7.

(ii) This follows directly from Lemma 4.6. $\square$

Proof of Corollary 4.2.

This follows from Remark 4.1 and Proposition 4.5 together with a situation in which a rational agent forms at least one inter-regional link and a slight decrease in the inter-regional linking cost leads to a change in a naive agent's link formation. While the naive agent can be strictly (and bounded away from zero) worse off according to Proposition 4.5, a rational agent is slightly better off.

$\square$

Proof of Proposition 4.6.

Both parts follow directly from Lemma 4.6.

$\square$

Proof of Lemma 4.8.

If the type-specific linking cost increases, the number of domestic links will obviously not increase. An agent might replace a domestic link with an inter-regional link. In this case, the number of inter-regional links increases. It will not decrease because the costs are constant and the benefits weakly increasing for $z_\theta^* \notin \mathbb{Z}$. $\square$

Bibliography

Abramowitz, Alan I. and Kyle L. Saunders, "Is Polarization a Myth?," *Journal of Politics*, 2008, *70* (02), 542–555.

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar, "Dynamics of Information Exchange in Endogenous Social Networks," *Theoretical Economics*, 2014, *9* (1), 41–97.

___ **, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar**, "Bayesian Learning in Social Networks," *Review of Economic Studies*, 2011, *78* (4), 1201–1236.

___ **, Victor Chernozhukov, and Muhamet Yildiz**, "Fragility of Asymptotic Agreement under Bayesian Learning," *Theoretical Economics*, 2016, *11* (1), 187–225.

Allen, Franklin and Douglas Gale, "Financial Contagion," *Journal of Political Economy*, 2000, *108* (1), 1–33.

Andreoni, James and Tymofiy Mylovanov, "Diverging Opinions," *American Economic Journal: Microeconomics*, 2012, *4* (1), 209–232.

Ariely, Dan and Klaus Wertenbroch, "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science*, 2002, *13* (3), 219–224.

Armstrong, Mark and John Vickers, "Consumer Protection and Contingent Charges," *Journal of Economic Literature*, 2012, *50* (2), 477–493.

___ **and Yongmin Chen**, "Inattentive Consumers and Product Quality," *Journal of the European Economic Association: Papers and Proceedings*, 2009, *7* (2-3), 411–422.

___ **, John Vickers, and Jidong Zhou**, "Consumer Protection and the Incentive to Become Informed," *Journal of the European Economic Association: Papers and Proceedings*, 2009, *7* (2-3), 399–410.

Asker, John and Estelle Cantillon, "Properties of Scoring Auctions," *The RAND Journal of Economics*, 2008, *39* (1), 69–85.

___ **and** ___ , "Procurement When Price and Quality Matter," *The RAND Journal of Economics*, 2010, *41* (1), 1–34.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger, "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.

Aumann, Robert J., "Agreeing to Disagree," *Annals of Statistics*, 1976, *4* (6), 1236–1239.

Ausubel, Lawrence M., "The Failure of Competition in the Credit Card Market," *American Economic Review*, 1991, *81* (1), 50–81.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein, "Behavioral Hazard in Health Insurance," *Quarterly Journal of Economics*, 2015, *130* (4), 1623–1667.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel, "Social Networks and Housing Markets," Technical Report 2016.

Bajari, Patrick and Steven Tadelis, "Incentives versus Transaction Costs: A Theory of Procurement Contracts," *The RAND Journal of Economics*, 2001, *32* (3), 87–407.

____, Stephanie Houghton, and Steven Tadelis, "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," *American Economic Review*, 2014, *104* (4), 1288–1319.

Bala, Venkatesh and Sanjeev Goyal, "Learning from Neighbours," *Review of Economic Studies*, 1998, *65* (3), 595–621.

____ and ____, "A Noncooperative Model of Network Formation," *Econometrica*, 2000, *68* (5), 1181–1229.

Beaudry, Paul and Michel Poitevin, "Contract Renegotiation: A Simple Framework and Implications for Organization Theory," *The Canadian Journal of Economics / Revue canadienne d'Economique*, 1995, *28* (2), pp. 302–335.

Bénabou, Roland, "Groupthink: Collective Delusions in Organizations and Markets," *Review of Economic Studies*, 2013, *80* (2), 429–462.

Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky, "The Nash Bargaining Solution in Economic Modelling," *The RAND Journal of Economics*, 1986, *17* (2), 176–188.

Birulin, Oleksii, "Optimal Auctions of Procurement Contracts," 2014.

____ and Sergei Izmalkov, "On Cost Overruns in Procurement," 2013.

Bohren, J Aislinn, "Informational Herding with Model Misspecification," *Journal of Economic Theory*, 2016, *163*, 222–247.

Brunnermeier, Markus K., Alp Simsek, and Wei Xiong, "A Welfare Criterion For Models With Distorted Beliefs," *Quarterly Journal of Economics*, 2014, *129* (4), 1753–1797.

Burguet, Roberto, Juan-José Ganuza, and Esther Hauk, "Limited Liability and Mechanism Design in Procurement," *Games and Economic Behavior*, 2012, *76* (1), 15 – 25.

Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin, "Regulation for Conservatives: Behavioral Economics and the Case for "Asymmetric Paternalism"," *University of Pennsylvania Law Review*, 2003, pp. 1211–1254.

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics*, 2009, *124* (4), 1639–1674.

Chang, Wei-Shiun, Timothy C. Salmon, and Krista J. Saral, "Procurement Auctions with Renegotiation and Wealth Constraints," *Economic Inquiry*, 2016, *54* (3), 1684–1704.

Che, Yeon-Koo, "Design Competition through Multidimensional Auctions," *RAND Journal of Economics*, 1993, *24* (4), 668–680.

Chen-Ritzo, Ching-Hua, Terry P. Harrison, Anthony M. Kwasnica, and Douglas J. Thomas, "Better, Faster, Cheaper: An Experimental Analysis of a Multiattribute Reverse Auction Mechanism with Restricted Information Feedback," *Management Science*, 2005, *51* (12), 1753–1762.

Chetty, Raj, "Behavioral Economics and Public Policy: A Pragmatic Perspective," *American Economic Review: Papers & Proceedings*, 2015, *105* (5), 1–33.

——— **, John N. Friedman, Søren Leth-Petersen, Torben H. Nielsen, and Tore Olsen**, "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *Quarterly Journal of Economics*, 2014, *129* (3), 1141–1219.

Crocker, Keith J. and Kenneth J. Reynolds, "The Efficiency of Incomplete Contracts: An Empirical Analysis of Air Force Engine Procurement," *The RAND Journal of Economics*, 1993, *24* (1), 126–146.

Damgaard, Mette T. and Christina Gravert, "The Hidden Costs of Nudging: Experimental Evidence from Reminders in Fundraising," 2016.

Dasgupta, Sudipto and Daniel F. Spulber, "Managing Procurement Auctions," *Information Economics and Policy*, 1989-1990, *4* (1), 5 – 29.

de Clippel, Geoffroy, Kfir Eliaz, and Kareen Rozen, "Competing for Consumer Inattention," *Journal of Political Economy*, 2014, *122* (6), 1203–1234.

de la Rosa, Leonidas E., "Relational Incentives and Moral Hazard in Teams," *Games and Economic Behavior*, 2011, *73* (2), 429–451.

Decarolis, Francesco, "Awarding Price, Contract Performance, and Bids Screening: Evidence from Procurement Auctions," *American Economic Journal: Applied Economics*, 2014, *6* (1), 108–132.

DeGroot, Morris H., "Reaching a Consensus," *Journal of the American Statistical Association*, 1974, *69* (345), 118–121.

DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 2009, *47* (2), 315–372.

——— **and Ulrike Malmendier**, "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics*, 2004, *119* (2), 353–402.

——— **and** ——— , "Paying Not to Go to the Gym," *American Economic Review*, 2006, *96* (3), 694–719.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions," *Quarterly Journal of Economics*, 2003, *118* (3), 909–968.

Dixit, Avinash K. and Jörgen W. Weibull, "Political Polarization," *Proceedings of the National Academy of Sciences*, 2007, *104* (18), 7351–7356.

Drew, Philip, *The Masterpiece: Jørn Utzon, A Secret Life*, Richmond (VIC), Australia: Hardie Grant Books, 1999.

Duarte, Fabian and Justine S. Hastings, "Fettered Consumers and Sophisticated Firms: Evidence from Mexico's Privatized Social Security Market," 2012.

Eliaz, Kfir and Ran Spiegler, "Contracting with Diversely Naive Agents," *Review of Economic Studies*, 2006, *73* (3), 689–714.

Englmaier, Florian, Matthias Fahn, and Marco A. Schwarz, "Long-Term Employment Relations When Agents are Present Biased," *working paper*, 2016.

Enke, Benjamin, "What You See is All There is," *working paper*, 2015.

——— **and Florian Zimmermann**, "Correlation Neglect in Belief Formation," *working paper*, 2015.

Ericson, Keith M., "When Consumers Do Not Make an Active Decision: Dynamic Default Rules and their Equilibrium Effects," 2014.

Eyster, Erik and Georg Weizsäcker, "Correlation Neglect in Financial Decision-Making," *working paper*, 2011.

____ and Matthew Rabin, "Naïve Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, 2010, *2* (4), 221–43.

____ and ____ , "Extensive Imitation is Irrational and Harmful," *Quarterly Journal of Economics*, 2014, *129* (4), 1861–1898.

Fahn, Matthias, "Minimum Wages and Relational Contracts," *working paper*, 2016.

fai Fong, Yuk and Jin Li, "Relational Contracts, Limited Liability, and Employment Dynamics," *working paper*, 2016.

Falck, Oliver, Robert Gold, and Stephan Heblich, "E-lections: Voting Behavior and the Internet," *American Economic Review*, 2014, *104* (7), 2238–2265.

Farrell, Joseph and Paul Klemperer, "Coordination and Lock-In: Competition with Switching Costs and Network Effects," in "Handbook of Industrial Organization," Vol. 3 2007, chapter 31, pp. 1967–2072.

Fiorina, Morris P., "Has the American Public Polarized?," *Essay Series on Contemporary American Politics*, 2016.

____, Samuel J. Abrams, and Jeremy C. Pope, *Culture War? The Myth of a Polarized America*, New York: Pearson Longman, 2005.

Fudenberg, Drew, Bengt Holmstrom, and Paul Milgrom, "Short-Term Contracts and Long Term Agency Relationships," *Journal of Economic Theory*, 1990, *51*, 1–31.

Gabaix, Xavier and David Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Gagnon-Bartsch, Tristan and Matthew Rabin, "Naive Social Learning, Mislearning, and Unlearning," *working paper*, 2016.

Gale, Douglas and Shachar Kariv, "Bayesian Learning in Social Networks," *Games and Economic Behavior*, 2003, *45* (2), 329–346.

Galeotti, Andrea and Sanjeev Goyal, "The Law of the Few," *American Economic Review*, 2010, *100* (4), 1468–1492.

Galperti, Simone, "Commitment, Flexibility, and Optimal Screening of Time Inconsistency," *Econometrica*, 2015, *83* (4), 1425–1465.

Ganuza, Juan-José, "Competition and Cost Overruns in Procurement," *Journal of Industrial Economics*, 2007, *55* (4), 633–660.

Garavan, Thomas N. and Michael Morley, "The Socialization of High-Potential Graduates into the Organization," *Journal of Managerial Psychology*, 1997, *12* (2), 118–137.

Gentzkow, Matthew and Jesse M. Shapiro, "Ideological Segregation Online and Offline," *Quarterly Journal of Economics*, 2011, *126* (4), 1799–1839.

Gervais, Simon, J.B. Heaton, and Terrance Odean, "Overconfidence, Compensation Contracts, and Capital Budgeting," *Journal of Finance*, 2011, *66* (5), 1735–1777.

Gilpatric, Scott M., "Present-Biased Preferences, Self-Awareness and Shirking," *Journal of Economic Behavior & Organization*, 2008, *67* (3), 735–754.

Glaeser, Edward L. and Cass R. Sunstein, "Why Does Balanced News Produce Unbalanced Views?," *working paper*, 2013, p. 18975.

Golub, Benjamin and Matthew O. Jackson, "Naive Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2010, *2* (1), 112–149.

___ and ___ , "How Homophily Affects the Speed of Learning and Best-Response Dynamics," *Quarterly Journal of Economics*, 2012, *127* (3), 1287–1338.

Gottlieb, Daniel, "Competition over Time-Inconsistent Consumers," *Journal of Public Economic Theory*, 2008, *10* (4), 673–684.

Grubb, Michael D., "Consumer Inattention and Bill-Shock Regulation," *Review of Economic Studies*, 2015, *82* (1), 219–257.

___ and Matthew Osborne, "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," *American Economic Review*, 2015, *105* (1), 234–271.

Handel, Benjamin R., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

Harbridge, Laurel and Neil Malhotra, "Electoral Incentives and Partisan Conflict in Congress: Evidence from Survey Experiments," *American Journal of Political Science*, 2011, *55* (3), 494–510.

Harstad, Ronald M. and Michael H. Rothkopf, "Withdrawable Bids as Winner's Curse Insurance," *Operations Research*, 1995, *43* (6), 983–994.

Heidhues, Paul and Botond Kőszegi, "Exploiting Naïvete about Self-Control in the Credit Market," *American Economic Review*, December 2010, *100* (5), 2279–2303.

___ and ___ , "Naivete-Based Discrimination," 2015.

___ , ___ , and Philipp Strack, "Unrealistic Expectations and Misguided Learning," *working paper*, 2016.

___ , ___ , and Takeshi Murooka, "Exploitative Innovation," *American Economic Journal: Microeconomics*, 2016, *8* (1), 1–23.

___ , ___ , and ___ , "Inferior Products and Profitable Deception," 2016.

Herweg, Fabian and Daniel Müller, "Performance of Procrastinators: On the Value of Deadlines," *Theory and Decision*, 2011, *70* (3), 329–366.

___ and Klaus M. Schmidt, "Auctions vs. Negotiations: The Effects of Inefficient Renegotiation," 2016.

Hetherington, Marc J. and Jonathan D. Weiler, *Authoritarianism and Polarization in American Politics*, New York: Cambridge University Press, 2009.

Hoffman, Mitchell and Stephen Burks, "Training Contracts, Worker Overconfidence, and the Return to Firm-Sponsored General Training," *working paper*, 2015.

Hojman, Daniel A. and Adam Szeidl, "Core and Periphery in Networks," *Journal of Economic Theory*, 2008, *139* (1), 295–309.

Holmstrom, Bengt, "Managerial Incentives Problems - A Dynamic Perspective," *Review of Economic Studies*, 1999, *66*, 169–182.

Jackson, Matthew O. and Alison Watts, "The Evolution of Social and Economic Networks," *Journal of Economic Theory*, 2002, *106* (2), 265–295.

___ **and Asher Wolinsky**, "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 1996, *71* (1), 44–74.

___ **and Brian W. Rogers**, "The Economics of Small Worlds," *Journal of the European Economic Association*, 2005, *3* (2-3), 617–627.

Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi, "Non-Bayesian Social Learning," *Games and Economic Behavior*, 2012, *76* (1), 210–225.

___ **,** ___ **, and Alireza Tahbaz-Salehi**, "Information Heterogeneity and the Speed of Learning in Social Networks," *working paper*, 2013.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan, "Self -Control and the Development of Work Arrangements," *American Economic Review Papers and Proceedings*, 2010, *100* (2), 624–628.

___ **,** ___ **, and** ___ , "Self-Control at Work," *Journal of Political Economy*, 2015, *123* (6), 1227–1277.

Keller, Punam A., Bari Harlam, George Loewenstein, and Kevin G. Volpp, "Enhanced Active Choice: A New Method to Motivate Behavior Change," *Journal of Consumer Psychology*, 2011, *21*, 376–383.

Kőszegi, Botond, "On the Feasibility of Market Solutions to Self-Control Problems," *Swedish Economic Policy Review*, 2005, *12* (2), 71–94.

___ , "Behavioral Contract Theory," *Journal of Economic Literature*, 2014, *52* (4), 1075–1118.

Kosfeld, Michael and Ulrich Schüwer, "Add-On Pricing in Retail Financial Markets and the Fallacies of Consumer Education," 2014.

Kostka, Genia and Niklas Anzinger, "Large Infrastructure Projects in Germany: Between Ambition and realities," 2015.

Krishna, Vijay, *Auction Theory*, Heidelberg: Academic Press, Elsevier, 2010.

Laffont, Jean-Jacques and Jean Tirole, *A Theory of Incentives in Procurement and Regulation*, Cambridge: MIT Press, 1993.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

Layman, Geoffrey C. and Thomas M. Carsey, "Party Polarization and "Conflict Extension" in the American Electorate," *American Journal of Political Science*, 2002, *46* (4), 786–802.

Le Yaouang, Yves, "A Model of Ideological Thinking," *working paper*, 2016.

Lelkes, Yphtach, Gaurav Sood, and Shanto Iyengar, "The Hostile Audience: Selective Exposure to Partisan Sources and Affective Polarization," *American Journal of Political Science*, forthcoming.

Levy, Gilat and Ronny Razin, "Correlation Capacity, Ambiguity and Group Shifts," *working paper*, 2015.

—— **and** ——, "Correlation Neglect, Voting Behavior, and Information Aggregation," *American Economic Review*, April 2015, *105* (4), 1634–45.

—— **and** ——, "Does Polarisation of Opinions Lead to Polarisation of Platforms? The Case of Correlation Neglect," *Quarterly Journal of Political Science*, 2015, *10* (3), 321–355.

—— **and** ——, "The Coevolution of Segregation, Polarised Beliefs and Discrimination: The Case of Private vs. State Education," *working paper*, 2016.

—— **and** ——, "Information Diffusion in Networks with the Bayesian Peer Influence Heuristic," *working paper*, 2016.

Malcomson, James M. and Frans Spinnewyn, "The Multiperiod Principal-Agent Problem," *Review of Economic Studies*, 1988, *55*, 391–408.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal, *Polarized America: The Dance of Ideology and Unequal Riches (Second Edition)*, Cambridge: MIT Press, 2016.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook, "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, 2001, *27*, 415–444.

Miller, David A. and Joel Watson, "A Theory of Disagreement in Repeated Games with Bargaining," *Econometrica*, 2013, *81* (6), 2303–2350.

Möbius, Markus M., Tuan Phan, and Adam Szeidl, "Treasure Hunt," *working paper*, 2015.

Mossel, Elchanan, Allan Sly, and Omer Tamuz, "Strategic Learning and the Topology of Social Networks," *Econometrica*, 2015, *83* (5), 1755–1794.

Mueller-Frank, Manuel, "The Robustness of Consensus in Network-Based Non-Bayesian Learning Models," *working paper*, 2011.

Mullainathan, Sendhil, Markus Nöth, and Antoinette Schoar, "The Market for Financial Advice: An Audit Study," 2012.

Murooka, Takeshi, "Zero Prices: Optimal Pricing of Experience Goods under Consumer Dynamic Loss Aversion," 2016.

Muthoo, Abhinay, *Bargaining Theory with Applications*, Cambridge: Cambridge University Press, 1999.

Myerson, Roger B., "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

Newton, Sidney, Martin Skitmore, and Peter Love, "Managing Uncertainty to Improve the Cost Performance of Complex Infrastructure Projects," in D. Amaratunga, R. Haigh, L. Ruddock, K. Keraminiyage, U. Kulatunga, and C. Pathirage, eds., *Construction in a Changing World*, University of Salford Sri Lanka 2014.

Nocke, Volker and Martin Peitz, "Hyperbolic Discounting and Secondary Markets," *Games and Economic Behavior*, 2003, *44* (1), 77–97.

Noel, Hans, *Political Ideologies and Political Parties in America*, New York: Cambridge University Press, 2014.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 1999, *89* (1), 103–124.

—— **and** ——, "Incentives for Procrastinators," *Quarterly Journal of Economics*, 1999, *114* (3), 769–816.

—— and ——, "Choice and Procrastination," *Quarterly Journal of Economics*, 2001, *116* (1), 121–160.

—— and ——, "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes," *American Economic Review: Papers & Proceedings*, 2003, *93* (2), 186–191.

—— and ——, "Optimal Sin Taxes," *Journal of Public Economics*, 2006, *90*, 1825–1849.

Ortoleva, Pietro and Erik Snowberg, "Overconfidence in Political Behavior," *American Economic Review*, 2015, *105* (2), 504–535.

Pariser, Eli, *The Filter Bubble: What the Internet Is Hiding from You*, New York: Penguin Press, 2011.

Pew Research Center, "Political Polarization in the American Public," `www.people-press.org/files/2014/06/6-12-2014-Political-Polarization-Release.pdf` 2014. Accessed: Sept. 18th, 2016.

Phelps, Edmund S. and Robert A. Pollak, "On Second-best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 1968, *35* (2), 185–199.

Piccione, Michele and Ran Spiegler, "Price Competition under Limited Comparability," *Quarterly Journal of Economics*, 2012, *127* (1), 97–135.

Poole, Keith T. and Howard Rosenthal, "The Polarization of American Politics," *Journal of Politics*, 1984, *46* (04), 1061–1079.

—— and ——, "A Spatial Model for Legislative Roll Call Analysis," *American Journal of Political Science*, 1985, *29* (2), 357–384.

—— and ——, *Congress: A Political-economic History of Roll Call Voting*, New York: Oxford University Press, 1997.

Poy, Samuele and Simone Schüller, "Internet and Voting in the Web 2.0 Era: Evidence from a Local Broadband Policy," *working paper*, 2016.

Rabin, Matthew and Joel L. Schrag, "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics*, 1999, *114* (1), 37–82.

Ramey, Garey and Joel Watson, "Contractual Fragility, Job Destruction, and Business Cycles," *Quarterly Journal of Economics*, 1997, *112* (3), 873–911.

Rindfuss, Ronald R., Elizabeth C. Cooksey, and Rebecca L. Sutterlin, "Young Adult Occupational Achievement: Early Expectations Versus Behavioral Reality," *Work and Occupations*, 1999, *26* (2), 220 – 263.

Roelofs, Matthew R., "Common Value Auctions with Default: An Experimental Approach," *Experimental Economics*, 2002, *5* (3), 233–252.

Rogerson, William P., "Repeated Moral Hazard," *Econometrica*, 1985, *53* (1), 69–76.

Rosenblat, Tanya S. and Markus M. Möbius, "Getting Closer or Drifting Apart?," *Quarterly Journal of Economics*, 2004, *119* (3), 971–1009.

Salop, Steven C., "Monopolistic Competition with Outside Goods," *The Bell Journal of Economics*, 1979, *10* (1), pp. 141–156.

Shachat, Jason and Lijia Tan, "An Experimental Investigation of Auctions and Bargaining in Procurement," *Management Science*, 2015, *61* (5), 1036–1051.

Shapiro, Carl, "Optimal Pricing of Experience Goods," *Bell Journal of Economics*, 1983, *14* (2), 497–507.

Shapiro, Jesse M., "Special Interests and the Media: Theory and an Application to Climate Change," *working paper*, 2015.

Shui, Haiyan and Lawrence M. Ausubel, "Time Inconsistency in the Credit Card Market," *working paper*, 2005.

Skiba, Paige M. and Jeremy Tobacman, "Payday Loans, Uncertainty and Discounting: Explaining Patterns of Borrowing, Repayment, and Default," 2008.

Skrzypacz, Andrzej, "Auctions with Contingent Payments – An Overview," *International Journal of Industrial Organization*, 2013, *31* (5), 666 – 675.

Song, Yangbo, "Social Learning with Endogenous Observation," *Journal of Economic Theory*, 2016, *166*, 324–333.

Spear, Stephen E. and Sanjay Srivastava, "On Repeated Moral Hazard with Discounting," *Review of Economic Studies*, 1987, *54* (4), 599–617.

Spiegler, Ran, "On the Equilibrium Effects of Nudging," *Journal of Legal Studies*, 2015, *44* (2), 389–416.

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 1955, *23* (3), 165–180.

Sunstein, Cass R., *Republic.com*, Princeton: Princeton University Press, 2001.

——, *Republic.com 2.0*, Princeton: Princeton University Press, 2009.

Tan, Xu, "A Dynamic Opinion and Network Formation Model," *working paper*, 2015.

Tirole, Jean, "Procurement and Renegotiation," *Journal of Political Economy*, 1986, *94* (2), 235–259.

Van Alstyne, Marshall and Erik Brynjolfsson, "Global Village or Cyberbalkans: Modeling and Measuring the Integration of Electronic Communities," *Management Science*, 2005, *51* (6), 851–868.

Waehrer, Keith, "A Model of Auction Contracts with Liquidated Damages," *Journal of Economic Theory*, 1995, *67* (2), 531 – 555.

Wang, Ruqu, "Bidding and Renegotiation in Procurement Auctions," *European Economic Review*, 2000, *44* (8), 1577 – 1597.

Watts, Alison, "A Dynamic Model of Network Formation," *Games and Economic Behavior*, 2001, *34* (2), 331–341.

Yan, Jianye, "Contracting with Heterogeneous Quasi-hyperbolic Agents: Investment Goods Pricing under Asymmetric Information," *working paper*, 2011.

Yılmaz, Murat, "Repeated Moral Hazard with a Time-Inconsistent Agent," *Journal of Economic Behavior & Organization*, 2013, *95*, 70–89.

——, "Contracting with a Naive Time-Inconsistent Agent: To Exploit or Not to Exploit?" *Mathematical Social Sciences*, 2015, *77*, 46–51.