Essays in Economic Theory

Inaugural-Dissertation

zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Benjamin Schickner

aus Frankfurt am Main

2017

Dekan:	Prof. Dr. Daniel Zimmer
Erstreferent:	Prof. Dr. Benny Moldovanu
Zweitreferent:	JProf. Francesc Dilmé, Ph.D.
Tag der mündlichen Prüfung:	31.01.2017
Ersatzprüfer:	Prof. Dr. Daniel Krähmer

Acknowledgements

Many people supported me in writing this thesis. First, I wish to express my sincere gratitude to my supervisor Benny Moldovanu for his guidance, for fruitful discussions, for enduring support, and for the freedom to pursue my own research interests in various areas of economics. Chapters 2 and 4 of this thesis originate from term papers written for his courses. Further, I would like to thank my supervisor Francesc Dilmé for his extraordinary advice and support in general, and for detailed comments and suggestions on Chapter 3 of this thesis in particular. It was a pleasure to have the two of them as my supervisors. For patiently listening to my research ideas and discussing them with me, I owe many thanks to Dirk Bergemann, Audrey Hu, Andy Kleiner, Daniel Krähmer, Stephan Lauermann, and Sven Rady.

For our joint work, I would like to thank my three friends and coauthors Tim Frommeyer, Holger Herbst, and Linda Schilling. I enjoyed the hours of discussions, derivations, frustration, and breakthroughs and learned a lot from each of them - academically and non-academically. I wish to thank Albin Erlanson, Christian Hilpert, and Thomas Kruse for discussing ideas that ought to be part of another thesis.

For making graduate life much more enjoyable, I thank my fellow (BGSE) students Jan Bergerhoff, Nina Bobkova, Bo Chen, Chunli Cheng, Anna Grodecka, Niccolò Lomys, Ariel Mecikovsky, Matthias Meier, Mirjam and Nazar Salish, Alexander Scheer, and Martin Wolf.

The BGSE provides an excellent and stimulating research environment, as is shown by the fact that all essays of this thesis are based on projects with students from my cohort. I am grateful in particular to Britta Altenburg, Silke Kinzig, Pamela Mertens, and Urs Schweizer for running the graduate school so smoothly and effectively. Moreover, I wish to thank the BGSE for financial support.

For support and for diverting my attention from work whenever necessary, I would like to thank my friends outside of economics.

Last but not least, a big thanks goes to my family and Anastasia for the opportunities they gave me and their continuous support and encouragement.

Contents

1

Introduction

The first-year textbook definition of economics is that it is the study of the allocation of scarce resources. At first, this sounds like an ambitiously broad aspiration. However, when thinking more carefully, it is not only in the areas that are commonly associated with economics, like growth and industrial organization, in which it is important to think about markets, incentives, and the interaction of economic agents, but also in areas that are less often associated with economics, such as health care and kidney exchange. As opposed to other social sciences, economics, and in particular economic theory, draws heavily on a rich set of analytic tools to study these topics; an approach better known from the natural sciences. Although this approach has come under increased scrutiny in recent years, formal models help to cleanly link assumptions with their implications. This can be particularly fruitful in areas in which a rigorous formal analysis has been absent so far.

This thesis comprises three essays which deal with three distinct topics from the aforementioned broad scope of economics. Each essay develops a formal model to contribute to a better understanding of its topic. Chapter 2 analyzes mechanism design in dynamic matching markets; Chapter 3 studies dynamic pricing under heterogeneous social consumption externalities, and Chapter 4 analyzes redistribution in the German health insurance system.

Chapter 2, which is based on joint work with Holger Herbst (see Herbst and Schickner, 2016), deals with mechanism design. In mechanism design, a designer faces strategic agents who have an informational advantage, for example about their ability or their preferences. The designer's task is to design a mechanism which induces a game that has an equilibrium which maximizes the designer's objective, for example his revenue or the agents' welfare. Over the last four decades, mechanism design has successfully been applied in many economic contexts. The context in which we apply it in Chapter 2 is dynamic matching. Dynamic matching plays an important role in many markets, ranging from more traditional markets, such as labor markets, to markets of the new economy, such as online exchanges and even dating websites. Our specific one-sided model fits particularly well to the examples of education in groups, group-lending in the microcredit market, formation of teams and task assignment in firms, and the establishment of partnerships which themselves constitute organizations and firms.

In our model, heterogeneous agents (for example productive and unproductive workers) arrive dynamically over time to a matching market. A designer matches agents pairwise and these subsequently produce output according to a supermodular production function. Because agents are discounting, it is costly for agents to wait in the market. The designer wants to maximize the agents' welfare. In a static setting, in which all agents are present simultaneously, positive assortative matching is both stable and welfare-maximizing. However, the dynamic arrival of agents challenges the optimality of positive assortative matching both from the designer's perspective and the agents' perspective.

To obtain a welfare benchmark, we start by studying the case when the designer can observe the agents' characteristics, i.e. whether they are productive and when they arrive. As the importance of private information may vary across markets, deriving the welfare-maximizing policy is of interest in itself. We show that if complementarities are sufficiently strong or the value of a productive pair is not too high compared to the value of an unproductive pair, then one of three policies is optimal: the Positive Assortative Policy, the Provident Impatient Policy, or the Myopic Impatient Policy. The Positive Assortative Policy only matches agents with the same productivity and stores agents if necessary, i.e. it is the dynamic version of positive assortative matching. In contrast, the Myopic Impatient Policy matches any two subsequent agents independent of their productivity. The Provident Impatient Policy combines elements of both policies. To solve for the optimal matching policy in closed form without imposing any restriction on the policy, we develop a tool, the "State Space Reduction", which we hope is more generally applicable when dealing with an infinite, discrete state space. We show that, for low discount factors, it is optimal to apply the Myopic Impatient Policy. As the discount factor increases there is a gradual transition to the Positive Assortative Policy. Similarly, there is a gradual transition from the Myopic Impatient Policy to the Positive Assortative Policy as complementarities increase. Our findings imply that, first, the fact that Positive Assortative Matching is welfare-maximizing in static matching models is robust to small time frictions and, second, the breakdown of positive assortative matching in search and matching models (cf. Shimer and Smith, 2000) can in fact be welfare-increasing. When the value of a productive pair exceeds the value of an unproductive pair by far and complementarities are weak, it can be optimal to store unproductive agents in the market to ensure that arriving productive agents are paired immediately.

Turning to the case when agents have private information about their productivity, we argue that the time friction also challenges incentive compatibility of positive assortative matching. Nevertheless, we show that, in the parameter region in which each policy is optimal, it is also implementable with a mechanism satisfying a no-deficit condition, incentive compatibility, individual rationality, and efficient exit. For our setup there exists no general implementation result. Furthermore, the impossibility result of Jehiel and Moldovanu (2001) shows that in static models with interdependent values, in particular matching models, it is generically impossible to implement the welfare-maximizing policy. Thus, our positive implementation result extends to a dynamic setting the insight of Dizdar and Moldovanu (2016) that the matching model is a non-generic model in the sense of Jehiel and Moldovanu (2001). We show that if complementarities are sufficiently strong or the discount factor sufficiently high, the welfare-maximizing policy can be implemented with payments that are reminiscent of those that implement the welfare-maximizing policy in the static model. In this sense, not only are the welfare-maximizing policies robust with respect to small time frictions but also the implementation of these policies. Moreover, we identify situations in which the designer can implement the welfare-maximizing policy by adjusting the share that each agent obtains in a match rather than using monetary transfers. Last, we argue that the welfare-maximizing policy is also implementable if agents can hide their arrival in addition to their productivity.

It was noticed as early as Veblen (1899) that when buying a good, buyers not only care about the utility that they derive from the good itself but also about what it symbolizes to others. In turn, what a good symbolizes to others also depends on who buys the good, that is, oftentimes a good is associated with its clientele. This observation is the starting point for Chapter 3 of this thesis, which is based on joint work with Tim Frommeyer (see Frommeyer and Schickner, 2016). Economically, if buyers care about who else buys a good, past buyers impose an externality on current buyers. This externality can be positive if current buyers want to be associated with past buyers, or negative if current buyers do not want to be associated with past buyers. We refer to a summary statistic of the externality that past buyers impose on current buyers as the public reputation of a seller. Whereas public reputation characterizes what a good symbolizes to others, we call the reputation that is traditionally studied in economics private reputation because it refers to the market's belief about private information of the seller, for example the good's quality. In Chapter 3, we are interested in how a seller dynamically manages her (public) reputation and how reputation, clientele, and price dynamics depend on the market's characteristics.

Specifically, we consider a monopolist who repeatedly sells a good to a sequence of continuum populations of heterogeneous buyers. Buyers differ in their type, which captures their effect on the seller's reputation. High-type buyers increase the seller's reputation, whereas low-type buyers decrease the seller's reputation. Thus, in terms of Leibenstein (1950) there are bandwagon effects with respect to high-type buyers and snob effects with respect to low-type buyers. Because the type of a buyer is positively correlated with his willingness to pay, the seller can choose her clientele through prices.

In a general setup, we show that reputation is valuable for the seller and that there exists a trade-off between current-period profits and a higher reputation in subsequent periods. On the one hand, the seller wants to set a very high price to attract a small, exclusive clientele. This drives up reputation in the next period and consequently increases the buyers' willingness to pay, which in turn increases future profits. On the other hand, the seller wants to set an intermediate price, the static monopoly price, to realize high current-period profits. At the optimum, the seller balances these two effects. To study reputation dynamics more explicitly, we use a linear-uniform specification of our model. First, we show that reputation converges to a stable level in the long run. Intuitively, in the long run it is optimal for the seller to have a constant reputation and a fixed clientele. Second, we show that the convergence behavior towards the stable level depends significantly on the good's durability. If the good is less durable, then reputation oscillates towards the long-run reputation: A period of high reputation is followed by a period of low reputation and vice versa. Thus, our model predicts substantial fluctuations in prices, reputation, and clientele if the durability of the good is low. In contrast, if the good is of comparatively high durability, reputation, price, and clientele dynamics are monotone. If reputation is high initially, reputation decreases monotonically to the stable long-run level, and if reputation is low initially, it increases monotonically to the stable long-run level. These findings seem to be in line with anecdotal evidence from various markets. For example, in the market for luxury watches, a market characterized by comparatively high durability, we argue that the reputation and prices of the Swiss manufacturer *Rolex* have steadily increased over the last sixty years. In contrast, we observe that the reputation of producers of fashionable apparel, a good with comparatively low durability, such as *Abercrombie & Fitch* and *Ed Hardy*, fluctuates significantly. In the last part of the chapter, we show that a seller who has a low reputation initially obtains higher profits if her good is of lower durability because this allows her to increase reputation more quickly, which yields higher profits in the future. Therefore, in this case a seller has incentives to produce a good with lower durability. Our finding provides an additional explanation for planned obsolescence from a reputational perspective which complements the traditional demand-driven explanation.

The last chapter, which is based on joint work with Linda Schilling (see Schickner and Schilling, 2016), analyzes redistribution in the German health insurance market. As health costs increase due to technological progress, higher life expectancy, and an aging population, the question arises of how to distribute the burden of financing the health insurance system, and thereby redistribute income, among the population. To understand how the health insurance market redistributes income between different parts of the population, it is key to understand how the health insurance market is organized. Moreover, a careful inspection of the organization of the German health insurance market is needed as there are several policy proposals on the table for changing the health insurance market, with commentators praising the advantage of one proposal or the other in terms of efficiency or equity. While economic theory has extensively studied abstract health insurance markets, in particular the implications of private information, there is relatively little work that analyzes the organizational structure of a particular health insurance system for the purpose of evaluating policy proposals and substantiating descriptive analysis.

Chapter 4 addresses this issue by proposing a formal model of the German health insurance market for analyzing redistribution streams within the system and for evaluating policy proposals for changing the system. In our model, a continuum population of customers who are characterized by health and income seeks health insurance from one of two insurers: first, a budget-balancing public insurer or, second, a more flexible, revenue-maximizing private insurer. Both insurers charge their customers fees to finance their expenditures. Whereas the public insurer charges a fixed percentage, the contribution rate, of income, the private insurer's fee may depend more generally on the customer's income and health.

We argue first that, if health insurance is voluntary, the health insurance market collapses because profitable customers choose to bear their health costs themselves. When health insurance is obligatory, we prove existence of equilibrium. In equilibrium, the private insurer tries to cream skim, i.e. it tries to attract profitable customers and deter unprofitable customers. Redistribution occurs across health and income; profitable public insurance customers with intermediate income subsidize all other public insurance customers. The surplus from profitable, high-income customers is transformed into profits for the private insurer and not available for redistribution. We argue, therefore, that, if a larger part of the population is forced to insure with the public insurer, then the entire population enjoys higher utility. Having analyzed the current German health insurance system, we compare it to a system in which there is only one type of insurer and insurers compete in premia, i.e. fees that are independent of income and health. In this premium-based system redistribution occurs only across health types. Consequently, customers with low income have a lower utility and customer with high income have a higher utility in the premium-based system compared to the current German system. We argue that combining the introduction of a premium-based system with an appropriate, budget-neutral redistribution of income, e.g. through an adjustment of income taxation, increases the population's welfare. Turning to a more careful welfare comparison of the current German system and the premium-based system, we identify two effects that determine the welfare ranking of the two systems. First, redistributing income from richer parts of the population to poorer parts of the population increases welfare. This redistribution effect speaks in favor of the current German system in which fees depend on customers' income. Second, competition among insurers increases the population's welfare. This competition effect speaks in favor of the premium-based system in which insurers compete on equal grounds for customers from the entire population. In the last part of the chapter we study the properties of welfare-maximizing fee schedules by considering a single benevolent insurer who has to finance the population's aggregate health costs. This analysis sheds further light on how to adjust insurers' fees in order to increase the population's welfare.

References

Dizdar, Deniz and Benny Moldovanu (2016): "On the importance of uniform sharing rules for efficient matching." *Journal of Economic Theory*, 165, 106–123. [3]

Frommeyer, Tim and Benjamin Schickner (2016): "Target mass or class? Dynamic reputation management with heterogeneous consumption externalities." *mimeo*. [3]

Herbst, Holger and Benjamin Schickner (2016): "Dynamic formation of teams: When does waiting for good matches pay off?" *mimeo*. [1]

Jehiel, Philippe and Benny Moldovanu (2001): "Efficient design with interdependent valuations." *Econometrica*, 1237–1259. [2, 3]

Leibenstein, Harvey (1950): "Bandwagon, snob, and Veblen effects in the theory of consumers' demand." *The Quarterly Journal of Economics*, 183–207. [3]

Schickner, Benjamin and Linda Schilling (2016): "Redistributional effects of health insurance in Germany: Private and public insurance, premia and contribution rates." *mimeo*. [4]

Shimer, Robert and Lones Smith (2000): "Assortative matching and search." *Econometrica*, 68 (2), 343–369. [2]

Veblen, Thorstein (1899): *The theory of the leisure class: An economic study of institutions*. London: Unwin Books. [3]

2

Dynamic Formation of Teams: When Does Waiting for Good Matches Pay Off?*

This chapter studies the trade-off between realizing match values early and waiting for good matches that arises in a dynamic matching model with discounting. We consider heterogeneous agents that arrive stochastically over time to a centralized matching market. First, we derive the welfare-maximizing assignment rule, which displays the subtle trade-off between matching agents early and accumulating agents to form assortative matches. Second, we show that the welfare-maximizing policy is implementable when agents have private information about their types. The corresponding mechanism satisfies natural requirements. Furthermore, we identify situations in which the designer can abstain from using monetary incentives.

2.1 Introduction

We study a canonical situation in which agents that arrive gradually over time join forces in order to generate output. Agents are heterogeneous and when forming a group their characteristics are complements in the production function. In a static world, when all agents are present from the beginning, positive assortative matchings are both stable and efficient with this kind of production function.[1] The dynamic arrival of agents com-

---

* Based on joint work with Holger Herbst (see Herbst and Schickner, 2016). We wish to thank Dirk Bergemann, Francis Bloch, Francesc Dilme, Jan Eeckhout, Daniel Garrett, Johannes Hörner, Andreas Kleiner, Daniel Krähmer, Stephan Lauermann, Sangmok Lee, Birgit Mauersberger, Benny Moldovanu, Sven Rady, Larry Samuelson, Philipp Strack, Takuro Yamashita, Leeat Yariv, and workshop participants in Bonn, Yale, the TSE Student Workshop (Toulouse), the 26th International Conference on Game Theory (Stony Brook), the Annual Congress of the EEA (Mannheim), and the SFB Young Researchers Workshop (Munich) for helpful comments and suggestions.

[1] This is a well-established result in the matching literature. For a study of necessary and sufficient conditions for positive assortative matchings see Legros and Newman (2002).

bined with impatience, however, poses a challenge to positive assortativeness. If future outcomes are discounted, the desirability of early matches increases both from a social welfare as well as a participating individual's perspective.

This chapter analyzes the emerging trade-off between realizing match values early and waiting for good matches. For this purpose it tackles the question of assortativeness in a centralized dynamic matching market. We address both the welfare-maximizing matching procedures under complete information, and socially optimal mechanisms when agents have private information. We develop a tool that allows us to solve for the welfare-maximizing matching policy in closed form without imposing any restriction on the policy. This provides clear insights into the effects involved. Then we prove implementability of the welfare-maximizing matching policy when agents have private information about their types. Furthermore, we identify situations in which the market organizer can abstain from using monetary incentives. Finally, we address the case in which the agents can, in addition to their private type, hide their arrival.

Applications cover a wide range of situations, including the formation of teams and task assignment within firms, as well as the establishment of partnerships that constitute organizations themselves. Nowadays, output within organizations is mostly created by teams: Examples are consultancy in firms, coauthoring at universities, or team sports in clubs. Complementarities of experts' skill in production processes is clearly illustrated in the well-known O-Ring Theory in Kremer (1993). Employees arrive over time when having finished previous projects or being newly hired. Furthermore, there are various industries in which entrepreneurs team up to found companies; an example is doctors who found group practices to share the burden of the large investment in medical equipment; complementarity arises when patients are risk-averse concerning the quality of treatment and the dynamic friction is that doctors arrive over time in a local market. Our framework also fits situations of education in groups, e.g. language courses, as group member's skills typically are complements and participants arrive over time. Finally, the model applies to the wide-spread practice of group-lending in the market for microcredits: Borrowers who cannot offer collateral obtain loans for individual projects only if they pool default risk with peers conducting independent projects. In this case, a borrower's type is the individual default risk. Following the standard assumption that the default risks of the group members are independent leads to the complementarity of the individual default risks.[2]

Our work aims at closing the gap between static matching models and the literature on search and matching. The growing literature on search and matching studies matching patterns using search models. Each agent from a continuous population meets random fellows one by one and then decides whether to match with that partner or to continue searching. Major contributions are Shimer and Smith (2000), Smith

[2] Explicit models of group-lending with complementarities in individual default risks are for example given in the publications on microcredit group-lending by Ghatak (1999) and Ghatak and Guinnane (1999).

(2006), and Atakan (2006).[3] While search and matching models modify the static matching model by introducing a time and a search friction jointly, we isolate the effect of the time friction. To that end, we study the centralized organization of matching markets using a mechanism design approach.[4]

We consider a population of heterogeneous agents that differ in a binary characteristic. Matched agents jointly produce socially valuable output according to a production function which is supermodular in the agents' characteristics. Once matches are made, they are irrevocable. Following most of the literature, we assume that matchings are pairwise.[5] Agents arrive according to a Poisson process and types are drawn independent of past arrivals. This model is flexible with respect to four key features: The degree of complementarity of the partners' characteristics in the output function, the relative size of absolute values of output generated by two possible matchings of similar agents, the probability distribution of arriving agents' types, and the patience represented by discounting.[6]

The irrevocability of matches may originate from the matched group's need to initially make sunk investments, which make any later split economically unprofitable. Investments may for example be capital investments in medical equipment or advertisement for a newly founded company, or social investments like trust and social arrangements within a group of workers. An alternative view is that pairs simply leave the market and do not return even in case of a split.

For the sake of tractability, we assume characteristics to be binary. This enables us to find a closed-form solution to an otherwise still unsolved optimization problem. Apart from the insights provided by the solution to this problem, both the knowledge of the exact shape of this solution and the solution technique developed in this chapter may help to tackle more comprehensive problems. We refer to the type that generates the higher output when being paired with itself as "productive".

In the first part of this chapter, we derive the welfare-optimal matching policy when the designer can observe both arrivals and the arriving agents' types. As opposed to the literature on search and matching, we do not impose any restrictions on the technology the central organizer may use.[7] This allows us to analyze the role of assortativeness from

---

[3] See also the early contributions by Sattinger (1995), Lu and McAfee (1996), and Burdett and Coles (1997) on two-sided matching. For a literature survey on search and matching models see Smith (2011).

[4] The environment can be interpreted as small in the sense that the arrival process is discrete, representing single agents arriving. In small matching markets central organization is naturally more appropriate than decentralized search models.

[5] For an exception see Ahlin (2015).

[6] More precisely, the combination of the discount rate and the frequency with which arrivals are expected. While in the introduction little discounting between two arrivals is interpreted as patience, a higher frequency of arrivals, which can be interpreted as a larger market, has the same effect.

[7] In the literature on search and matching there is little work on social optimality. Shimer and Smith (2001b) and Shimer and Smith (2001a) study socially optimal policies, however, the search friction is taken as given.

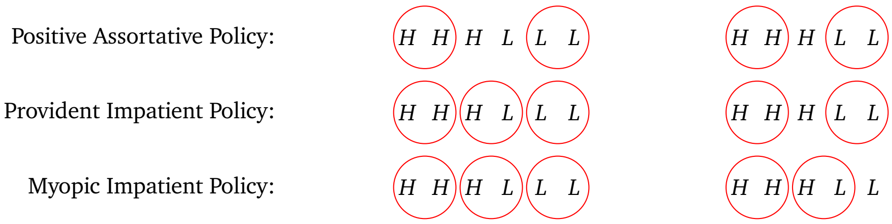

Figure 1. H stands for a productive agent, L represents an unproductive agent, and circles indicate matches. The policies are illustrated by means of two freely chosen sets of agents. The first set of agents consists of three productive agents and three unproductive agents. The second set consists of three productive agents and two unproductive agents.

an efficiency perspective. We obtain the welfare-optimal policy in closed form depending on the four key characteristics. This allows for a clear-cut analysis of the dynamic friction on efficiency.

First, assume that the outputs produced by two productive agents and two unproductive agents do not differ too drastically. Depending on the remaining three key features, always one of three matching policies is optimal. The *Positive Assortative Policy* matches agents with equal types whenever possible and lets every agent wait otherwise. The *Provident Impatient Policy* matches two similar agents with priority whenever possible, but it also matches two unequal types if only those are left. Finally, the *Myopic Impatient Policy* always matches pairs of productive agents. If the number of productive agents is uneven, the remaining one is then matched with priority to an "unproductive" agent. Only then remaining unproductive agents are matched. The differences between the matching policies are illustrated in Figure 1.

The relation between patience and the optimal policies is monotone in the sense that there are two cut-off levels: When discounting is weak, the Positive Assortative Policy is optimal, for intermediate levels of patience the Provident Impatient Policy maximizes welfare, and in an impatient environment it is best to apply the Myopic Impatient Policy. The intuition for this result is perspicuous: The stronger future payoffs are discounted, the less willing is the designer to give up immediate output for the option of realizing gains from positive assortativeness in the future. Whereas the Positive Assortative Policy always respects these options, the Myopic Impatient Policy only maximizes current payoff.

An immediate insight is that positive assortative matchings are not always welfare-maximizing. This means in particular that a failure of positive assortativeness in search models triggered by agents that are "too" impatient to wait for good matches may indeed be welfare-enhancing.[8]

Considered from the opposite point of view, the result tells us that for small rates of discounting the efficient matching pattern resembles the standard pattern from frictionless matching. The dynamic model approaches the static frictionless version when

[8] Examples from the search and matching literature in which supermodular output functions are not sufficient for positive assortativeness are Shimer and Smith (2000) and Smith (2006).

discounting becomes negligible. This implies that the result that positive assortative matchings are efficient in static environments is robust to small dynamic frictions.

The role of the degree of complementarity of the matched agents' characteristics is closely related to the degree of patience. The stronger the complementarity, the greater are the gains from positive assortativeness. Consequently, *ceteris paribus*, for little complementarities the Myopic Impatient Policy is optimal, for intermediate levels the Provident Impatient Policy, and for strong complementarities the Positive Assortative Policy maximizes welfare.

Surprisingly, the relation between the distribution of arriving agents' types and optimal policies may be non-monotone. There are situations in which *ceteris paribus* the Positive Assortative Policy is optimal for intermediate probabilities of arriving agents to be productive but for both small and very high probabilities the Provident Impatient Policy is optimal. For illustrative purposes, consider the state in which there is one agent of either type in the market. When applying the Positive Assortative Policy, both agents are matched with the next arriving peer. When the probability of arriving agents' types takes extreme values, abstaining from the creation of mixed matches implies that one type has large expected waiting costs. The designer might be more willing to enforce the Positive Assortative Policy if he knows that arrivals of both types happen such that all agents get matched in near future.

Next, we consider the situation in which the output produced by two unproductive agents is drastically smaller than the one generated by two productive agents. If in addition complementarities are weak such that the Myopic Impatient Policy is strongly preferred to the Positive Assortative Policy, different matching policies may be welfare-maximizing. It may become optimal to store unproductive types on the market only for the purpose of matching productive agents with them immediately upon arrival and thereby avoiding losses from letting productive agents wait.[9] The reason for the optimality of these policies is the strongly heterogeneous waiting cost induced by the big differences in productivity, together with discounting.

Besides the insights gained from the design of the optimal policies, there is a technical contribution in this chapter. We solve for the welfare-maximizing matching policy using methods from dynamic programming. As the problem has discrete but infinitely many states, a guess and verify approach implies the need to check all possible deviations on infinitely many states. We develop a tool to do an involved form of induction, which we call a "State Space Reduction". This tool enables us to solve the problem via a guess and verify approach by only checking deviations on a small set of states. This tool might well be applicable in a broader range of problems, in particular, it might be used to solve the problem with richer types spaces.

[9] This aspect relates to the basic thought of optimal inventory. See Arrow et al. (1951) for the fundamental thought and Whitin (1954) and Veinott Jr (1966) for early literature surveys.

In the second part of this chapter, we treat implementability of the welfare-maximizing matching policies by a designer who faces agents that have private information and care only about their own matches' output. Match values are split equally. We follow Bergemann and Välimäki (2010) and consider mechanisms that satisfy "efficient exit" and are interim incentive compatible. We show that the welfare-maximizing policy is always implementable if agents have private information about their type but the designer can observe their arrival. This holds even under the most disadvantageous information structure for implementation: Reports are public such that agents in equilibrium have all information about the set of agents in the market when arriving.[10] Note that in our setup no general implementation theorem applies, as agents' values are interdependent and types are uncorrelated in this dynamic setting. While with observable arrivals the implementation of the Provident Impatient Policy and the Myopic Impatient Policy turns out to be generally unproblematic, the possibility result is surprising concerning the Positive Assortative Policy. In static environments, positive assortativeness can be implemented using a single-crossing property with respect to each other agent's type. By this we mean that the gain from matching with a productive agent instead of an unproductive one is higher for agents which are productive themselves. However, the time friction is not only a friction to efficiency but also to the incentive constraints. When the expected time until getting matched depends on the reported type, it might be more attractive for unproductive agents to report being productive than it is for productive agents. Whenever this is the case, the Positive Assortative Policy is not implementable. However, it turns out that whenever this happens, the Positive Assortative Policy is not welfare-maximizing either.

In addition, we show that if the complementarity of the match value function is sufficiently strong, or equivalently the environment is sufficiently patient, the welfare-maximizing policy can be implemented with transfers that only depend on the agent's reported type. This simple structure of payments is reminiscent of the one that implements the positive assortative matching in the static model. Thus, besides the optimal policy, also the implementation in the static matching model is robust to small dynamic frictions. We prove further that whenever this is possible, there exists a splitting rule for the match value of mixed matches such that the optimal policy can be implemented without transfers.

Finally, we address implementation when agents' arrivals are unobservable to the designer. In this case the agents' private information is two-dimensional: It consists of the type and the arrival time. Deviations from truthful reporting may, hence, consist of misreports about the type combined with strategic delays of the report about the arrival. We prove that even in this environment the optimal policy is always implementable. The contract that implements the Positive Assortative Policy when arrivals are observable is also incentive compatible with unobservable arrivals. Concerning the Provident Impatient Policy and the Myopic Impatient Policy, we crucially exploit that for their

[10] In our model this is equivalent to the notion of periodic ex post equilibrium as defined in Bergemann and Välimäki (2010).

implementation the authority does not always need to elicit information about arriving agents' types.

Besides the literature on search and matching, our work relates to further contributions that study the efficiency of positive assortativeness in dynamic matching models. Shi (2005) considers two-sided matching with a supermodular production function where matches get split after random durations. Shi, however, endogenizes one side's quality choice which turns out to make the trade-off that lies at the heart of our work uninteresting in his setting. The focus of his paper is on an additionally introduced coordination friction that cannot be overcome by the central authority. A. Anderson and Smith (2010) analyzes the trade-off between creating a payoff-maximizing positive assortative matching in the current period and having an advantageous distribution over characteristics in the next period. Agents can be rematched each period, but the trade-off arises as the evolution of agents characteristics depends on their current match.

Similar to our model, Baccara et al. (2015) considers a matching market with dynamically arriving agents that is organized by a benevolent central planner. There are three key differences in their setup: The market is two-sided, waiting cost are homogeneous, and each period two agents arrive. A consequence of these differences is that the analytical problem and hence both the solution technique and the optimal matching policies differ remarkably from ours. Furthermore, the authors do not consider implementation with private information. Instead, they are interested in a welfare-comparison between the outcome of a decentralized organization of the market and the socially optimal outcome.

Dynamic matching markets that are organized by a central authority are further studied in the growing literature on dynamic kidney exchange. Such papers include Ünver (2010), Ashlagi et al. (2013), R. Anderson et al. (2015), Akbarpour et al. (2016), and Ashlagi et al. (2016). The objective in these papers is to minimize waiting times and therefore maximize the number of matches respecting restrictions on feasible matches that are exogenously given on medical grounds. Contrary to this literature, we focus on maximizing total match value in an environment in which any match is feasible.

Fershtman and Pavan (2015) studies a centralized two-sided matching market in which the agent's private valuations for partners change over time. The profit-maximizing intermediary faces restrictions on rematching agents. Besides the invariant set of agents in the market, an important difference to our study is that agents have idiosyncratic valuations for partners. First, this means that their paper is not about assortativeness, and second, values are not interdependent.

In a broader perspective, our work adds to the literature on dynamic assignment problems. One strand of this literature considers the assignment of dynamically arriving agents to goods which are present from the beginning. Examples are Gallien (2006), Gershkov and Moldovanu (2009), Mierendorff (2015), Board and Skrzypacz (2015), Dizdar et al. (2011), and Pai and Vohra (2013). Another strand treats the assignment of dynamically arriving goods to agents that are queuing for these goods. Examples are

Leshno (2015) and Bloch and Cantala (2016). The housing literature combines these two strands: Agents arrive over time and are matched with houses that return to the market when the assigned agents have moved out. Examples are Kurino (2014), Bloch and Houy (2012), and Bloch and Cantala (2013). The housing literature and our work share the property that both matching partners arrive over time. There are, however, two substantial differences. First, whereas in the housing literature the arriving stream of houses is determined by the allocation, in our work it is entirely exogenous. Second, in our environment both matching partners have preferences over partners and both have private information.

Finally, we relate to a small literature that treats the implementability of welfare-maximizing social choice functions in general dynamic environments. Bergemann and Välimäki (2010) presents a dynamic VCG mechanism that implements socially optimal social choice functions in dynamic environments with private values. Liu (2014) and Noda (2016) develop dynamic versions of payment schemes found in Cremer and McLean (1985) to implement welfare-optimal social choice functions in environments with interdependent values and correlated types. Nath et al. (2015) extends the idea of Mezzetti (2004) to dynamic settings. Their payment scheme enables implementation of the optimal social choice in general environments if payments are made in periods subsequent to the agent's report. Because in our environment values are interdependent, types are uncorrelated and payments have to be made immediately upon reports, none of the latter results can be applied.

The rest of this chapter is organized as follows: Section 2.2 presents the setup, Section 2.3 derives the welfare-maximizing matching policies, Section 2.4 addresses implementability of the policies under private information, and Section 2.5 concludes.

2.2 Model

We consider agents that arrive over time to a matching market. Time is continuous, and the time horizon is infinite, $t \in [0, \infty)$. Having arrived to the market, agents remain in the market until they are matched, i.e., agents are long-lived. Agents are characterized by the tuple (θ, a), where θ is the agent's type and $a \in [0, \infty)$ is his arrival time. An agent's type reflects his productivity; he is either productive H or unproductive L, $H > L > 0$.

Arrivals are described by a Poisson process $(N_t)_{t\geq 0}$ with arrival rate λ. A Poisson process is a counting process and thus describes discrete arrivals. The random variable N_t describes the number of arrivals up to time t. Let t_n be the time of the n-th arrival. Arriving agents' types are drawn from a Bernoulli distribution that is independent of the process $(N_t)_{t\geq 0}$ and i.i.d. across time; we denote the probability of the productive type by $p \in (0, 1)$. We refer to the process induced by $(N_t)_{t\geq 0}$ joint with the Bernoulli distributions as arrival process.

There exists a central authority, the designer, which organizes the market. Once an agent arrives, the designer may assign him to another agent that is in the market. After being assigned a partner, an agent cannot be reassigned. This could be, for example, because agents leave the market after forming a group and are thus no longer available for the designer or because they make sunk investments that are too costly to forfeit. Together agents produce a match value depending on the pair's types. Formally,

$$\begin{aligned} m \,:\, \mathbb{R}_{>0} \times \mathbb{R}_{>0} &\longrightarrow \mathbb{R}_{\geq 0}, \\ \theta_1 \times \theta_2 &\longmapsto m(\theta_1, \theta_2). \end{aligned}$$

In accordance with the literature, the match value function m is assumed to be symmetric and strictly increasing in both arguments. Given the binary type space, m can be described by three match values m_{LL}, m_{HL}, and m_{HH}, where $m_{\theta_1\theta_2} := m(\theta_1, \theta_2)$.[11] We refer to pairs where both agents have the same type as homogeneous matches; pairs of agents with different types are termed mixed matches. We assume that the match value function is supermodular, which in our setup boils down to $2m_{HL} \leq m_{HH} + m_{LL}$. Supermodularity implies that in a static model, where all agents are present simultaneously, positive assortative matching maximizes the sum of match values.[12] Alternatively, if we regard the two match partners as contributing to the match value, the types, interpreted as input factors, may vary from negligible degrees of complementarity to perfect complementarity. In addition, we require that the match value of the unproductive pair is not too small compared to the match value of the productive pair, $3m_{LL} \geq m_{HH}$.[13]

In the absence of additional payments, an agent's utility from a match, his premuneration value cf. Mailath et al. (2013), equals his share of the match value. In the first part of this chapter, the precise share and the way it is determined, endogenously or exogenously, may be arbitrary. All agents discount future payoffs with a common discount rate $r \in (0, \infty)$.

The designer seeks to maximize the expected discounted sum of match values, i.e., the expected sum of discounted utilities. If we assume that all agents are present from $t = 0$ but only enter the market at their arrival time a, our objective corresponds to maximizing the expected sum of utilities. In yet another, less benevolent, interpretation the designer maximizes output. For a formal description of the designer's objective denote by $a^t = (a_1, a_2, \dots, a_{N_t})$ the history of arrival times up to time t and by $\theta^t = (\theta_1, \theta_2, \dots, \theta_{N_t})$ the corresponding history of types. Let $\varphi_s = (\varphi_s^{HH}, \varphi_s^{HL}, \varphi_s^{LL}) \in \mathbb{Z}^3_{\geq 0}$ be the action taken by the designer at time s, where φ_s^{HH} is the number of homogeneous pairs of productive agents, φ_s^{HL} is the number of mixed pairs, and φ_s^{LL} is the number of homogeneous pairs of unproductive agents formed at time s. Denote by φ^t the history

[11] As will become clear from the analysis below, one could normalize only one of these three values.

[12] A positive assortative matching is a pairing of all agents in the market, in which productive types pair with productive types and unproductive types pair with unproductive types.

[13] In Section 2.3.2 we analyze the case $3m_{LL} < m_{HH}$.

of the designer's actions up to time t, $\varphi^t = \{\varphi_s\}_{0\leq s<t}$. Altogether, a history up to time t, h^t, is given by $h^t = \big(t, a^t, \theta^t, \varphi^t\big)$. Let $\mathscr{H}^t$ be the set of all histories up to time t. A matching policy ρ is a family of functions $\rho = (\rho_t)_{t\in\mathbb{R}_{\geq 0}}$ where ρ_t is defined as

$$\begin{aligned} \rho_t \ : \ \mathscr{H}^t &\longrightarrow \mathbb{Z}^3_{\geq 0} \\ \rho_t(h^t) &\longmapsto \varphi_t. \end{aligned}$$

We write $\rho_t = (\rho_t^{HH}, \rho_t^{HL}, \rho_t^{LL})$, where $\rho_t^{\theta_1\theta_2}$ maps the history h^t into the number of $\theta_1\theta_2$-pairs created at time t. The value generated by this policy at time t after history h^t is

$$v_t^{\rho}(h^t) \ = \ \rho_t^{HH}(h^t)m_{HH} + \rho_t^{HL}(h^t)m_{HL} + \rho_t^{LL}(h^t)m_{LL}.$$

Denote by ϑ_t the corresponding realized value at time t. Let $x(h^t)$ and $y(h^t)$ be the number of productive and unproductive types that are still available in the market given history h^t. Formally, $x(h^t)$ and $y(h^t)$ are given by[14]

$$x(h^t) \ = \ \#\{i \ | \ \theta_i \ \in \ \theta^t, \ \ \theta_i \ = \ H\} - \sum_{\substack{s:\vartheta_s>0 \\ s<t}} (2\varphi_s^{HH} + \varphi_s^{HL}),$$

and

$$y(h^t) \ = \ \#\{i \ | \ \theta_i \ \in \ \theta^t, \ \ \theta_i \ = \ L\} - \sum_{\substack{s:\vartheta_s>0 \\ s<t}} (2\varphi_s^{LL} + \varphi_s^{HL}).$$

We call a matching policy ρ feasible if it never matches more agents than available in the market, i.e., for all t and h^t,

$$\begin{aligned} 2\rho_t^{HH}(h^t) + \rho_t^{HH}(h^t) \ &\leq \ x(h^t), \\ 2\rho_t^{LL}(h^t) + \rho_t^{HL}(h^t) \ &\leq \ y(h^t). \end{aligned}$$

The designer's expected payoff from ρ can be written as

$$\mathbb{E}\Big[\int_{s=0}^{\infty} e^{-rs} v_s^{\rho}(h^s) \ \ \mathrm{d}s\Big], \tag{1}$$

where the expectation is taken over histories with respect to the probability measure induced by ρ and the arrival process. We refer to a feasible matching policy ρ that maximizes (1) as welfare-maximizing. Define the expected payoff at time t after history h^t, i.e., the continuation value, from ρ

[14] Here we slightly abuse notation. By $\theta_i \in \theta^t$ we refer to the components of the vector θ^t.

$$V_\rho(h^t) \;=\; \mathbb{E}[\int_{s=t}^{\infty} e^{-r(s-t)} v_s^\rho(h^s) \;\; \mathrm{d}s|h^t], \tag{2}$$

where the expectation is taken with respect to the probability measure induced by the arrival process and $(\rho_s)_{s\geq t}$. We call a feasible matching policy ρ optimal if it maximizes (2) for any time t and any history h^t. Therefore, an optimal matching policy is a refinement of a welfare-maximizing matching policy, requiring maximization of the designer's payoff not only ex ante but after any path of play. In particular, every optimal matching policy is welfare-maximizing.

Recursive Formulation. As a first step, we derive a recursive formulation. Start by noting that at every point in time, after any history, x and y are a sufficient statistic to summarize the maximization problem: Firstly, feasibility depends only on x and y. Secondly, the arrival process is independent of the history. This is a consequence of the memorylessness of the Poisson process, the independence of interarrival times, and the independence of the Bernoulli distributions from the arrival times. Thus, the state of our problem is (x, y) with state space $S := \{(x, y) \in \mathbb{N}^2_{\geq 0} | \, x \geq 0, y \geq 0, x + y > 0\}$.[15]

In particular, time is not part of the state. Thus, there exists an optimal policy that matches only upon agents' arrival: Whenever an action is optimal on a given state, the policy that takes this action on this state is optimal and only matches upon arrival.

A policy that conditions the action, for some state (x, y), on variables different from x and y, e.g. time, is only optimal if all actions that the policy might take in that state are optimal. Vice versa, if there is a unique optimal policy among those that match only upon arrival, then at every state exactly one action is optimal which implies that this policy is the unique optimal policy in the set of all policies.[16]

In the following, we restrict attention to feasible matching policies that condition solely on (x, y). As we will prove generic uniqueness of the optimal policy, the restriction will turn out to be without loss, and we will obtain a sharp characterization of all optimal matching policies. Accordingly, we suppress the dependence on time and history in the notation for continuation values and policies and write them merely as a function of x and y.

It is convenient to define the expected discount factor until the arrival of the next agent

$$\delta \;:=\; \mathbb{E}[e^{-rt_1}] \;=\; \frac{\lambda}{\lambda + r}. \tag{3}$$

[15] Note, that this formulation implies that a state represents the set of agents in the market including the new arrival.

[16] Whenever the optimal policy is unique, there does not exist a stochastic policy which yields a higher payoff.

Building upon the preceding insights, we are thus led to study the Bellman equation

$$\begin{aligned} V(x,y) \;=\; \max_{\varphi^{HH},\ \varphi^{HL},\ \varphi^{LL}} & \big\{\varphi^{HH} m_{HH} + \varphi^{HL} m_{HL} + \varphi^{LL} m_{LL} \\ & + \delta\big[pV(x-2\varphi^{HH}-\varphi^{HL}+1, y-2\varphi^{LL}-\varphi^{HL}) \\ & + \ (1-p)V(x-2\varphi^{HH}-\varphi^{HL}, y-2\varphi^{LL}-\varphi^{HL}+1)\big]\big\}, \\ \text{subject to:} & \\ & 2\varphi^{HH}+\varphi^{HL} \le x, \quad 2\varphi^{LL}+\varphi^{HL} \le y, \quad \forall x,y \in S. \end{aligned} \tag{4}$$

Given the recursive formulation of the designer's problem (4), the difference between optimal matching policies and welfare-maximizing matching policies has an intuitive interpretation. Every matching policy defines, together with the arrival process, a Markov chain over states in S. As matching policies are deterministic, there are two possible successors for each state with transition probabilities p and $1-p$. Depending on the policy, this Markov chain might have a recurrent set. If the Markov chain has a recurrent set and the market is initially, i.e., at the point in time at which the policy is enforced, in a state within the recurrent set, then no state outside the recurrent set is reached with positive probability. A policy is then welfare-maximizing if it maximizes the designer's expected payoff at all states in the recurrent set. The criterion welfare maximization does not restrict policies on states outside of the recurrent set. An optimal matching policy maximizes the designer's expected payoff at every state in S. In the following, we solve for optimal policies.[17] Firstly, this allows us to solve the designer's problem for any initial state of the market. Secondly, an optimal policy maximizes the designer's payoff even if in the past several agents arrived at the same point in time. Also, as we will see below, focusing on optimal policies elucidates the economic trade-offs connected to welfare maximization and their resolution.

2.3 Optimal Policy

We start by considering the welfare maximization problem under complete information. This means at the point in time of an agent's arrival, the authority observes both the agent's arrival and his type. Incentive constraints arising from the agents' informational advantage are added in Section 2.4.

2.3.1 Regular Case

By means of the assumption $m_{HH} \le 3m_{LL}$ we have guaranteed that the match value of a pair of productive agents is not extremely high compared to the output generated by two unproductive agents. In this case there are three important policies to be considered, which are portrayed in the following.

[17] As will become clear from the analysis below, solving for optimal policies is indispensable even if we are only interested in welfare-maximizing policies. See the outline of the proof of Theorem 1 for an explanation.

Definition 1. ***Positive Assortative Policy***
The Positive Assortative Policy creates in each state (x, y) *the maximal number of homogeneous pairs of both kinds of agents. Mixed matches are never created.*

The matching pattern produced by this policy is positive assortative. In the absence of discounting, this policy maximizes the overall match value. Whenever there is exactly one agent of either kind left in the market, the policy lets both agents wait despite waiting costs.

Definition 2. ***Provident Impatient Policy***
The Provident Impatient Policy creates in each state (x, y) *the maximal number of homogeneous pairs of both kinds of agents. If both* x *and* y *are uneven, one mixed match is created in addition.*

The matching pattern produced by this policy is not positive assortative but contains a relatively large number of homogeneous matches. Only when there is exactly one mixed pair left in the market after all homogeneous pairs have been created, the mixed match is created as well. In that sense homogeneous matches are given priority over mixed matches. However, the policy never lets two agents wait.

Definition 3. ***Myopic Impatient Policy***
The Myopic Impatient Policy creates in each state (x, y) *the maximal number of pairs of productive agents. If* x *is uneven and* $y \geq 1$*, one mixed match is created. The maximal number of pairs from the pool of remaining unproductive agents is formed.*

The matching pattern produced by this policy is not positive assortative and contains a relatively little number of homogeneous matches. Again, the policy first creates the maximal number of productive matches. If afterwards there is a productive agent left, it is matched to an unproductive agent with priority. Only then pairs of unproductive agents are formed. In that sense productive agents are given priority over homogeneous matches. The policy acts entirely myopic in the sense that it maximizes the sum of immediate match values. The three matching policies are illustrated in Figure 1 in the introduction.

In order to state the main theorem, we define two functions that partition the space of parameter constellations $(p, \delta, m_{HH}, m_{LL}, m_{HL})$. Define

$$m_{HL}^1 \;:=\; m_{HH}\frac{\delta p}{1-\delta(1-2p)} + m_{LL}\frac{\delta(1-p)}{1+\delta(1-2p)} \tag{5}$$

and

$$m_{HL}^2 \;:=\; m_{HH}\frac{\delta p}{1-\delta(1-2p)} + m_{LL}\frac{1-\delta(1-p)}{1-\delta(1-2p)}. \tag{6}$$

Note that for any $m_{HH}, m_{LL}, m_{HL}, \delta$ and p it holds that $m_{HL}^1 < m_{HL}^2$.

Theorem 1. *For any given parameter constellation* $(p, \delta, m_{HH}, m_{LL}, m_{HL})$ *one of three matching policies is optimal:*

If $m_{HL} \leq m_{HL}^1$, *the Positive Assortative Policy is optimal.*
If $m_{HL} \in [m_{HL}^1, m_{HL}^2]$, *the Provident Impatient Policy is optimal.*
If $m_{HL} \geq m_{HL}^2$, *the Myopic Impatient Policy is optimal.*
The optimal policy is unique if $m_{HL} \notin \{m_{HL}^1, m_{HL}^2\}$.

In the following we first elaborate on the statement of Theorem 1. Then we provide an economic intuition for the result and the underlying trade-offs. The treatment of the theorem is completed with an outline of the proof.

Immediate implications. First, Theorem 1 is surprisingly simple in the sense that only three policies are generically optimal. In particular, it is generically not optimal to let two unproductive agents wait when there is nobody else in the market. As we show in the extension to this section, this crucially hinges on the regularity assumption $3m_{LL} \geq m_{HH}$.

The second comment addresses the difference between optimal and welfare-maximizing policies. The Provident Impatient Policy and the Myopic Impatient Policy do not differ on states which are reached once one of the two policies is established. Once they are installed, the total number of agents in the market never exceeds two and both policies form a pair whenever possible, which is each second arrival. We call the set of policies which always form a pair as soon as two agents are present the set of *Impatient Policies*. For a discussion why to focus on optimal policies we refer the reader to Section 2.2.

Intuition. There are two major driving forces. We refer to the first one as the *gain of assortative matching*. The value of matching two productive agents and two unproductive agents is higher than the value of two mixed matches. In order to achieve positive assortative matchings, it might be necessary to accumulate agents. This is the case whenever pairing agents in the order of their arrivals induces creating mixed pairs. Hence, the gain of assortative matching is effective in favor of waiting with agents in the market. The second force is the *loss from deferring matches*. Having agents wait in the market is costly because of discounting. This force is effective in favor of creating match values early; this may include creating mixed pairs.

We first comment on the influence of complementarity of the match value function on optimal policies. Given all other parameters, the match value of a mixed pair characterizes the degree of supermodularity of the match value function in the agents' types. The larger m_{HL}, the more are the partners' types substitutes, and perfect substitutability is achieved at the upper bound $m_{HL} = 1/2(m_{HH} + m_{LL})$. The theorem states that the relation between the match value of a mixed pair and the optimal policies is monotone in the sense that there are two cut-off levels: When complementarity is strong, the Positive Assortative Policy is optimal, for intermediate levels the Provident Impatient

Policy is optimal, and for high degrees of substitutability it is optimal to apply the Myopic Impatient Policy. The intuition is perspicuous: The higher the value of mixed matches, the smaller is the gain from assortative matching. As all other parameters are kept fixed, the loss from deferring matches is unchanged. Hence, with increasing match values for mixed pairs the losses from deferring matches more and more outweigh the gains from positive assortativeness. Whereas the Positive Assortative Policy fully realizes gains from positive assortativeness, the Myopic Impatient Policy solely prevents losses from deferring matches.

Second, for given match values we discuss how the choice of optimal policies depends on the impatience, represented by the discount factor δ. Recall that δ is a compound expression of the arrival rate and the discounting rate. The discount factor increases when the discount rate decreases, or when the arrival rate increases. The arrival rate of agents can be interpreted as the size of the market. We obtain the following consequence of Theorem 1:

Corollary 1. *There exist two functions δ^1 and δ^2 that map any parameter constellation $(p, m_{HH}, m_{LL}, m_{HL})$ into $[0,1]$ with $\delta^1 > \delta^2$ such that:*

If $\delta \geq \delta^1$, the Positive Assortative Policy is optimal.

If $\delta \in [\delta^2, \delta^1]$, the Provident Impatient Policy is optimal.

If $\delta \leq \delta^2$, the Myopic Impatient Policy is optimal.

We call δ^1 and δ^2 cut-off levels. The existence of two cut-off levels follows from showing that $\partial m^1_{HL}/\partial\delta \geq 0$ and $\partial m^2_{HL}/\partial\delta \geq 0$ independent of the specific choice of parameters. This means that when fixing all parameters but δ, the monotonicity in m_{HL} carries over to monotonicity in δ. $1 \geq \delta^1 \geq \delta^2 \geq 0$ implies that given any parameter constellation $(p, m_{HH}, m_{LL}, m_{HL})$, for each of the three policies there exists a δ such that the policy is optimal.

Implication 1: The relation between patience and the optimal policies is also monotone. The stronger discounting, the greater are the losses from deferring matches. Hence, the stronger future payoffs are discounted, the less willing is the designer to give up immediate output for the option of realizing gains from positive assortativeness in the future.

Implication 2: The result tells us that for little rates of discounting the efficient matching pattern resembles the standard pattern from frictionless matching. More precisely, note that for $\delta = 1$ holds $m^1_{HL} = 1/2(m_{HH} + m_{LL})$, which is the maximum value for m_{HL}, and whenever m^1_{HL} is interior, then δ^1 is interior as well. m^1_{HL} and thus δ^1 is interior for all $p, \delta \neq \{0,1\}$. From an economic perspective, the dynamic model approaches the static frictionless model when discounting gets negligible. Combining the above statements, this implies that the efficiency result of positive assortative matchings in static environments is robust to small dynamic frictions.

Third, we examine the role of the distribution of arriving agents' types represented by p. Surprisingly, there are situations in which *ceteris paribus* the Provident Impatient Policy is optimal for small and high probabilities but for intermediate values of p the Positive Assortative Policy is optimal. This non-monotonicity arises from the presence of two different effects. For an illustration, consider the state in which there is one agent of either type in the market. First, a higher probability of productive arrivals decreases the expected time until the next productive arrival. For the productive agent in the market this means that the likelihood that he can be matched with a productive peer in the near future increases. This raises the expected value of letting the productive agent wait instead of creating a mixed match. This effect implies that the attractiveness of the Positive Assortative Policy is increasing in p, implying that the cut-off δ^1 should decrease in p. However, there is an opposite second effect originating from the unproductive agent. For high levels of p abstaining from the creation of mixed matches implies a large expected waiting time for the unproductive agent. This implies that the value the unproductive agent contributes to is strongly discounted. In isolation, δ^1 should increase in p . In particular, if the values m_{LL} and m_{HH} do not differ much, the second effect is strong enough such that as p approaches one, δ^1 increases in p. Then δ^1 has its minimum at some interior value, decreases for small p, and increases for high values of p. In this case the designer is more willing to enforce the Positive Assortative Policy because he knows that arrivals of both types happen regularly such that all agents get matched in near future.

The dependence of the optimal matching policy on p and δ is graphically illustrated for match values that represent three canonical match value functions: The case of perfect complements (Figure 2), the multiplicative case (Figure 3), and the case of (almost) perfect substitutes (Figure 4). The red line depicts δ^1, the boundary between the parameter regions in which the Positive Assortative Policy and the Provident Impatient Policy are optimal. The blue line depicts δ^2, the boundary between the Provident Impatient Policy and the Myopic Impatient Policy. As mentioned, for large values of δ the Positive Assortative Policy is optimal, for intermediate values the Provident Impatient Policy, and for small values the Myopic Impatient Policy.

In the case of perfect complements, Figure 2, the match value of a mixed match equals the match value of a homogeneous match of unproductive agents. A consequence is that the Provident Impatient Policy dominates the Myopic Impatient Policy in the sense that the expected welfare from the Provident Impatient Policy is weakly higher on each possible state. The reason is that, starting on a given state, both policies generate the same sum of match values in the first period using the same number of agents. However, the Myopic Impatient policy uses weakly more productive agents for it than the Provident Impatient Policy.[18] Furthermore, as the match value function is strongly

[18] Consider, for example, the state $(1, 2)$. The Provident Impatient Policy matches the two unproductive agents, whereas the Myopic Impatient Policy creates a mixed match. Both match values are equal, but the Provident Impatient Policy leaves a productive agent in the market whereas the Myopic Impatient Policy leaves an unproductive agent in the market.

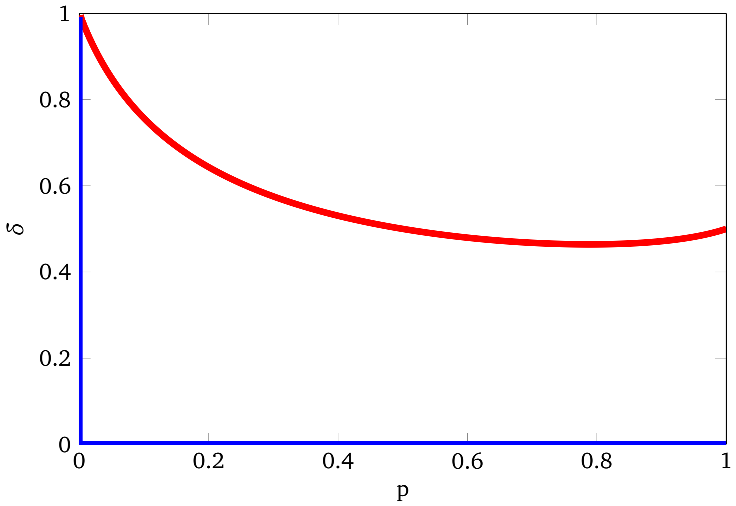

Figure 2. Perfect complements: $m(H, L) = \min\{H, L\}$; Here $H = 3$ and $L = 1$. The red line depicts δ^1; the blue line depicts δ^2.

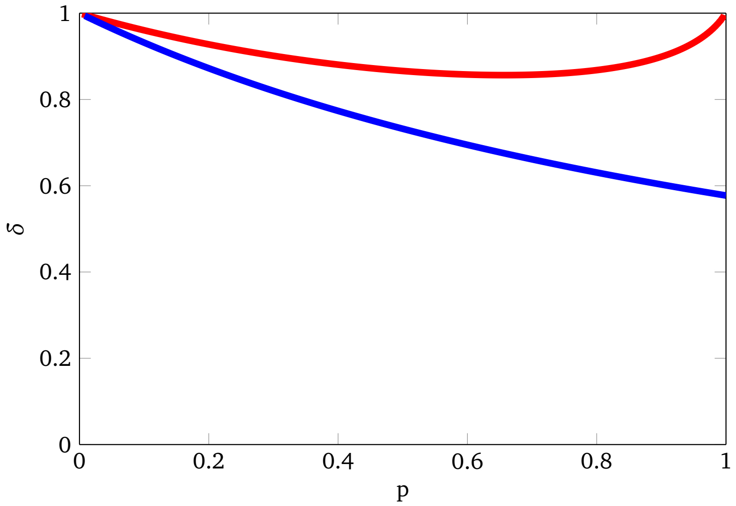

Figure 3. Product: $m(H, L) = H \cdot L$; Here $H = \sqrt{3}$ and $L = 1$.

supermodular, the gain of assortative matching is high, and the parameter region in which the Positive Assortative Policy is optimal is large.

The case of multiplication, illustrated in Figure 3, is regularly used to model complementarities in the match value function. All three policies from Theorem 1 are optimal on a parameter region with positive measure. Note that for p close to one, the Positive Assortative Policy is never optimal - the designer always wants to avoid leaving unproductive agents unmatched.

In the case of perfect substitutes, approximated by Figure 4, there are no complementarities in the match value function. As there are no gains from assortative

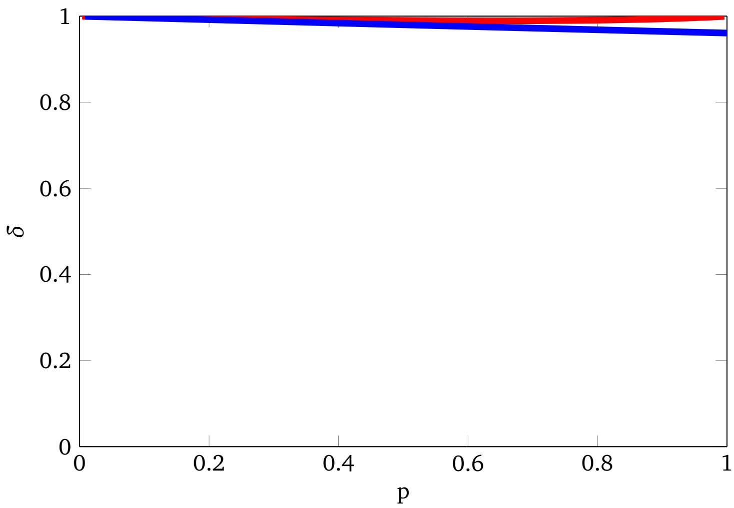

Figure 4. Almost perfect substitutes: Almost $m(H, L) = H + L$; Here $H = 1.5$, $L = 0.5$ and $m(H, L) = 1,98 < H + L$.

matching, the parameter regions, for which the Positive Assortative Policy and the Provident Impatient Policy are optimal, vanish.[19]

Outline of the proof. The problem is solved using a Guess & Verify method. We guess three candidate matching policies, which turn out to be optimal on some subset of the parameter space. As a side product of the verification, we obtain both the precise parameter region where the verified policy is optimal and its (generic) uniqueness on that parameter region. It turns out that the respective parameter regions of the three matching policies constitute a partitioning of the parameter space. The challenging step is the verification, as it involves checking deviations on a discrete but infinite state space. We cope with the situation by developing a procedure we call "State Space Reduction".

Guess. We guess the Positive Assortative Policy, the Provident Impatient Policy, and the Myopic Impatient Policy as candidate matching policies. Remember that joint with the arrival process a matching policy defines a Markov chain over future states S. For all three candidates, the induced Markov chains jumps to a finite recurrent set after the first period. This means that (apart from the initial state) only a finite number of states realize. For the Provident Impatient Policy and the Myopic Impatient Policy only the five states $(1, 0), (0, 1), (0, 2), (2, 0),$ and $(1, 1)$ can occur. For the Positive Assortative Policy the recurrent set is $\{(1, 0), (0, 1), (1, 1), (0, 2), (2, 0), (2, 1), (1, 2)\}$.[20]

[19] From $\partial m^1_{HL}/\partial\delta \geq 0$ and $\partial m^2_{HL}/\partial\delta \geq 0$ joint with $m^1_{HL} = m^2_{HL} = 1/2(m(H, H) + m(L, L))$ at $\delta = 1$ follows that when $m(H, L)$ approaches $1/2(m(H, H) + m(L, L))$, $m(H, L) > m^1_{HL}, m^2_{HL}$ for almost all δ.

[20] Note that the state represents the set of agents in the market including the new arrival.

As these recurrent sets are finite, the value function of each candidate at every state in the respective recurrent set can be computed as the solution to a finite system of equations. The reason is that the value function at an element in the recurrent set only depends on payoffs generated from states in that set and the transition probabilities. For illustrative reasons, we state here the system of equations for the Myopic Impatient Policy. Denote the corresponding values by V_{MIP}:

$$\begin{aligned}
V_{MIP}(0,1) &= \delta[pV_{MIP}(1,1) + (1-p)V_{MIP}(0,2)] \\
V_{MIP}(1,0) &= \delta[pV_{MIP}(2,0) + (1-p)V_{MIP}(1,1)] \\
V_{MIP}(0,2) &= m_{HH} + \delta[pV_{MIP}(1,0) + (1-p)V_{MIP}(0,1)] \\
V_{MIP}(2,0) &= m_{LL} + \delta[pV_{MIP}(1,0) + (1-p)V_{MIP}(0,1)] \\
V_{MIP}(1,1) &= m_{HL} + \delta[pV_{MIP}(1,0) + (1-p)V_{MIP}(0,1)].
\end{aligned} \tag{7}$$

The value functions on states outside the respective recurrent sets can then easily be computed as they only differ from the ones on the recurrent set by the payoff generated in the starting period. We illustrate this by a short example.

Example 1. *The example shows how the value function of the Myopic Impatient Policy at state* $(6,5)$ *can be written as a function of the immediate payoff and the value at a state in the recurrent set:*

$$\begin{aligned}
V_{MIP}(6,5) &= 3m_{HH} + 2m_{LL} + \delta[pV_{MIP}(1,1) + (1-p)V_{MIP}(0,2)] \\
&= 3m_{HH} + 2m_{LL} + \delta V_{MIP}(0,1).
\end{aligned}$$

Policies that differ only on states outside the recurrent set have the same value function at states inside the recurrent set and the recurrent sets coincide. In particular, the values of the Provident Impatient Policy (V_{PIP}) equal those of the Myopic Impatient Policy on the recurrent set. If the initial state lies in the recurrent set, a policy is welfare-maximizing if it maximizes the value on each state in the recurrent set. Any policy that differs outside of the recurrent set is welfare-maximizing as well.[21]

Verification. Verifying the optimality of the three candidates can potentially be cumbersome. We describe the procedure that we apply to all three candidate policies. Fix one candidate. In the previous step, we determined the value on each state. The verification consists of showing for any state that following the policy is better than deviating on the particular state and subsequently following the candidate solution. By the principle of optimality, if all conditions are checked, we have shown both the matching policy, defined as a course of actions, and the associated value function to be optimal. When following the candidate is strictly better than any deviation on any state, the optimal policy is unique.

[21] For the difference of optimality and welfare maximization see Section 2.2.

The difficulty in the verification is that there is a large number of potential deviations to be checked: First, the number of states on which deviations have to be checked is (countably) infinite; second, on states with many agents in the market there is a large number of deviations possible.

This problem can not be avoided by restricting attention to welfare-maximizing policies that start with the very first arrival. Even though in the latter case welfare maximization demands optimality only on a finite set of states, the verification still demands guessing the optimal policy on each possible state. The reasoning is as follows: On each state in the recurrent set, applying the candidate policy must be better than any alternative action. These alternative actions might, however, lead to states outside the recurrent set as illustrated in Example 2.

Example 2. *The example illustrates how checking deviations from the Myopic Impatient Policy at state* $(1, 1)$*, a state in the respective recurrent state, involves the value on states* $(2, 1)$ *and* $(2, 1)$*, which are not in the recurrent set for that candidate. The (only possible) deviation considered here is to not create the mixed pair:*

$$\begin{aligned} V_{MIP}(1,1) \ =& m_{HL} + \delta[pV_{MIP}(1,0) + (1-p)V_{MIP}(0,1)] \\ >& \delta[pV_{MIP}(2,1) + (1-p)V_{MIP}(1,2)]. \end{aligned}$$

In order to evaluate whether the value associated with a deviation is strictly lower than the candidate course of action, we must make statements about the upper bound of the value when being in this "off-path" state. As we want to determine for which parameter configurations precisely the candidate policy is optimal, we need to know exactly the maximal value of the "off-path" state. Therefore, we need to know what the optimal policy does on that state. In order to find the optimal policy on that "off-path" state we guess it and again check deviations, which may lead to more "off-path" states. As there is always the possibility of not matching anything, any state in the state space is reached by the procedure. Hence, it is necessary to set up a candidate that is optimal and verify it.

We tackle the problem of checking the large number of deviations by a proof strategy that involves what we call "State Space Reduction". The State Space Reduction is an elaborate induction argument and is the crucial step of the proof. We believe that the concept of the "State Space Reduction" can be used to address comprehensive maximization problems, in particular the related problem with more than two types of agents. The State Space Reduction works as follows: Instead of checking deviations on each state, we set up a number of general statements. On an arbitrary state, these statements identify deviations, which are not optimal, given the candidate policy will be continued in the following states and given the candidate policy is optimal on smaller states.[22] We call these statements "principles" that specify which kind of deviations do not have to

[22] State (x, y) is smaller than state (x', y') if and only if $x \leq x'$ and $y \leq y'$ with at least one inequality being strict.

be considered. Then we identify states on which these principles capture every possible deviation. We show that there is only a finite number of states including the smallest ones, on which the principles do not capture every possible deviation. When explicitly showing that on this finite set there are no profitable deviations, we have shown that there are no profitable deviations at all. Note, that this set of states does not equal the recurrent set and also differs between the Provident Impatient Policy and the Myopic Impatient Policy. To state the principles, we use one further definition.

Definition 4. *A policy ρ is consistent iff for any state $(x, y) \in S$ and any $(\theta_1, \theta_2) \in \Theta^2$ holds:* $\rho_{\theta_1\theta_2}(x, y) > 0 \Rightarrow \rho_{\theta_1\theta_2}(x - \mathbf{1}_{\{\theta_1=L\}} - \mathbf{1}_{\{\theta_2=L\}}, y - \mathbf{1}_{\{\theta_1=H\}} - \mathbf{1}_{\{\theta_2=H\}}) = \rho_{\theta_1\theta_2}(x, y) - 1$ *and* $\rho_{\theta_3\theta_4}(x - \mathbf{1}_{\{\theta_1=L\}} - \mathbf{1}_{\{\theta_2=L\}}, y - \mathbf{1}_{\{\theta_1=H\}} - \mathbf{1}_{\{\theta_2=H\}}) = \rho_{\theta_3\theta_4}(x, y)$ *for all $(\theta_3, \theta_4) \notin \{(\theta_1, \theta_2), (\theta_2, \theta_1)\}$.*

The meaning of consistency is best illustrated by an example: Suppose on a given state the policy creates matches with at least one pair of productive agents amongst them, $\theta_1 = \theta_2 = H$. Then, on the state with two productive agents less, the policy creates the same matches except for one pair of productive agents. This definition is used in order to formulate the first principle. Each of our candidates is consistent.

Lemma 1. ***(Principle 1)*** *Assume that the candidate policy is consistent. Then in every state, deviations that form a pair that is also formed under the candidate policy, do not have to be checked.*

For example, consider a consistent candidate policy and a state on which it forms a homogeneous pair of productive types. On that state no deviations have to be considered that also match two productive agents.

The reason is based on two observations. First, note that the match value of the pair that is created under the candidate and the deviation can be canceled out from the inequality that corresponds to checking the deviation. Second, due to consistency, the candidate policy creates the same matches except for one pair of productive agents on the state with two productive agents less. Combining the two observations, the deviation is not profitable given there was no profitable deviation on the state with two productive agents less. The same holds for homogeneous pairs of unproductive agents and mixed pairs. The following example illustrates this point:

Example 3. *The example shows how a deviation from the Positive Assortative Policy as in Principle 1 on state $(6, 5)$ can be traced back to a deviation on a smaller state. Denote V_{PAP} the value under the Positive Assortative Policy. The deviation considered matches exactly* *one pair of productive agents:*

$$\begin{aligned} V_{PAP}(6,5) &= 3m_{HH} + 2m_{LL} + \delta[pV_{PAP}(1,1) + (1-p)V_{PAP}(0,2)] \\ &> m_{HH} + \delta[pV_{PAP}(5,5) + (1-p)V_{PAP}(4,6)] \\ \Leftrightarrow \quad & 2m_{HH} + 2m_{LL} + \delta[pV_{PAP}(1,1) + (1-p)V_{PAP}(0,2)] \\ &> \delta[pV_{PAP}(5,5) + (1-p)V_{PAP}(4,6)] \\ \Leftrightarrow \quad & V_{PAP}(4,5) > \delta[pV_{PAP}(5,5) + (1-p)V_{PAP}(4,6)]. \end{aligned}$$

Lemma 2. ***(Principle 2)*** *Assume that the candidate policy is consistent. Consider a state and a deviation on it that leaves two agents with types θ_1 and θ_2 unmatched. If after the following arrival the candidate policy creates a pair (θ_1, θ_2) independent of the arriving agent's type, this deviation does not have to be checked.*

For example, consider a deviation which lets two productive types in the market. The deviation does not have to be checked, if in the next period (upon the next arrival) a match of two productive agents is created independently of the type of arrival. The reason is that a better deviation exists. The more profitable deviation equals the excluded deviation except that the match, which would be made later for sure, is created immediately. Because of discounting, this deviation has a higher value. It suffices to check against the best deviation instead of all deviations.

Lemma 3. ***(Principle 3)*** *In every state, deviations that create more than one mixed pair do not have to be considered.*

The reason is similar to the one of Principle 2. There exists another deviation which is more profitable. The more profitable deviation exploits the supermodularity of the match value function and creates one homogeneous pair of either type instead of two mixed pairs.

In the next step, we apply the principles to all three candidate matching policies. For each candidate we identify the set of states for which all deviations can be excluded.

Lemma 4. *When the Positive Assortative Policy is the candidate, there is no need to verify deviations on all states that contain more than two agents of the same type.*

The proof consists of applying Principles 1 to 3 to the Positive Assortative Policy and then combining them. By Principle 1, deviations that create a homogeneous pair do not have to be checked. By Principle 2, deviations that leave more than one agent of the same type in the market do not have to be checked. The application of Principle 3 to the Positive Assortative Policy is straightforward.

Finally, we combine the applications of the principles: Consider a state with three productive types. Principle 1 implies that there is no need to consider deviations that match two productive types. Principle 2 states that we do not need to treat deviations that leave two or more productive types unmatched. The only deviations left to consider match two or more productive types with unproductive ones. For those deviations, Principle 2 applies. The proof for unproductive agents is analogous.

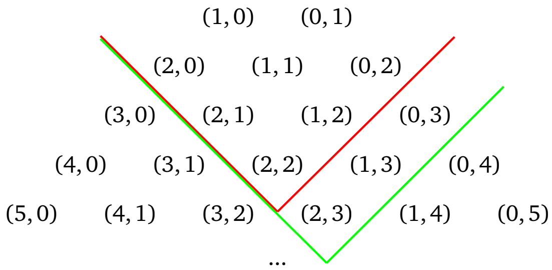

Figure 5. This graph represents the state space. It visualizes Lemmas 4, 5 and 6. Deviations on states below the lines do not have to be verified. The red line corresponds to Lemma 4 and 5, the green line to Lemma 6.

Lemma 5. *When the Provident Impatient Policy is the candidate, there is no need to verify deviations on all states that contain more than two agents of the same type.*

The application of the three principles to the Provident Impatient Policy follows similar thoughts, even though there are differences. There are states in which the candidates do not create homogeneous matches even though this is possible and in addition there are states in which mixed matches are created.

Lemma 6. *When the Myopic Impatient Policy is the candidate, there is no need to verify deviations on all states that contain more than two productive agents or more than three unproductive agents.*

When the Myopic Impatient Policy is the candidate solution, the sets of remaining states that are to be verified one by one is slightly larger than for the other two candidates. The reason is that if two unproductive agents stay in the market, it might happen that one of them is matched with a productive arrival in the next period. Hence, the statement that two unproductive agents get matched in the next period anyways if they are not matched, does not hold.

The deviations on the remaining states are verified one by one. Some deviations are unprofitable only under certain conditions on parameters. These conditions define the region of the parameter space in which a candidate is optimal. It turns out that these regions constitute a partitioning of the entire parameter space. The condition $m_{HL} \leq m^1_{HL}$ ensures that creating the mixed match in $(1,1)$ is unprofitable if the candidate is the Positive Assortative Policy. $m_{HL} \geq m^1_{HL}$ ensures that creating the mixed match in $(1,1)$ is profitable if the candidate is the Provident Impatient Policy. m^2_{HL} takes the corresponding role for the question whether to create the homogeneous match or the mixed match in $(1,2)$: If the candidate is the Provident Impatient Policy, the condition $m_{HL} \leq m^2_{HL}$ ensures that the homogeneous match is optimal, and if the candidate is the Myopic Impatient Policy, the condition $m_{HL} \geq m_{HL}^2$ ensures that the mixed match is value-maximizing.

2.3.2 Extension: Non-regular Case

In this extension we drop the assumption $m_{HH} \leq 3m_{LL}$. This means we allow for extreme differences in the productivity of agents. When the value of unproductive agents is very low compared to the value of productive agents, new matching policies can be optimal. Optimality sometimes requires two unproductive agents to wait in the market, when nobody else is in the market. Therefore we need to define a new class of matching policies.

Definition 5. ***Matching Policy*** $\mathscr{P}_k$
The Matching Policy $\mathscr{P}_k$ creates in each state (x, y) the maximal number of pairs of productive agents. If x is uneven and $y \geq 1$, one mixed matches is created. If a mixed match is created, the maximal number of pairs from the pool of $\max\{0, y-1-k\}$ *unproductive agents is formed. If no mixed match is created, the maximal number of pairs from the pool of* $\max\{0, y-k\}$ *unproductive agents is formed.*

Policy $\mathscr{P}_k$ has similarities to the Myopic Impatient Policy. The difference is that less than the maximum number of homogeneous matches of unproductive agents is created. When there are only unproductive agents in the market, the policy always keeps at least k of them in the market. For example, Policy $\mathscr{P}_1$ creates no match in state $(0, 2)$.

Proposition 1. *For any given m_{LL}, m_{HH} such that $m_{HH} > 3m_{LL}$, there exist parameter constellations (p, δ, m_{HL}) for which matching policy $\mathscr{P}_1$ is optimal.*

In order to proof Proposition 1 we apply the same strategy as for proving Theorem 1. An application of the three principles reduces the verification to a finite set of states on which deviations are verified by hand.

At first glance it might be surprising that it can be optimal to abstain from creating the homogeneous match but to let two unproductive agents wait in the market. To gain intuition for this result, we reconsider the two basic effects in the maximization problem. Consider a situation in which the gain of assortative matching is very small such that the sorting into matches is not decisive. The effect of the loss from deferring matches is then to match agents early with priority given to productive agents. The implication, however, requires careful consideration: Strong discounting does not only make it attractive to match agents which are currently present in the market. It also provides incentives to lay the foundation to quickly match productive agents which arrive in the future. The latter can be achieved by storing an unproductive agent in the market such that a productive agent that arrives in the future can be matched immediately upon arrival. If the value of this availability exceeds the cost of deferring the match of two unproductive agents, it is optimal to abstain from creating a match in order to keep an unproductive agent as stock in the market. Proposition 1 states that there are parameter constellations for which this does happen.

Policy $\mathscr{P}_1$ is optimal when complementarities are low, the differences in the productivity of agents are large, arrivals of productive agents are likely, and discounting is intermediate. A prerequisite for Policy $\mathscr{P}_1$ to be optimal is that there is a strong preference to match a single productive agent with an unproductive agent even if there are two unproductive agents in the market. Therefore, the gain of assortative matching and hence complementarities must be small. Under Policy $\mathscr{P}_1$, two unproductive types incur waiting costs in order to save the waiting cost of a potentially arriving productive type in the future. This is optimal if the productive agent's loss from waiting significantly exceeds the unproductive agents' loss, which is the case only for strong differences in productivity. Furthermore, letting two unproductive agents wait in the market is only optimal if the probability that it pays off, namely that a productive agent arrives in the next period, is sufficiently high. Finally, Policy $\mathscr{P}_1$ is optimal for intermediate values of δ. When discounting is very little, there is no desire to create mixed matches anyways and clearly Policy $\mathscr{P}_1$ is not optimal. The smaller δ, the higher is the value of having unproductive agents available when an productive agent arrives. This dominates the increased cost of accumulating unproductive agents. However, when discounting is very strong, Policy $\mathscr{P}_1$ is not optimal either: The designer does not care about the option to match productive agents earlier, because the option only increases payoffs in the strongly discounted future.

In Figure 6 we fix match values such that for each possible (p, δ) one of the policies listed in Theorem 1 or the Policy $\mathscr{P}_1$ is optimal. The figure illustrates the respective parameter regions (p, δ). The red and the blue line are δ^1 and δ^2 as before. The black line depicts the boundary between the parameter regions on which the Myopic Impatient Policy and Matching Policy $\mathscr{P}_1$ are optimal. Matching Policy $\mathscr{P}_1$ has no boundary to the Provident Impatient Policy. Policy $\mathscr{P}_1$ is optimal for large values of p.

Policy $\mathscr{P}_1$ is optimal for low levels of complementarity. Corollary 2 shows that the monotonicity of optimal policies with respect to m_{HL} extends to Matching Policy $\mathscr{P}_1$.

Corollary 2. *If $m_{HH} > 3m_{LL}$, there exist two continuous functions m_{HL}^3 and m_{HL}^4 that map any parameter constellation $(p, \delta, m_{HH}, m_{LL})$ into $\mathbb{R}_{\geq 0}$ satisfying $m_{HL}^3 < \frac{1}{2}(m_{LL} + m_{HH})$ and $m_{HL}^2 \leq m_{HL}^3 \leq m_{HL}^4$ such that:*

If $m_{HL} \leq m_{HL}^1$, the Positive Assortative Policy is optimal.
If $m_{HL} \in [m_{HL}^1, m_{HL}^2]$, the Provident Impatient Policy is optimal.
If $m_{HL} \in [m_{HL}^2, m_{HL}^3]$, the Myopic Impatient Policy is optimal.
If $m_{HL} \in [m_{HL}^3, min\{m_{HL}^4, \frac{1}{2}(m_{LL} + m_{HH})\}]$, the Matching Policy $\mathscr{P}_1$ is optimal.

A consequence of Corollary 2 is that the condition $m_{HH} \leq 3m_{LL}$ is not only sufficient for the statement of Theorem 1 to hold but also necessary. Whenever $m_{HH} > 3m_{LL}$, there exist parameters (p, δ) such that none of the three matching policies from Theorem 1 is optimal if $m_{HL} = \frac{1}{2}(m_{LL} + m_{HH})$.

Note, that for some parameter constellations $m_{HL}^4 < \frac{1}{2}(m_{LL} + m_{HH})$. This may happen for some values of (p, δ) if the ratio m_{HH} / m_{LL} is extremely high and m_{HL} is close

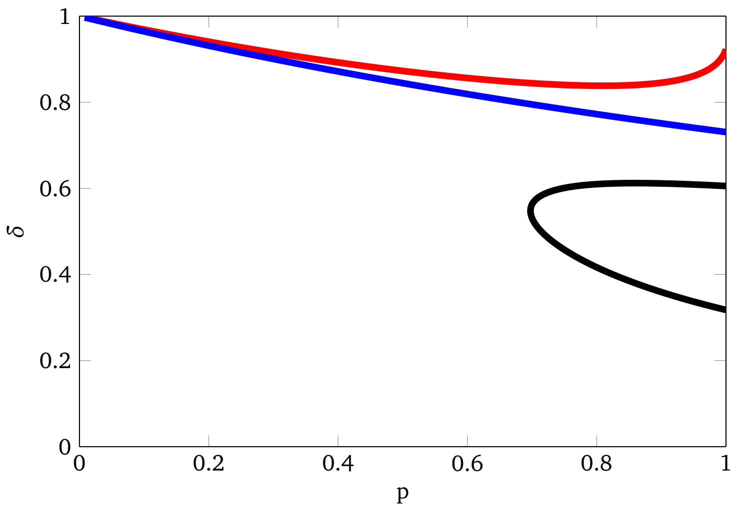

Figure 6. Parameter choice: $m(H,H) = 10, m(L,L) = 1, m(H,L) = 4.8$. The black line depicts the boundary between the parameter regions on which the Myopic Impatient Policy and Matching Policy $\mathscr{P}_1$ are optimal. Matching Policy $\mathscr{P}_1$ has no boundary to the Provident Impatient Policy.

to its upper limit. In that case none of the four matching policies is optimal. Following the logic presented in this extension, our guess is that in these special cases it would be optimal to hold even more than one unproductive agent on stock to prepare for the case that several productive agents arrive to the market in row. Proposition 2 describes an extreme case in which this stock is even infinite.

Proposition 2. *Policy $\mathscr{P}_\infty$ is optimal only if $m_{LL} = 0$.*

Policy $\mathscr{P}_\infty$ never matches two unproductive agents. Potentially, it accumulates an unbounded stock of unproductive agents. When two unproductive agents generate no value, there is no loss of deferring homogeneous matches of unproductive agents. If in addition the creation of mixed pairs is strictly profitable when a single productive agent is in the market, Policy $\mathscr{P}_\infty$ is uniquely optimal.

However, we show that apart from the extreme case $m_{LL} = 0$, Policy $\mathscr{P}_\infty$ is never optimal: This means that generically Policy $\mathscr{P}_\infty$ is not optimal. Hence, if a policy that keeps unproductive agents on stock is optimal and there is a positive cost of waiting with the agents, there is a maximum number of unproductive agents above which two of them get matched. The reason for keeping k unproductive agents on stock is to prepare for the event that k productive agents arrive in row. The probability for this event is exponentially decreasing in k; the cost of holding an additional agent on stock is, however, not decreasing in k. Therefore, at some number of agents the additional cost from accumulating a larger stock exceeds the additional expected profit.

The previous propositions have identified parameter constellations on which none of the three initially introduced policies is optimal. Recall that we interpret the three match values as describing the possible outcomes of a function that maps tuples of types (θ_1, θ_2) into match values $m(\theta_1, \theta_2)$. Despite the results of this extension, there is a large number of natural match value functions for which on each parameter constellation either the Positive Assortative Policy, the Provident Impatient Policy, or the Myopic Impatient Policy is optimal. An important functional form, which is regularly used to represent complementarities in matching, is the product case.

Proposition 3. *Assume $m(\theta_1, \theta_2) = \theta_1 \cdot \theta_2$. For any values of H, L, p and δ one of the following matching policies is optimal: The Positive Assortative Policy, the Provident Impatient Policy, or the Myopic Impatient Policy.*

Note, that this result holds irrespective of the ratio H / L. This means that even when the ratio is large such that $m_{HH} > 3m_{LL}$ unproductive agents are never accumulated.

2.4 Incomplete Information and Implementation

In our model, agents are characterized by their productivity θ and their arrival time a. In Section 2.3, we assumed that the designer can observe agents' characteristics. We now consider situations in which the designer cannot observe agents' entry into the market or their productivity, which are thus private information to the agents. Therefore, the designer needs to elicit private information from agents. As the designer's and agents' interests are not aligned, e.g. either type of agent wants to be assigned a productive partner, the presence of private information gives rise to an incentive problem. In this section we analyze ways of implementing the welfare-maximizing policies under various information structures.

Henceforth, we assume that the match value is divided equally among the two partners. This splitting-rule can be justified as the Nash Bargaining Solution: In our model, the designer assigns two agents to a pair. Once the match is formed, both agents leave the market and can not return. This implies that once they are matched, both partners have an outside option of zero. If both agents have equal bargaining power, they share the surplus of their cooperation, the match value, equally.[23] We allow the designer to use monetary transfers and assume that agents have quasilinear utility. Thus, they maximize (half of the) match value minus payments.[24] We study the market beginning with the first arrival. Therefore, initially the market is in the recurrent set of all optimal policies, and we may focus on the implementation of welfare-maximizing policies. In particular, we only distinguish between the Positive Assortative Policy and

[23] For a strong justification of uniform sharing rules in static settings see also Dizdar and Moldovanu (2016).

[24] Recall that agents discount the future.

the Impatient Policy.

In the following, we will prove a possibility result for the implementation of the optimal policies. To strengthen this result, we consider a setting which impedes implementation. Firstly, we will impose strong requirements on the mechanism. Secondly, agents will draw on a rich information structure, allowing for many deviations. Finally, regarding the designer's information, we assume that he does not observe arriving agents' types. We consider both observable and unobservable arrivals.

2.4.1 Observable Arrivals

First, we analyze the case in which the designer observes arrivals. We consider direct mechanisms in which agents report their type. Upon arrival, an agent observes the past reports of all other agents in the market and reports his type. With a slight abuse of notation we denote by $\mathscr{S}$ the set of agents that are already in the market and by Θ_S the vector containing their types. We adopt the convention to denote reported types with hats. We call the vector $\hat{\Theta}_S$ market report. Given policy ρ, market report $\hat{\Theta}_S$, and an agent entering the market and reporting type $\hat{\theta}$, we denote the (random) variable describing the type of that agent's partner by $\tilde{\theta}^{\rho(\hat{\Theta}_S,\hat{\theta})}$ and the random variable describing the time when the agent will be matched by $t_{\rho(\hat{\Theta}_S,\hat{\theta})}$. A mechanism maps the market report into the allocation given by the welfare-maximizing policy and a payment.

We begin by stating the properties of the mechanism. We concentrate on mechanisms that support *efficient exit* meaning that an agent who stops to affect current and future matches also stops to receive and pay transfers.[25] In particular, payments do not condition on realized match values and agents cannot reveal their partner's type to the designer after being matched. This is in accordance with our interpretation that agents leave the market after forming a group. Thus, we focus on payments $\tau^{\hat{\Theta}_S}(\hat{\theta})$ that are charged upon arrival and depend on the market report and the agent's report.

We study direct mechanisms that have a truthful equilibrium in which welfare is maximized. Agents arrive to the market, observe past reports of all agents in the market, form Bayesian expectations with respect to the future, and maximize their utility given that all other agents report truthfully. Observe that this notion of incentive compatibility coincides with interim incentive compatibility in Bergemann and Välimäki (2010). Formally, the *incentive compatibility* constraints are given by:

$$\frac{1}{2}\mathbb{E}[e^{-rt_{\rho(\Theta_S,\theta)}}m(\tilde{\theta}^{\rho(\Theta_S,\theta)},\theta)] - \tau^{\Theta_S}(\theta) \;\geq\; \frac{1}{2}\mathbb{E}[e^{-rt_{\rho(\Theta_S,\hat{\theta})}}m(\tilde{\theta}^{\rho(\Theta_S,\hat{\theta})},\theta)] - \tau^{\Theta_S}(\hat{\theta}), \;\; \forall\theta,\hat{\theta},\Theta_S, \tag{8}$$

where the expectation is taken with respect to the partner's type and the matching time.[26] There are other perceivable specifications of the agents' information structure

---

[25] See also Bergemann and Välimäki (2010). In the dynamic assignment literature an analogous condition is the requirement that mechanisms are *online*, cf. Gershkov and Moldovanu (2010).

[26] As types do not change over time, (8) implies that the mechanism is periodic ex post incentive compatible in the sense of Bergemann and Välimäki (2010). Thus, the mechanism exhibits the no-regret property

in which agents observe only the number of reports, i.e., the number of agents in the market, or do not observe reports at all. If the designer can implement the welfare-maximizing policies under this information structure, he can implement the welfare-maximizing policies under any information structure in which agents have less information.[27]

As agents participate in the mechanism voluntarily, the following *individual rationality* constraints have to be satisfied

$$\frac{1}{2}\mathbb{E}[e^{-rt_{\rho(\Theta_S,\theta)}}m(\tilde{\theta}^{\rho(\Theta_S,\theta)},\theta)] - \tau^{\Theta_S}(\theta) \ \geq \ 0, \ \ \forall\theta,\Theta_S. \tag{9}$$

Observe that (9) entails a strong notion of individual rationality because, in addition to observing his type, an agent also observes the market report before he decides whether to participate.[28]

As last condition, we impose that the mechanism requires no external injection of money:

$$\tau^{\Theta_S}(\theta) \ \geq \ 0, \ \ \forall\theta,\Theta_S \tag{10}$$

i.e., all payments are positive, which implies that the mechanism runs *no deficit* at any point in time.

Implementation in our setting is not straightforward. Consider the Positive Assortative Policy: An agent that reports a productive type will be assigned a productive partner; an agent that reports an unproductive type will be assigned an unproductive partner. Consider an agent that arrives to a market with one productive agent and one unproductive agent. This situation resembles the static model: Independently of his report, the third agent will be matched immediately with a partner whose report coincides with his report. Either type of arriving agent would like to form a group with the productive agent. Supermodularity of the match value is equivalent to increasing differences, $m_{HH} - m_{HL} \geq m_{HL} - m_{LL}$.[29] Intuitively, increasing differences implies that a match with a productive partner instead of a match with an unproductive partner is valued more by a productive agent than by an unproductive agent. This gap allows the designer to construct payments that make truthful revelation incentive compatible.[30]

---

with respect to past agents' types. Also, as noted by Dizdar and Moldovanu (2016), ex post implementation is intuitively closer to the complete information environment of traditional matching models.

[27] See also Myerson (1986). For example, the designer can reveal any information that is missing to agents and use the original mechanism.

[28] As will become clear from the analysis below the outside option could be any sufficiently small positive value.

[29] This is the discrete analogon of the single-crossing property from mechanism design with continuous types.

[30] If the unproductive agent would have a larger incentive to report the productive type than the productive agent, separation would be possible but would induce both types of agents to lie.

Next, consider an agent that arrives to a market with one unproductive agent. Reporting the productive type is less attractive because of the waiting costs incurred until the arrival of the next productive agent. If he reports the unproductive type, he is matched immediately. Therefore, in this case, supermodularity does not necessarily imply increasing differences. The theorem below establishes that, despite these time constraints, the optimal policy is implementable whenever it is welfare-maximizing.

As mentioned in the introduction, there is no general implementation result for our environment as we have a dynamic setting with interdependent values, independent types, and payments that satisfy efficient exit.[31] Our positive implementation result is also surprising in light of the impossibility result for implementing efficient allocations in static settings with interdependent values, cf. Jehiel and Moldovanu (2001).

Theorem 2. *There exist payments such that the implementation of the welfare-maximizing policies is incentive compatible, individual rational, runs no deficit, and supports efficient exit.*

Impatient Policy. If the designer observes arrivals, the implementation of the Impatient Policy is straightforward. As the policy does not condition on the type but matches every two consecutive agents, irrespective of their types, there is no need to elicit agents' private information. Hence, it is possible to set all payments equal to zero. Furthermore, this is individual rational, runs no deficit, and satisfies efficient exit.

Positive Assortative Policy. In case of the Positive Assortative Policy, the situation is more intricate as the agents' report affects their match. It is useful to aggregate the market report Θ_S into a tuple. Under the Positive Assortative Policy there can be either no agent (0,0), an agent with a productive report (1,0), an agent with an unproductive report (0,1), or two agents with one productive and one unproductive report (1,1) in the market. Denote by $\Delta^{\theta}(p,\delta)$ the expected discount factor until the next arrival of type θ. In accordance with our intuition, $\Delta^{\theta}(p,\delta)$ is increasing in δ, $\Delta^{L}(p,\delta)$ is decreasing in p, and $\Delta^{H}(p,\delta)$ is increasing in p.

The incentive constraint for the productive type given market report $(0,1)$ can be written as

$$\frac{1}{2}\Delta^{H}(p,\delta)m_{HH}-\tau^{(0,1)}(H)\;\geq\;\frac{1}{2}m_{HL}-\tau^{(0,1)}(L). \tag{11}$$

Analogously, the incentive constraint for the unproductive type in $(0,1)$ is

$$\frac{1}{2}\Delta^{H}(p,\delta)m_{HL}-\tau^{(0,1)}(H)\;\leq\;\frac{1}{2}m_{LL}-\tau^{(0,1)}(L). \tag{12}$$

[31] For settings with: (i) private values cf. Bergemann and Välimäki (2010), (ii) correlated types cf. Liu (2014), and (iii) without efficient exit cf. Nath et al. (2015).

Rearranging yields the following condition on the payment difference

$$\frac{1}{2}\big(\Delta^H(p,\delta)m_{HH} - m_{HL}\big) \;\geq\; \tau^{(0,1)}(H) - \tau^{(0,1)}(L) \;\geq\; \frac{1}{2}\big(\Delta^H(p,\delta)m_{HL} - m_{LL}\big). \tag{13}$$

Therefore,

$$\Delta^H(p,\delta)m_{HH} - m_{HL} \;\geq\; \Delta^H(p,\delta)m_{HL} - m_{LL} \tag{14}$$

is a necessary and sufficient condition for the existence of an incentive compatible payment pair in $(0,1)$.

Observe that the left side of (14) decreases quicker than the right side as the excepted discount factor decreases. Especially for low discount factors δ or low values of p, (14) might be violated. Similarly, we can derive conditions for the existence of incentive compatible payments for all other market reports, $(0,0)$, $(1,0)$, $(1,1)$. It turns out, if (14) holds, the conditions for the other market reports are also satisfied. Comparing (14) to the boundary of the Positive Assortative Policy (5), shows that for all parameters for which the Positive Assortative Policy is welfare-maximizing (14) holds. Thus there exist incentive compatible payments that satisfy efficient exit, for all market reports.

We proceed by charging the unproductive type the maximum individual rational payment for all market reports,

$$\tau^{(0,0)}(L) \;=\; \tau^{(1,0)}(L) \;=\; \frac{1}{2}\Delta^L(p,\delta)m_{LL}, \tag{15}$$

$$\tau^{(0,1)}(L) \;=\; \tau^{(1,1)}(L) \;=\; \frac{1}{2}m_{LL}. \tag{16}$$

Given the unproductive type's payment, we choose, for every market report, the maximal payment for the productive type such that the payment pair is incentive compatible, cf. e.g. (13). Individual rationality of the unproductive type and incentive compatibility of the productive type yield individual rationality for the productive type. The proof concludes by verifying that all payments are positive.

The proof of Theorem 2 reveals that if positive assortative matching fails to be incentive compatible, it also fails to be welfare-maximizing. Intuitively, if (14) is violated, the incentive for an unproductive agent to report the productive type is stronger than the incentive for a productive agent. This means that an unproductive agent's gain of being matched with a productive agent instead of being matched with an unproductive agent outweighs the respective loss for a productive agent. Then, it is plausible that it is welfare-maximizing to create mixed matches.

Note that the mechanism constructed in the proof of Theorem 2 generates revenues: Firstly, we set the unproductive type's expected utility to zero for all market reports by charging the highest payment that is individual rational. Secondly, we choose the maximal incentive compatible payment for the productive type. By reducing payments, the mechanism could account for more lucrative outside options.

The payments of the unproductive type that implement the Positive Assortative Policy depend on the market report only through the presence or absence of an agent in the market whose type coincides with the agent's reported type. That is, payments are equal for market reports $(0,0)$, $(1,0)$ and for market reports $(0,1)$, $(1,1)$. This is a consequence of the agent's report fixing his (future) partner's type and his expected waiting costs under the Positive Assortative Policy.

2.4.2 Extension: Simple Payments

The payments that implement the Positive Assortative Policy in Theorem 2 depend on the market report. In applications it is often desirable to use simple mechanisms that condition on as few parameters as possible. In this section, we examine under which conditions the Positive Assortative Policy is implementable with transfers that depend solely on the reported type but not on the market report. We refer to these payments as simple payments. In addition, this elucidates the relation of implementation of positive assortative matching in our dynamic model and in the static analogon. In the static model payments condition only on the agent's own report. Therefore, our analysis investigates when positive assortative matching can be implemented in the dynamic model with "static" payments.

Proposition 4. *The Positive Assortative Policy is implementable with simple payments if*

$$m_{HL} \;\leq\; \Delta^H(p,\delta)\frac{m_{HH}}{2} + \Delta^L(p,\delta)\frac{m_{LL}}{2}. \tag{17}$$

This parameter region is a strict subset of the parameter region where the Positive Assortative Policy is optimal.

In Theorem 2 we proved that the conditions for the existence of incentive compatible payment differences hold. However, the conditions differ across market reports. Therefore, the main issue is to find a single payment pair $(\tau(H),\tau(L))$ that is incentive compatible for all possible market reports. To this end, it is instructive to consider Figure 7. Rows correspond to market reports. The green part of each line marks the region where payment differences are incentive compatible, whereas payment differences that lie within the red region are not incentive compatible. We are looking for a payment difference $\tau(H)-\tau(L)$ which lies in the green interval across all market reports. As the boundaries vary significantly with p and δ, existence of such a payment difference is not guaranteed. As illustrated in Figure 7, the left boundary of market report $(1,0)$ and the right boundary of market report $(0,1)$ are most restrictive. Intuitively, $(1,0)$ is the most attractive market report for reporting the productive type, whereas $(0,1)$ is the most attractive market report for reporting the unproductive type. Combining these two conditions yields (17).

Observe that (4) holds as δ approaches one, that is, as the time friction vanishes. This means that as the time constraints fades, we can draw on simple payments, which

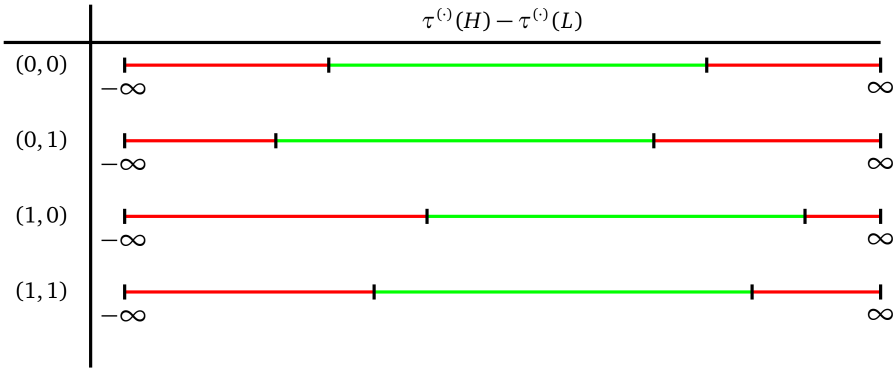

Figure 7. Incentive compatible payment differences

reflect the payments used in the static model to implement the Positive Assortative Policy. Similarly, an increase in complementarities, i.e., a decrease in m_{HL}, strengthens the increasing differences property and thus allows for an implementation with simple payments.

2.4.3 Extension: Asymmetric Match Value Splits

Hitherto, we assumed that partners share their match value equally. While this seems intuitive if partners are homogeneous, i.e., have the same type, one can imagine other sharing rules in case of mixed pairs.[32] An appropriate sharing rule might alleviate the incentive problem.[33] This section investigates under which conditions there exists a sharing rule which induces truthful revelation of types without further intervention, that is to say, without incentivizing agents with payments. In the following, denote by α the productive type's share of the match value when he forms a group with an unproductive type.

Proposition 5. *There exists a share α such that the welfare-maximizing policies are implementable without payments if and only if the welfare-maximizing policies are implementable with simple payments.*

Recall that the implementation of the Impatient Policy in Theorem 2 is straightforward. Hence, we may concentrate on the Positive Assortative Policy. Consider the incentive constraints for market report $(0, 1)$. The productive agent reports truthfully if

$$\frac{1}{2}\Delta^H(p, \delta)m_{HH} \;\geq\; \alpha m_{HL}. \tag{18}$$

[32] The precise way of how these "premuneration values" are determined may depend on the specific legal or institutional environment and may lie beyond the designer's control, see Mailath et al. (2015) for an exhaustive discussion.
[33] A different interpretation is that agents share the match value equally but the designer can prescribe internal transfers within matched pairs.

Analogously, the unproductive agent reports truthfully if

$$\frac{1}{2}m_{LL} \;\geq\; (1-\alpha)\Delta^H(p,\delta)m_{HL}. \tag{19}$$

Observe that the incentive constraint of the productive agent gives an upper bound on α, whereas the incentive constraint of the unproductive agent gives a lower bound on α. We proceed similar for the remaining market reports. As in the previous section, the most restrictive conditions arise in $(1,0)$ and $(0,1)$. In particular, (18) yields the lowest upper bound on α, whereas the incentive constraint of the unproductive type for market report $(1,0)$

$$\frac{1}{2}\Delta^L(p,\delta)m_{LL} \;\geq\; (1-\alpha)m_{HL} \tag{20}$$

yields the highest lower bound on α. Thus, an incentive compatible sharing rule exists if

$$\frac{1}{2}\frac{\Delta^H(p,\delta)m_{HH}}{m_{HL}} \;\geq\; \frac{m_{HL}-\frac{1}{2}\Delta^L(p,\delta)m_{LL}}{m_{HL}}. \tag{21}$$

Reformulating (21) shows that it coincides with (17), which concludes the proof.

To get some intuition for Proposition 5, we reformulate the crucial incentive constraints (18) and (20). The difference between the equal split and the α split can be interpreted as a substitute for payments:

$$\frac{1}{2}\Delta^H(p,\delta)m_{HH} \;\geq\; \frac{1}{2}m_{HL} + \left(\alpha m_{HL} - \frac{1}{2}m_{HL}\right), \tag{22}$$

$$\frac{1}{2}\Delta^L(p,\delta)m_{LL} \;\geq\; \frac{1}{2}m_{HL} - \left(\alpha m_{HL} - \frac{1}{2}m_{HL}\right). \tag{23}$$

Recall that the boundaries on the difference of incentive compatible, simple payments in the proof of Proposition 4 are determined by exactly the same incentive constraints:

$$\frac{1}{2}\Delta^H(p,\delta)m_{HH} \;\geq\; \frac{1}{2}m_{HL} + (\tau(H)-\tau(L)), \tag{24}$$

$$\frac{1}{2}\Delta^L(p,\delta)m_{LL} \;\geq\; \frac{1}{2}m_{HL} - (\tau(H)-\tau(L)). \tag{25}$$

Because α is contained in $[0,1]$, (22) and (23) provide less flexibility than (24) and (25). The value that can be redistributed through a sharing rule is bounded by the total match value that is generated in the mixed match, whereas there is no bound on the payment difference. Hence, if there exists an incentive compatible sharing rule, we can also find an incentive compatible, simple payment difference. Proposition 5 states, however, that the converse is true as well. A conclusion is that for any incentive compatible pair of simple payments, the payment difference never exceeds the total match value. Put differently, incentive compatibility does not require extreme transfer differences.

Inequality (20) implies that for any incentive compatible sharing rule α, it holds $\alpha > \frac{1}{2}$, i.e., the productive type receives a larger share of the match value when he forms a group with an unproductive type. In the Positive Assortative Policy, the mixed group never occurs on path. Thus, changes in the sharing rule only affect the attractiveness of deviations. When the market consists of one productive agent, an unproductive agent's misreport increases the match value that he creates with his partner from m_{LL} to m_{HL} and reduces the time until he is matched. Consequently, in the absence of payments, truth-telling is incentive compatible only if the unproductive agent's share of the output is lower when being matched with an productive agent. Under the condition identified in Proposition 5, also the productive agent's incentive constraints are satisfied, even though he receives a larger share of the match value when misreporting. The reason is that for the productive agent the total value that is shared is smaller if he misreports.

2.4.4 Unobservable Arrivals

Depending on the organizational details of the market, the designer might not observe agents' arrivals to the market. Instead, agents report their arrival to the designer. Given the welfare-maximizing policies, agents may want to exploit this additional source of private information by strategically delaying their arrival report. We maintain the assumption of private types. This renders implementation of the welfare-maximizing policy a multidimensional screening problem. The current section examines conditions under which the designer can overcome this additional challenge and implement the welfare-maximizing policies.

As in Section 2.4.1, we give arriving agents the informational advantage of past reports being public. We focus on incentive compatible, individual rational, direct mechanisms that run no deficit, satisfy efficient exit, and implement the welfare-maximizing policies. We modify the market report to contain the reported types of all agents in the market that have reported their type and arrival. We construct payments that depend on the reported type, the market report, and the reported arrival time. Agents can report their arrival only after actually arriving to the market, and only agents who have reported their arrival may report their type.

Because of substantially different issues, we discuss the Positive Assortative Policy and the Impatient Policy separately.

Positive Assortative Policy. When type spaces have more than one dimension, incentive constraints pose a severe challenge to the design of incentive compatible mechanisms as one has to account for double deviations, i.e., deviations in several dimensions at the same time. Surprisingly, the mechanism constructed in the proof of Theorem 2 also implements the Positive Assortative Policy with unobservable arrivals.

Proposition 6. *There exists an incentive compatible, individual rational mechanism that runs no deficit, supports efficient exit, and implements the Positive Assortative Policy when both, arrivals and types, are private information to the agents.*

Proof. To prove Proposition 6, we show that agents report their type truthfully and that also the timing of the report remains unchanged, which means that agents reveal their arrival immediately, i.e., truthfully.

We naturally adjust the payments constructed in the proof of Theorem 2 to account for the two-dimensional type space: If an agent reveals his arrival and his type at the same time, payments are as in the proof of Theorem 2. If an agent reveals his arrival strictly before his type, he is punished by a payment of m_{HH}.

To tackle the issue of double deviations in the framework of our model, we divide the problem of showing incentive compatibility into two steps:

(i) First, we show that whenever agents report their type, they report truthfully.

(ii) Second, we argue that given agents report their type truthfully, agents report their arrival time truthfully.

By the memorylessness of the Poisson process, the incentive problem faced by an agent at an arbitrary point in time is the same as the incentive problem at the time of the last arrival. This latter problem, however, resembles the incentive problem with observable arrivals. As the payments solve the incentive problem with observable arrivals, we deduce that (i) holds.

Now, we prove (ii). Given our specification of payments, agents report arrival time and type simultaneously. By the first step, agents report their type truthfully. It remains to be shown that agents want to report their arrival as early as possible. Under the Positive Assortative Policy, an agent's report fixes his match partner's type. The agent's partner is, depending on the market report, either an agent with the same type that is already in the market, or the next agent of his type that arrives to the market. By memorylessness of the Poisson process, delaying an arrival may only be profitable for an agent if the market report changes compared to the market report at the arrival time. By our choice of payments, the unproductive type receives zero expected utility upon arrival for every market report. Therefore, it is an optimal strategy for the unproductive agent to report his arrival time truthfully. Given our payments, the expected utility of a productive type at the point of his arrival is $\Delta^L(p,\delta)(m_{HL} - m_{LL})$ for market reports $(0,0)$ and $(1,0)$, and $m_{HL} - m_{LL}$ for market reports $(0,1)$ and $(1,1)$. We see that the productive agent's expected utility is highest if an unproductive agent is already in the market. Thus, if the productive agent arrives in $(1,1)$ or $(0,1)$, he reports his arrival immediately. On the other hand, if the productive agent arrives in $(0,0)$ or $(1,0)$, he might consider waiting for the arrival of an unproductive agent before he reports his arrival to get a higher level of expected utility. Yet, the waiting time until the next arrival of an unproductive agent discounts future payoffs with an expected discount factor of at least $\Delta^L(p,\delta)$ thereby mitigating the advantage of waiting. Hence, also in $(0,0)$ and $(1,0)$ it is unprofitable for the productive type to delay his arrival report.[34] This concludes the proof of Step (ii).

Jointly, (i) and (ii) imply that the Positive Assortative Policy together with our payments is incentive compatible even when arrivals are unobservable. Individual rationality, no deficit, and efficient exit remain satisfied, completing the construction of the mechanism. ∎

Impatient Policy. Implementing the Impatient Policy in a market where the designer can observe arrivals turned out to be straightforward. As the designer may ignore agents' private information to implement the Impatient Policy, he can abstain from using payments. Yet, if the designer cannot observe agents' arrivals, information relevant for implementing the welfare-maximizing policies, implementation of the Impatient Policy becomes more difficult.

In contrast to the Positive Assortative Policy, the agent's reported type does not fix his match partner's type in the Impatient Policy. The agent's partner is, depending on the market report, either the only agent that is present in the market or the next agent that arrives to the market, irrespective of his type. If the designer asks an agent for his type, future agents may condition their reporting strategy on that report. Consider, for example, a productive agent that arrives to a market with one agent that has reported an unproductive type. If the productive agent reveals his arrival immediately, he will form a group with the unproductive agent. If the productive agent delays his arrival report until after the next arrival, he has the opportunity to be matched with a productive agent. Therefore, depending on the parameter constellation, it might be profitable for the productive agent to delay his arrival report.

The designer can circumvent this problem by separating the agents' arrival report from their type report. To implement the Impatient Policy, the designer only needs agents' arrival times but not their types. If the designer asks agents only for their arrival time, future agents only observe arrival reports. Given that agents only observe arrival reports, it is optimal for the agents to report their arrival as early as possible, i.e.,

[34] To avoid issues with large states that occur because several agents report their arrival simultaneously, we punish agents reporting the same arrival time with a sufficiently high payment, say, m_{HH}. In equilibrium this entails no welfare loss. The deviations checked are, thus, an upper bound for the most profitable deviation.

truthfully. Therefore, anticipating the agents' informational advantage from reported types, the designer strategically chooses not to ask the agents for their type in order to implement the Impatient Policy.

Combining the insights of the last two sections, we find that even if the designer does not observe arrivals to the market, the welfare-maximizing policies are implementable.

2.4.5 Concluding Remarks

Remark 1. When implementing the Impatient Policy with unobservable arrivals, we demonstrated that it can be beneficial for the designer to strategically not ask agents for their type. Transferring this thought, we can construct another mechanism which implements the Positive Assortative Policy with observable arrivals. As opposed to the Impatient Policy, the Positive Assortative Policy exploits information about agents' types; there exists exactly one situation in which the designer does not need this information upon an agent's arrival: When an agent arrives to an empty market. In this case, the designer needs the agent's information only upon arrival of the next agent, as there is no decision to be taken beforehand. Thus, the designer could set up a mechanism in which an agent that arrives to an empty market does not report his type immediately but only upon arrival of the next agent. The difference to the mechanism studied in Section 2.4.1 is that the second agent arriving to the market does not know the first agent's type. Recall that the most critical situation when implementing the Positive Assortative Policy in Section 2.4.1 arose when the market consisted of one unproductive agent and for small values of p. Under the new mechanism, the subtle difference is that in this situation the second agent does not know that the first agent is unproductive but attaches a high probability to this event.

Remark 2. Observe that throughout Section 2.4 we did not use the assumption of Section 2.3 that the value of the productive pair is not too large compared to the value of the unproductive pair. Hence, our implementation results carry over to the case $m_{HH} > 3m_{LL}$ whenever the Positive Assortative Policy and the Impatient Policy are welfare-maximizing.

2.5 Conclusion

This chapter studies a dynamic matching market organized by a central authority. Agents of different types that arrive to the market according to a discrete process are matched by a social planner. The model is flexible with respect to four key features: The degree of complementarity of the partners' characteristics in the match value function, the relative size of absolute values of output generated by the two possible homogeneous matchings, the probability distribution of arriving agents' types and the patience represented by discounting. In the first part of this chapter, we address the optimal matching policies under complete information. We develop a tool that helps us to solve for the optimal matching policy in closed form without imposing any restriction on the policy. Whenever the agents' productivities do not differ too much, one of three policies is optimal: The Positive Assortative Policy, the Provident Impatient Policy, or the Myopic Impatient Policy. The social planner is more willing to abstain from creating mixed matches in order to wait for positive assortative matchings when discounting is weak or complementarities are strong. This has two immediate implications: a) the optimality of positive assortative matchings in static matching is robust to small discounting frictions, and b) when mixed matches are created because of impatience in models of search and matching, this might be welfare-enhancing. The role of the distribution of arriving agents' types is more sophisticated: The designer might abstain from mixed matches only for intermediate probabilities of productive arrivals. When the match value of two productive agents exceeds by far the match value of the unproductive counterpart, it is sometimes optimal to stock unproductive agents in the market in order to ensure that arriving productive agents can get paired immediately. In the second part of this chapter, we consider implementability of the optimal policy in the presence of private information. We prove implementability of the optimal matching policy when agents have private information about their types and can hide their arrival to the market. We show that if the complementarity of the match value function is sufficiently strong or the environment is sufficiently patient, the welfare-maximizing policy can be implemented with payments that are reminiscent of those that implement the welfare-maximizing policy in the static model. Finally, we identify situations in which the market organizer can abstain from using monetary incentives.

The simple structure of our model helps to expose the trade-off between accumulating agents to achieve positive assortative matchings and matching agents early in order to avoid waiting costs. We conjecture that the State Space Reduction developed in this chapter can also be employed to find optimal policies in the extended model with an arbitrary but finite number of types. While we focus on a supermodular match value, we conjecture that in the submodular case it is optimal to form exclusively mixed pairs as the time friction vanishes. Therefore, policies which store homogeneous groups of both productive and unproductive agents might be optimal. Analyzing a model with a continuum of types, but discrete arrivals, would allow for a more detailed comparison between the centralized matching market and decentralized search and matching models. This is an interesting avenue for future research.

References

Ahlin, Christian (2015): "Matching patterns when group size exceeds two." *mimeo*. [9]

Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan (2016): "Thickness and information in dynamic matching markets." *mimeo*. [13]

Anderson, Axel and Lones Smith (2010): "Dynamic matching and evolving reputations." *The Review of Economic Studies*, 77 (1), 3–29. [13]

Anderson, Ross, Itai Ashlagi, David Gamarnik, and Yash Kanoria (2015): "Efficient dynamic barter exchange." *mimeo*. [13]

Arrow, Kenneth J, Theodore Harris, and Jacob Marschak (1951): “Optimal inventory policy.” *Econometrica*, 19 (3), 250–272. [11]

Ashlagi, Itai, Maximillien Burq, Patrick Jaillet, and Vahideh H Manshadi (2016): “On matching and thickness in heterogeneous dynamic markets.” *mimeo*. [13]

Ashlagi, Itai, Patrick Jaillet, and Vahideh H Manshadi (2013): “Kidney exchange in dynamic sparse heterogeneous pools.” *mimeo*. [13]

Atakan, Alp E (2006): “Assortative matching with explicit search costs.” *Econometrica*, 74 (3), 667–680. [9]

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv (2015): “Optimal dynamic matching.” *mimeo*. [13]

Bergemann, Dirk and Juuso Välimäki (2010): “The dynamic pivot mechanism.” *Econometrica*, 78 (2), 771–789. [12, 14, 34, 36]

Bloch, Francis and David Cantala (2013): “Markovian assignment rules.” *Social Choice and Welfare*, 40 (1), 1–25. [14]

Bloch, Francis and David Cantala (2016): “Dynamic assignment of objects to queuing agents.” *mimeo*. [14]

Bloch, Francis and Nicolas Houy (2012): “Optimal assignment of durable objects to successive agents.” *Economic Theory*, 51 (1), 13–33. [14]

Board, Simon and Andy Skrzypacz (2015): “Revenue management with forward-looking buyers.” *Journal of Political Economy, forthcoming*. [13]

Burdett, Ken and Melvyn G Coles (1997): “Marriage and class.” *The Quarterly Journal of Economics*, 112 (1), 141–168. [9]

Cremer, Jacques and Richard P McLean (1985): “Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent.” *Econometrica*, 53 (2), 345–361. [14]

Cremer, Jacques and Richard P McLean (1988): “Full extraction of the surplus in Bayesian and dominant strategy auctions.” *Econometrica*, 56 (6), 1247–1257. [14]

Dizdar, Deniz, Alex Gershkov, and Benny Moldovanu (2011): “Revenue maximization in the dynamic knapsack problem.” *Theoretical Economics*, 6 (2), 157–184. [13]

Dizdar, Deniz and Benny Moldovanu (2016): “On the importance of uniform sharing rules for efficient matching.” *Journal of Economic Theory*, 165, 106–123. [33, 35]

Fershtman, Daniel and Alessandro Pavan (2015): “Re-matching, experimentation, and cross-subsidization.” *mimeo*. [13]

Gallien, Jérémie (2006): “Dynamic mechanism design for online commerce.” *Operations Research*, 54 (2), 291–310. [13]

Gershkov, Alex and Benny Moldovanu (2009): “Dynamic revenue maximization with heterogeneous objects: A mechanism design approach.” *American Economic Journal: Microeconomics*, 1 (2), 168–198. [13]

Gershkov, Alex and Benny Moldovanu (2010): “Efficient sequential assignment with incomplete information.” *Games and Economic Behavior*, 68 (1), 144–154. [13, 34]

Ghatak, Maitreesh (1999): “Group lending, local information and peer selection.” *Journal of Development Economics*, 60 (1), 27–50. [8]

Ghatak, Maitreesh (2000): “Screening by the company you keep: Joint liability lending and the peer selection effect.” *The Economic Journal*, 110 (465), 601–631. [8]

Ghatak, Maitreesh and Timothy W Guinnane (1999): “The economics of lending with joint liability: theory and practice.” *Journal of Development Economics*, 60 (1), 195–228. [8]

Herbst, Holger and Benjamin Schickner (2016): “Dynamic formation of teams: When does waiting for good matches pay off?” *mimeo*. [7]

Jehiel, Philippe and Benny Moldovanu (2001): "Efficient design with interdependent valuations." *Econometrica*, 1237–1259. [36]

Kremer, Michael (1993): "The O-ring theory of economic development." *The Quarterly Journal of Economics*, 108 (3), 551–575. [8]

Kurino, Morimitsu (2014): "House allocation with overlapping agents: A dynamic mechanism design approach." *American Economic Journal: Microeconomics*, 6 (1), 258–289. [14]

Legros, Patrick and Andrew F Newman (2002): "Monotone matching in perfect and imperfect worlds." *The Review of Economic Studies*, 69 (4), 925–942. [7]

Leshno, Jacob D (2015): "Dynamic matching in overloaded waiting lists." *mimeo*. [14]

Liu, Heng (2014): "Efficient dynamic mechanisms in environments with interdependent valuations." *mimeo*. [14, 36]

Lu, Xiaohua and R Preston McAfee (1996): "Matching and expectations in a market with heterogeneous agents." *Advances in applied microeconomics*, 6, 121–156. [9]

Mailath, George J, Andrew Postlewaite, and Larry Samuelson (2013): "Pricing and investments in matching markets." *Theoretical Economics*, 8 (2), 535–590. [15]

Mailath, George J, Andrew Postlewaite, and Larry Samuelson (2015): "Premuneration values and investments in matching markets." *Economic Journal, forthcoming*. [15, 39]

Mezzetti, Claudio (2004): "Mechanism design with interdependent valuations: Efficiency." *Econometrica*, 72 (5), 1617–1626. [14]

Mierendorff, Konrad (2015): "Optimal dynamic mechanism design with deadlines." *Journal of Economic Theory, forthcoming*. [13]

Myerson, Roger B (1986): "Mulitstage games with communication." *Econometrica*, 54 (2), 323–358. [35]

Nath, Swaprava, Onno Zoeter, Y Narahari, and Christopher R Dance (2015): "Dynamic mechanism design with interdependent valuations." *Review of Economic Design*, 19 (3), 211–228. [14, 36]

Noda, Shunya (2016): "Full surplus extraction and within-period ex post implementation in dynamic environments." *mimeo*. [14]

Pai, Mallesh M and Rakesh Vohra (2013): "Optimal dynamic auctions and simple index rules." *Mathematics of Operations Research*, 38 (4), 682–697. [13]

Sattinger, Michael (1995): "Search and the efficient assignment of workers to jobs." *International Economic Review*, 36 (2), 283–302. [9]

Shi, Shouyong (2005): "Frictional assignment, Part II: Infinite horizon and inequality." *Review of Economic Dynamics*, 8 (1), 106–137. [13]

Shimer, Robert and Lones Smith (2000): "Assortative matching and search." *Econometrica*, 68 (2), 343–369. [8, 10]

Shimer, Robert and Lones Smith (2001a): "Matching, search, and heterogeneity." *Advances in Macroeconomics*, 1 (1). [9]

Shimer, Robert and Lones Smith (2001b): "Nonstationary search." *mimeo*. [9]

Smith, Lones (2006): "The marriage model with search frictions." *Journal of Political Economy*, 114 (6), 1124–1144. [8, 10]

Smith, Lones (2011): "Frictional matching models." *Annual Review of Economics*, 3 (1), 319–338. [9]

Ünver, M Utku (2010): "Dynamic kidney exchange." *The Review of Economic Studies*, 77 (1), 372–414. [13]

Veinott Jr, Arthur F (1966): "The status of mathematical inventory theory." *Management Science*, 12 (11), 745–777. [11]

Whitin, Thomson M (1954): "Inventory control research: A survey." *Management Science*, 1 (1), 32–40. [11]

Appendix 2.A Appendix: Proofs

The proof of Theorem 1 uses Lemma 1 to 6, which are, hence, proven first.

Preliminaries for Lemma 1 to 3. We denote the candidate policy by ρ, fix an arbitrary state $(x, y) \in S$, and denote by $d = (d_{HH}, d_{HL}, d_{LL})$ a one-period deviation that matches on (x, y) d_{HH} homogeneous pairs of productive agents, d_{HL} mixed pairs, and d_{LL} homogeneous pairs of unproductive agents. The value of ρ on (x, y) can be written as

$$\begin{aligned} V_\rho(x, y) \;=\; & \rho^{HH}(x, y)m_{HH} + \rho^{HL}(x, y)m_{HL} + \rho^{LL}(x, y)m_{LL} \\ & + \delta[pV_\rho(x' + 1, y') + (1 - p)V_\rho(x', y' + 1)] \end{aligned} \tag{26}$$

with

$$x' \;=\; x - 2\rho^{HH}(x, y) - \rho^{HL}(x, y), \quad y' \;=\; y - 2\rho^{LL}(x, y) - \rho^{HL}(x, y).$$

Similarly, the value of deviation d from ρ on (x, y) can be expressed as

$$\begin{aligned} V_\rho^d(x, y) \;=\; & d_{HH}m_{HH} + d_{HL}m_{HL} + d_{LL}m_{LL} \\ & + \delta[pV_\rho(x'' + 1, y'') + (1 - p)V_\rho(x'', y'' + 1)] \end{aligned} \tag{27}$$

with

$$x'' \;=\; x - 2d_{HH}(x, y) - d_{HL}(x, y), \quad y'' \;=\; y - 2d_{LL}(x, y) - d_{HL}(x, y).$$

In each of the Lemmas we will argue that the value of ρ exceeds the value of a certain class of deviations.

Proof of Lemma 1. Assume that $\rho^{HH}(x, y) > 0$ and $d_{HH} > 0$. Subtracting m_{HH} from $V_\rho(x, y)$ and $V_\rho^d(x, y)$, we observe that the deviation is unprofitable if and only if

$$\begin{aligned} & (\rho^{HH}(x, y) - 1)m_{HH} + \rho^{HL}(x, y)m_{HL} + \rho^{LL}(x, y)m_{LL} + V_\rho(x', y') \\ & \geq \; (d_{HH} - 1)m_{HH} + d_{HL}m_{HL} + d_{LL}m_{LL} + \delta[pV_\rho(x'' + 1, y'') + (1 - p)V_\rho(x'', y'' + 1)]. \end{aligned} \tag{28}$$

By consistency the left hand side of the above inequality coincides with $V_\rho(x - 2, y)$. As the deviation $d_{HH} - 1$, d_{HL}, and d_{LL} is feasible on $(x - 2, y)$, (28) describes a deviation from ρ on $(x - 2, y)$. Given that no deviation is profitable on smaller states, the inequality holds. Thus, deviation (d_{HH}, d_{HL}, d_{LL}) is not profitable on (x, y) either. The proof is analogous for the cases $\rho^{HL}(x, y), d_{HL} > 0$ and $\rho^{LL}(x, y), d_{LL} > 0$.■

Proof of Lemma 2. Consider a deviation (d_{HH}, d_{HL}, d_{LL}) such that $\rho^{HH}(x'' + 1, y'') > 0$ and $\rho^{HH}(x'', y'' + 1) > 0$. The proof constructs an auxiliary deviation $d' =$

$(d'_{HH}, d'_{HL}, d'_{LL})$ on (x, y) with $V_\rho^{d'}(x, y) \geq V_\rho^d(x, y)$. Hence, if $(d'_{HH}, d'_{HL}, d'_{LL})$ is not profitable, then (d_{HH}, d_{HL}, d_{LL}) is not profitable either.

Set $(d'_{HH}, d'_{HL}, d'_{LL}) = (d_{HH} + 1, d_{HL}, d_{LL})$. By our choice of (d_{HH}, d_{HL}, d_{LL}), $(d'_{HH}, d'_{HL}, d'_{LL})$ is feasible. From consistency of ρ, $\rho^{HH}(x'' + 1, y'') > 0$, and $\rho^{HH}(x'', y'' + 1) > 0$ follows $V_\rho(x'' + 1, y'') = m_{HH} + V_\rho(x'' - 1, y'')$ and $V_\rho(x'', y'' + 1) = m_{HH} + V_\rho(x'' - 2, y'' + 1)$. Together with $\delta < 1$ this implies for $V_\rho^d(x, y)$:

$$\begin{aligned} & d_{HH}m_{HH} + d_{HL}m_{HL} + d_{LL}m_{LL} + \delta m_{HH} \\ & \qquad + \delta[pV_\rho(x'' - 1, y'') + (1 - p)V_\rho(x'' - 2, y'' + 1)] \\ < \; & (d_{HH} + 1)m_{HH} + d_{HL}m_{HL} + d_{LL}m_{LL} \\ & \qquad + \delta[pV_\rho(x'' - 1, y'') + (1 - p)V_\rho(x'' - 2, y'' + 1)]. \end{aligned}$$

We conclude by observing that the latter term is $V_\rho^{d'}(x, y)$. The two remaining cases $\rho^{HL}(x'' + 1, y'')$, $\rho^{HL}(x'', y'' + 1) > 0$ and $\rho^{LL}(x'' + 1, y'')$, $\rho^{LL}(x'', y'' + 1) > 0$ follow from an analogous argument. ■

Proof of Lemma 3. Consider a deviation (d_{HH}, d_{HL}, d_{LL}) with $d_{HL} \geq 2$. As in Lemma 2, we construct an auxiliary deviation $d' = (d'_{HH}, d'_{HL}, d'_{LL})$ with higher value.

Set $(d'_{HH}, d'_{HL}, d'_{LL}) = (d_{HH} + 1, d_{HL} - 2, d_{LL} + 1)$. $(d'_{HH}, d'_{HL}, d'_{LL})$ is feasible because (d_{HH}, d_{HL}, d_{LL}) is feasible. As next period states are identical under both deviations, $V_\rho^{d'}(x, y) - V_\rho^d(x, y) = m_{HH} + m_{LL} - 2m_{HL} \geq 0$, where the inequality follows from the supermodularity of the match value function. Thus, $V_\rho^{d'}(x, y) \geq V_\rho^d(x, y)$. ■

Proof of Lemma 4. By construction, the Positive Assortative Policy ρ_{PAP} is consistent. Fix a state (x, y) with $x \geq 3$ and consider a deviation (d_{HH}, d_{HL}, d_{LL}) on (x, y). As $\rho_{PAP}^{HH}(x, y) > 0$, only deviations with $d_{HH} = 0$ have to be verified by Lemma 1. Subsequent to following (d_{HH}, d_{HL}, d_{LL}) there are at least $x - 2d_{HH} - d_{HL}$ productive agents in the market after the next arrival. As $\rho_{PAP}^{HH}(x', y') > 0$, $\forall x' \geq 2$, only deviations with $x - 2d_{HH} - d_{HL} < 2$ have to be checked by Lemma 2. By Lemma 3, only deviations with $d_{HL} < 2$ have to be checked. As $x \geq 3$, the set of deviations satisfying the latter three conditions is empty, i.e., no deviation on (x, y) with $x \geq 3$ has to be checked. Similarly, no deviation on (x, y) with $y \geq 3$ has to be checked. ■

Proof of Lemma 5. As all arguments used for the Positive Assortative Policy ρ_{PAP} also apply to the Provident Impatient Policy ρ_{PIP}, the proof is exactly the same as the proof of Lemma 4. ■

Proof of Lemma 6. By construction, the Myopic Impatient Policy ρ_{MIP} is consistent. The proof for states (x, y) with $x \geq 3$ parallels the proof of Lemma 4. For states with many unproductive agents, however, we need to alter the argument slightly. When applying Lemma 2, we can only exclude deviations with $y - 2d_{LL} - d_{HL} \geq 3$ because $\rho^{LL}_{MIP}(x', y') > 0$ holds only $\forall y' \geq 3$. We can apply Lemma 1 and Lemma 3 as before. Thus, no deviation has to checked on states (x, y) with $y \geq 4$. ■

Proof of Theorem 1. For each candidate policy, Lemmas 4 to 6 identify the set of states on which every possible deviation has to be verified for its unprofitability by hand. As is shown in the following, for all parameter constellations $(p, \delta, m_{HH}, m_{LL}, m_{HL})$ such that $m_{HL} \notin \{m^1_{HL}, m^2_{HL}\}$, deviations from the respective candidate policy give a strictly lower payoff. This implies uniqueness of the optimal policy.

Claim 1: The Positive Assortative Policy is optimal for all parameter constellations $(p, \delta, m_{HH}, m_{LL}, m_{HL})$ such that $m_{HL} \leq m^1_{HL}$.

The value function V_{PAP} at states in the recurrent set is determined by the following equations:

$$
\begin{aligned}
V_{PAP}(1,0) &= \delta[pV_{PAP}(2,0) + (1-p)V_{PAP}(1,1)], && (29)\\
V_{PAP}(0,1) &= \delta[pV_{PAP}(1,1) + (1-p)V_{PAP}(0,2)], && (30)\\
V_{PAP}(1,1) &= \delta[pV_{PAP}(2,1) + (1-p)V_{PAP}(1,2)], && (31)\\
V_{PAP}(2,0) &= m_{HH} + \delta[pV_{PAP}(1,0) + (1-p)V_{PAP}(0,1)], && (32)\\
V_{PAP}(0,2) &= m_{LL} + \delta[pV_{PAP}(1,0) + (1-p)V_{PAP}(0,1)], && (33)\\
V_{PAP}(2,1) &= m_{HH} + \delta[pV_{PAP}(1,1) + (1-p)V_{PAP}(0,2)], && (34)\\
V_{PAP}(1,2) &= m_{LL} + \delta[pV_{PAP}(2,0) + (1-p)V_{PAP}(1,1)]. && (35)
\end{aligned}
$$

Define $V_{PAP}(0,0) := \delta[pV_{PAP}(1,0) + (1-p)V_{PAP}(0,1)]$. The value at all remaining states (x, y) is given by

$$
\begin{aligned}
V_{PAP}(x,y) &= \rho^{HH}_{PAP}(x,y) + \rho^{LL}_{PAP}(x,y) + \rho^{HL}_{PAP}(x,y) + V_{PAP}(x',y')\\
\text{with} \quad x' &= x - 2\rho^{HH}_{PAP}(x,y) - \rho^{HL}_{PAP}(x,y) < 2\\
\text{and} \quad y' &= y - 2\rho^{LL}_{PAP}(x,y) - \rho^{HL}_{PAP}(x,y) < 2.
\end{aligned}
$$

Observe that this determines $V_{PAP}(x, y)$ uniquely on the entire state space. By Lemma 4, the set of states on which the unprofitability of deviations has to be verified is $\{(x, y) | x \leq 2, y \leq 2, x + y \geq 2)\}$.[35]

[35] Note that on states $(1, 0)$ and $(0, 1)$ there is no possible deviation and hence no profitable deviation.

On states $(2,0)$ and $(0,2)$ the only possible deviation is $(d_{HH}, d_{HL}, d_{LL}) = (0,0,0)$. In both states this deviation does not have to be considered by Lemma 2.
On state $(1,2)$ there are two possible deviations: (0, 0, 0) and (0, 1, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 1, 0) is not profitable if

$$V_{PAP}(1,2) \; = \; m_{LL} + V_{PAP}(1,0) \; \geq \; m_{HL} + V_{PAP}(0,1). \tag{36}$$

On state $(2,1)$ there are two possible deviations: (0, 0, 0) and (0, 1, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 1, 0) is not profitable if

$$V_{PAP}(2,1) \; = \; m_{HH} + V_{PAP}(0,1) \; \geq \; m_{HL} + V_{PAP}(1,0). \tag{37}$$

On state $(2,2)$ there are five possible deviations: (0, 0, 0), (0, 1, 0), (1, 0, 0), (0, 0, 1) and (0, 2, 0). Deviations (0, 0, 0), (1, 0, 0) and (0, 0, 1) do not have to be considered by Lemma 2. Deviation (0, 2, 0) does not have to be considered by Lemma 3. Deviation (0, 1, 0) is not profitable if

$$V_{PAP}(2,2) \; = \; m_{HH} + m_{LL} + V_{PAP}(0,0) \; \geq \; m_{HL} + V_{PAP}(1,1). \tag{38}$$

On state $(1,1)$ there is one possible deviation, which is (0, 1, 0). The condition for deviation (0, 1, 0) to be unprofitable is

$$V_{PAP}(1,1) \; \geq \; m_{HL} + V_{PAP}(0,0). \tag{39}$$

The final step is to observe that inequalities (36) to (39) hold if and only if $m_{HL} \leq m^1_{HL}$, in particular, (39) corresponds exactly to $m_{HL} \leq m^1_{HL}$. For computational details see Appendix 2.B.

Claim 2: The Provident Impatient Policy is optimal for all parameter constellations $(p, \delta, m_{HH}, m_{LL}, m_{HL})$ such that $m^1_{HL} \leq m_{HL} \leq m^2_{HL}$.

The value function V_{PIP} at states in the recurrent set is given by the following equations:

$$\begin{aligned}
V_{PIP}(1,0) \; &= \delta[pV_{PIP}(2,0) + (1-p)V_{PIP}(1,1)], && (40)\\
V_{PIP}(0,1) \; &= \delta[pV_{PIP}(1,1) + (1-p)V_{PIP}(0,2)], && (41)\\
V_{PIP}(1,1) \; &= m_{HL} + \delta[pV_{PIP}(1,0) + (1-p)V_{PIP}(0,1)], && (42)\\
V_{PIP}(2,0) \; &= m_{HH} + \delta[pV_{PIP}(1,0) + (1-p)V_{PIP}(0,1)], && (43)\\
V_{PIP}(0,2) \; &= m_{LL} + \delta[pV_{PIP}(1,0) + (1-p)V_{PIP}(0,1)]. && (44)
\end{aligned}$$

Define $V_{PIP}(0,0) := \delta[pV_{PIP}(1,0) + (1-p)V_{PIP}(0,1)]$. The value at all remaining states (x,y) is determined as follows:

$$\begin{aligned} V_{PIP}(x,y) &= \rho^{HH}_{PIP}(x,y) + \rho^{LL}_{PIP}(x,y) + \rho^{HL}_{PIP}(x,y) + V_{PIP}(x',y') \\ \text{with} \quad x' &= x - 2\rho^{HH}_{PIP}(x,y) - \rho^{HL}_{PIP}(x,y) < 2 \\ \text{and} \quad y' &= y - 2\rho^{LL}_{PIP}(x,y) - \rho^{HL}_{PIP}(x,y) < 2. \end{aligned}$$

Solving the system gives

$$V_{PIP}(0,0) = \frac{\delta^2}{1-\delta^2}[(p^2 + (1-p)^2)m_{HL} + (1-p)p(m_{HH} + m_{LL})], \tag{45}$$

$$V_{PIP}(1,0) = \delta V_{PIP}(0,0) + \delta(pm_{HH} + (1-p)m_{HL}), \tag{46}$$

$$V_{PIP}(0,1) = \delta V_{PIP}(0,0) + \delta(pm_{HL} + (1-p)m_{LL}), \tag{47}$$

$$V_{PIP}(1,1) = m_{HL} + \delta V_{PIP}(0,0), \tag{48}$$

$$V_{PIP}(2,0) = m_{HH} + \delta V_{PIP}(0,0), \tag{49}$$

$$V_{PIP}(0,2) = m_{LL} + \delta V_{PIP}(0,0). \tag{50}$$

By Lemma 5, the set of states on which the unprofitability of deviations has to be verified is $\{(x,y)|x \leq 2, y \leq 2, x+y \geq 2)\}$.

On states $(2,0)$ and $(0,2)$ the only possible deviation is $(d_{HH}, d_{HL}, d_{LL}) = (0,0,0)$. In both states this deviation does not have to be considered by Lemma 2.

On state $(2,1)$ there are two possible deviations: (0, 0, 0) and (0, 1, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. The inequality corresponding to deviation (0, 1, 0) is

$$V_{PIP}(2,1) = m_{HH} + V_{PIP}(0,1) \geq m_{HL} + V_{PIP}(1,0). \tag{51}$$

On state $(2,2)$ there are five possible deviations: (0, 0, 0), (0, 1, 0), (1, 0, 0), (0, 0, 1) and (0, 2, 0). Deviations (0, 0, 0), (1, 0, 0) and (0, 0, 1) do not have to be considered by Lemma 2. Deviation (0, 2, 0) does not have to be considered by Lemma 3. Deviation (0, 1, 0) is not profitable if

$$m_{HH} + m_{LL} + V_{PIP}(0,0) \geq m_{HL} + \delta[pV_{PIP}(2,1) + (1-p)V_{PIP}(1,2)]. \tag{52}$$

On state $(1,1)$ there is one possible deviation, which is (0, 0, 0). The condition for deviation (0, 0, 0) to be not profitable is

$$V_{PIP}(1,1) = m_{HL} + V_{PIP}(0,0) \geq \delta[pV_{PIP}(2,1) + (1-p)V_{PIP}(1,2)]. \tag{53}$$

On state $(1,2)$ there are two possible deviations: (0, 0, 0) and (0, 1, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 1, 0) is not profitable if

$$V_{PIP}(1,2) \;=\; m_{LL} + V_{PIP}(1,0) \;\geq\; m_{HL} + V_{PIP}(0,1). \tag{54}$$

The final step is to show that inequalities (51) to (54) hold if and only if $m^1_{HL} \leq m_{HL} \leq m^2_{HL}$. Inserting the explicit solution for V_{PIP} into (53) and solving for m_{HL} shows that (53) corresponds to $m_{HL} \geq m^1_{HL}$. Similarly, (54) yields $m_{HL} \leq m^2_{HL}$. See Appendix 2.B for calculatory details.

Claim 3: The Myopic Impatient Policy is optimal for all parameter constellations $(p, \delta, m_{HH}, m_{LL}, m_{HL})$ such that $m_{HL} \geq m^2_{HL}$.

The Myopic Impatient Policy has the same recurrent set as the Provident Impatient Policy, $R := \{(x,y)|1 \leq x + y \leq 2\}$. Define $V_{MIP}(0,0) := \delta[pV_{MIP}(1,0) + (1-p)V_{MIP}(0,1)]$. By construction, $\rho_{MIP}(x,y) = \rho_{PIP}(x,y)$, $\forall(x,y) \in R$, hence, $V_{MIP}(x,y) = V_{PIP}(x,y)$, $\forall(x,y) \in R$ with explicit solution (45) to (50). At all remaining states (x,y) the value is given by

$$\begin{aligned}
V_{MIP}(x,y) \; &= \rho^{HH}_{MIP}(x,y) + \rho^{LL}_{MIP}(x,y) + \rho^{HL}_{MIP}(x,y) + V_{MIP}(x',y') \\
\text{with} \quad x' \; &= \; x - 2\rho^{HH}_{MIP}(x,y) - \rho^{HL}_{MIP}(x,y) \; < \; 2 \\
\text{and} \quad y' \; &= \; y - 2\rho^{LL}_{MIP}(x,y) - \rho^{HL}_{MIP}(x,y) \; < \; 2.
\end{aligned}$$

By Lemma 6, the set of states on which the unprofitability of deviations has to be verified is $\{(x,y)|x \leq 2, y \leq 3, x + y \geq 2)\}$.

On state $(2,0)$ the only possible deviation is $(d_{HH}, d_{HL}, d_{LL}) = (0,0,0)$. This deviation does not have to be considered by Lemma 2.

On state $(2,1)$ there are two possible deviations: (0, 0, 0) and (0, 1, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 1, 0) is not profitable either: The corresponding inequality is

$$V_{MIP}(2,1) \;=\; m_{HH} + V_{MIP}(0,1) \;\geq\; m_{HL} + V_{MIP}(1,0). \tag{55}$$

As $V_{MIP}(0,1) = V_{PIP}(0,1)$ and $V_{MIP}(1,0) = V_{PIP}(1,0)$, (55) equals (51) which holds.

On state $(1,1)$ the only possible deviation is (0, 0, 0) which is unprofitable if

$$V_{MIP}(1,1) \;=\; m_{HL} + V_{MIP}(0,0) \;\geq\; \delta[pV_{MIP}(2,1) + (1-p)V_{MIP}(1,2)]. \tag{56}$$

On state $(2,2)$ there are five possible deviations: (0, 0, 0), (0, 1, 0), (1, 0, 0), (0, 0, 1) and (0, 2, 0). Deviations (0, 0, 0), and (1, 0, 0) do not have to be considered by Lemma 2. Deviation (0, 2, 0) does not have to be considered by Lemma 3. Deviation (0, 0, 1) does not have to be considered by Lemma 1. Deviation (0, 1, 0) is not profitable if

$$m_{HH} + m_{LL} + V_{MIP}(0,0) \;\geq\; m_{HL} + \delta[pV_{MIP}(2,1) + (1-p)V_{MIP}(1,2)]. \tag{57}$$

On state $(0,3)$ the only possible deviation is (0, 0, 0) which does not have to be considered by Lemma 2.

On state $(1,3)$ there are three possible deviations: (0, 0, 0), (0, 0, 1) and (0, 1, 0). Deviations (0, 0, 1) and (0, 1, 0) do not have to be considered by Lemma 1. Deviation (0, 0, 0) does not have to be considered by Lemma 2.

On state $(2,3)$ there are five possible deviations: (0, 0, 0), (0, 1, 0), (1, 0, 0), (0, 0, 1) and (0, 2, 0). Deviations (0, 0, 0), and (1, 0, 0) do not have to be considered by Lemma 2. Deviation (0, 2, 0) does not have to be considered by Lemma 3. Deviation (0, 0, 1) does not have to be considered by Lemma 1. Deviation (0, 1, 0) is not profitable if

$$m_{HH} + m_{LL} + V_{MIP}(0,1) \;\geq\; m_{HL} + \delta[pV_{MIP}(2,2) + (1-p)V_{MIP}(1,3)]. \tag{58}$$

On state $(0,2)$ the only possible deviation is (0, 0, 0) which is unprofitable if

$$V_{MIP}(0,2) \;=\; m_{LL} + V_{MIP}(0,0) \;\geq\; \delta[pV_{MIP}(1,2) + (1-p)V_{MIP}(0,3)]. \tag{59}$$

On state $(1,2)$ there are two possible deviations: (0, 0, 0) and (0, 0, 1). Deviation (0, 0, 1) is not profitable if

$$m_{HL} + V_{MIP}(0,1) \;\geq\; m_{LL} + V_{MIP}(1,0), \tag{60}$$

and deviation (0, 0, 0) is not profitable if

$$m_{HL} + V_{MIP}(0,1) \;\geq\; \delta[pV_{MIP}(2,2) + (1-p)V_{MIP}(1,3)]. \tag{61}$$

The final step is to show that inequalities (56) to (61) hold if and only if $m_{HL} \geq m^2_{HL}$. Using the explicit solution for V_{MIP} and solving for m_{HL}, we find that (60) is equivalent to $m_{HL} \geq m^2_{HL}$ and that (59) corresponds to $m_{HL} \leq m^3_{HL}$. Supporting calculations can be found in Appendix 2.B.

To conclude the proof, observe that $m_{HL} \leq m^3_{HL}$ for all $p, \delta, m_{HH}, m_{LL}, m_{HL}$ such that $m_{HH} \leq 3m_{LL}$. Thus, the parameter regions on which the three candidates are optimal span the entire parameter space. ■

Proof of Corollary 1. The existence of two cut-off levels follows from showing that $\frac{\partial m^1_{HL}}{\partial \delta} \geq 0$ and $\frac{\partial m^2_{HL}}{\partial \delta} \geq 0$, independent of the specific choice of parameters. Using the definitions of m^1_{HL} and m^2_{HL} from (5) and (6), we obtain

$$\frac{\partial m^1_{HL}}{\partial \delta} = m_{HH}\frac{p}{[1-\delta(1-2p)]^2} + m_{LL}\frac{1-p}{[1+\delta(1-2p)]^2} > 0$$

and

$$\begin{aligned}\frac{\partial m^2_{HL}}{\partial \delta} &= m_{HH}\frac{p}{[1-\delta(1-2p)]^2} + m_{LL}\frac{-p}{[1-\delta(1-2p)]^2} \\ &> m_{HH}\frac{p}{[1-\delta(1-2p)]^2} + m_{HH}\frac{-p}{[1-\delta(1-2p)]^2} = 0.\end{aligned}$$

Furthermore, $m^1_{HL} < m^2_{HL}$ implies $\delta^1 > \delta^2$. Finally, $m^1_{HL} = m^2_{HL} = 0$ for $\delta = 0$ and $m^1_{HL} = m^2_{HL} = 1/2(m_{HH} + m_{LL})$ for $\delta = 1$ and imply $\delta^1, \delta^2 \in [0,1]$. ■

Proof of Proposition 1. We follow the same steps as for the other candidate policies above. Denote Matching Policy $\mathscr{P}_1$ by ρ_{P1}.

Claim 1: To verify candidate policy ρ_{P1} it is sufficient to verify deviations on $\{(x,y)|x \le 2, y \le 4, x+y \ge 2)\}$.

By construction, ρ_{P1} is consistent. For states (x,y) with $x \ge 3$ the argument is the same as in Lemma 4. Analogously to Lemma 6, we need to adjust the proof slightly for states with many unproductive agents when applying Lemma 2. In this case, we can only exclude deviations on states (x,y) with $y \ge 5$.

Claim 2: There exists a parameter region on which there is no profitable deviation from ρ_{P1} on $\{(x,y)|x \le 2, y \le 4, x+y \ge 2)\}$.

Define $V_{P1}(0,0) := \delta[pV_{P1}(1,0) + (1-p)V_{P1}(0,1)]$. The value function V_{P1} at states in the recurrent set is determined by

$$V_{P1}(1,0) = \delta[pV_{P1}(2,0) + (1-p)V_{P1}(1,1)], \tag{62}$$

$$V_{P1}(0,1) = \delta[pV_{P1}(1,1) + (1-p)V_{P1}(0,2)] \tag{63}$$

$$V_{P1}(1,1) = m_{HL} + \delta[pV_{P1}(1,0) + (1-p)V_{P1}(0,1)], \tag{64}$$

$$V_{P1}(2,0) = m_{HH} + \delta[pV_{P1}(1,0) + (1-p)V_{P1}(0,1)], \tag{65}$$

$$V_{P1}(0,2) = \delta[pV_{P1}(1,2) + (1-p)V_{P1}(0,3)], \tag{66}$$

$$V_{P1}(0,3) = m_{LL} + \delta[pV_{P1}(1,1) + (1-p)V_{P1}(0,2)], \tag{67}$$

$$V_{P1}(1,2) = m_{HL} + \delta[pV_{P1}(1,2) + (1-p)V_{P1}(0,3)]. \tag{68}$$

The value at all remaining states (x, y) is

$$\begin{aligned} V_{P1}(x,y) &= \rho_{P1}^{HH}(x,y) + \rho_{P1}^{LL}(x,y) + \rho_{P1}^{HL}(x,y) + V_{P1}(x',y') \\ \text{with} \quad x' &= x - 2\rho_{P1}^{HH}(x,y) - \rho_{P1}^{HL}(x,y) < 2 \\ \text{and} \quad y' &= y - 2\rho_{P1}^{LL}(x,y) - \rho_{P1}^{HL}(x,y) < 3. \end{aligned}$$

On state $(2,0)$ the only possible deviation is $(d_{HH}, d_{HL}, d_{LL}) = (0,0,0)$ which does not have to be considered by Lemma 2.
On state $(2,1)$ there are two possible deviations: (0, 0, 0) and (0, 1, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 1, 0) is not profitable if

$$V_{P1}(2,1) = m_{HH} + V_{P1}(0,1) \geq m_{HL} + V_{P1}(1,0). \tag{69}$$

On state $(1,1)$ the only possible deviation is (0, 0, 0) which is unprofitable if

$$V_{P1}(1,1) = m_{HL} + V_{P1}(0,0) \geq \delta[pV_{P1}(2,1) + (1-p)V_{P1}(1,2)]. \tag{70}$$

On state $(2,2)$ there are five possible deviations: (0, 0, 0), (0, 1, 0), (1, 0, 1), (0, 0, 1) and (0, 2, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 2, 0) does not have to be considered by Lemma 3. Deviations (0, 0, 1) and (1, 0, 1) do not have to be considered by Lemma 1. Deviation (0, 1, 0) is not profitable if

$$m_{HH} + m_{LL} + V_{P1}(0,0) \geq m_{HL} + \delta[pV_{P1}(2,1) + (1-p)V_{P1}(1,2)]. \tag{71}$$

On state $(0,3)$ the only possible deviation is (0, 0, 0) which is unprofitable if

$$V_{P1}(0,3) = m_{LL} + V_{P1}(0,1) \geq \delta[pV_{P1}(1,3) + (1-p)V_{P1}(0,4)]. \tag{72}$$

On state $(1,3)$ there are three possible deviations: (0, 0, 0), (0, 0, 1) and (0, 1, 1). Deviation (0, 1, 1) does not have to be considered by Lemma 1. Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 0, 1) is not profitable if

$$m_{HL} + V_{P1}(0,2) \geq m_{LL} + \delta[pV_{P1}(2,1) + (1-p)V_{P1}(1,2)]. \tag{73}$$

On state $(0,2)$ the only possible deviation is (0, 0, 1) which is unprofitable if

$$V_{P1}(0,2) = \delta[pV_{P1}(1,2) + (1-p)V_{P1}(0,3)] \geq m_{LL} + V_{P1}(0,0). \tag{74}$$

On state $(1,2)$ there are two possible deviations: (0, 0, 0) and (0, 0, 1). Deviation (0, 0, 1) is not profitable if

$$m_{HL} + V_{P1}(1,0) \geq m_{LL} + V_{P1}(1,0), \tag{75}$$

and deviation (0, 0, 0) is not profitable if

$$m_{HL} + V_{P1}(0,1) \geq \delta[pV_{P1}(2,2) + (1-p)V_{P1}(1,3)]. \tag{76}$$

On state $(2,3)$ there are five possible deviations: (0, 0, 0), (0, 1, 0), (1, 0, 0), (0, 0, 1) and (0, 2, 0). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 2, 0) does not have to be considered by Lemma 3. Deviations (1, 0, 0) and (0, 0, 1) do not have to be considered by Lemma 1. Deviation (0, 1, 0) is not profitable if

$$m_{HH} + m_{LL} + V_{P1}(0,1) \geq m_{HL} + \delta[pV_{P1}(2,2) + (1-p)V_{P1}(1,3)]. \tag{77}$$

On state $(0,4)$ there are two possible deviations: (0, 0, 0) and (0, 0, 2). Deviation (0, 0, 0) does not have to be considered by Lemma 2. Deviation (0, 0, 2) does not have to be considered by Lemma 1.

On state $(1,4)$ there are four possible deviations: (0, 0, 0), (0, 1, 0), (0, 0, 1) and (0, 0, 2). Deviations (0, 0, 0) and (0, 1, 0) do not have to be considered by Lemma 2. Deviations (0, 0, 1) and (0, 0, 2) do not have to be considered by Lemma 1.

On state $(2,4)$ there are nine possible deviations: (0, 0, 0), (0, 1, 0), (1, 0, 0), (0, 1, 1), (0, 0, 2), (0, 2, 1), (1, 0, 2), (0, 0, 1) and (0, 2, 0). Deviations (0, 0, 0) and (0, 1, 0) do not have to be considered by Lemma 2. Deviations (0, 2, 0) and (0, 2, 1) do not have to be considered by Lemma 3. Deviations (1, 0, 0), (1, 0, 2), (0, 0, 2), (0, 1, 1), (0, 2, 1) and (0, 0, 1) do not have to be considered by Lemma 1.

(69) to (77) hold if and only if $m^3_{HL} \leq m_{HL} \leq m^4_{HL}$, where

$$m^4_{HL} = \frac{1}{1-\delta(1-2p)}\left[m_{HH}\delta p + m_{LL}\left(2-\delta(2-p)+\frac{1-\delta+\delta p(1-\delta)^2}{\delta^2 p^2}\right)\right]. \tag{78}$$

(72) coincides with $m_{HL} \leq m^4_{HL}$, and (74) is equivalent to $m_{HL} \geq m^3_{HL}$. All other inequalities are then implied. See Appendix 2.B for supporting calculations.

We conclude the proof by showing that P_1 actually arises. First, we argue that $m^3_{HL} \leq m^4_{HL}$. Note that in the definition of m^4_{HL} and m^3_{HL} the factors in front of m_{HH} coincide. Therefore, $m^3_{HL} \leq m^4_{HL}$ is equivalent to

$$2(1-\delta) + \delta p + \frac{1-\delta}{\delta^2 p^2} + \frac{(1-\delta)^2}{\delta p} \geq 1-\delta+\delta p + \frac{1-\delta}{\delta p} \tag{79}$$

which holds as $\delta, p \in (0,1)$. Second, as argued in the last part of the proof to Theorem 1, if $m_{HH} > 3m_{LL}$, then there exist p, δ, m_{HL} such that $m^3_{HL} < \frac{1}{2}m_{HH} + \frac{1}{2}m_{LL}$. ■

Proof of Corollary 2. The corollary follows from Theorem 1 and the proof of Proposition 1. ■

Proof of Proposition 2. Assume that there exists an optimal policy ρ that never matches two unproductive agents and denote its value function on (x,y) by $V_\rho(x,y)$. We derive a lower bound $\underline{a}$ and an upper bound $\overline{a}$ for $V_\rho(x,y)$ such that $\underline{a} > \overline{a}$, which yields a contradiction.

Lower bound. Observe that by optimality

$$V_\rho(0,k) \geq \left\lfloor \frac{k}{2} \right\rfloor m_{LL} + V_\rho(0, k - 2\left\lfloor \frac{k}{2} \right\rfloor) \geq \left\lfloor \frac{k}{2} \right\rfloor m_{LL} = \underline{a}, \tag{80}$$

where the second inequality holds because $V_\rho(0,0), V_\rho(0,1) \geq 0$.

Upper bound. As ρ never matches two unproductive agents, we can derive an upper bound on the number of matches created in each period when starting in state $(0,k)$ and following policy ρ. When being in state $(0,k)$ in period t, the maximal number of matches in period $t+s$ is bounded from above by s, for any $s \in \mathbb{N}$. Hence, $V_\rho(0,k)$ is bounded from above by the value generated from creating the highest match value as often as possible and as early as possible, i.e., matching two productive agents in every subsequent period which yields

$$V_\rho(0,k) \leq m_{HH} \frac{1}{1-\delta} = \overline{a}.$$

For every parameter constellation $(p, \delta, m_{HH}, m_{LL}, m_{HL})$ there exists a k such that $\underline{a} > \overline{a}$, which is a contradiction. ∎

Proof of Proposition 3. It is sufficient to proof that if $m(\theta_1,\theta_2) = \theta_1 \cdot \theta_2$ then

$$m_{HL} \leq m_{HL}^3 = m_{HH} \frac{\delta p}{1-\delta(1-2p)} + m_{LL} \frac{1-\delta(1-p) + \frac{1-\delta}{\delta p}}{1-\delta(1-2p)}, \quad \forall H, L, p, \delta$$

with $H > L > 0$. Inserting $m(\theta_1,\theta_2) = \theta_1 \cdot \theta_2$, dividing by $(L)^2$, and rearranging terms yields

$$0 \leq \left(\frac{H}{L}\right)^2 \delta p - \frac{H}{L}(1-\delta(1-2p)) + \left(1 - \delta(1-p) + \frac{1-\delta}{\delta p}\right). \tag{81}$$

The parabola in $\frac{H}{L}$ on the right side of (81) is minimized at

$$\frac{H}{L} = 1 + \frac{1-\delta}{2p\delta}. \tag{82}$$

Plugging (82) into (81) gives

$$0 \leq 3 - 2\delta - \delta^2,$$

which holds true as $\delta < 1$. This completes the proof. ■

Proof of Theorem 2. The Impatient Policy is implementable by setting $\tau^{\Theta_S}(\theta) = 0$, for all θ and Θ_S. The implementability of the Positive Assortative Policy requires a proof. We define

$$\Delta^H(p,\delta) = \frac{\delta p}{1-\delta(1-p)}, \qquad \Delta^L(p,\delta) = \frac{\delta(1-p)}{1-\delta p}. \tag{83}$$

The incentive constraint for the productive and the unproductive type are, for market report (0,0),

$$\begin{aligned}
\frac{1}{2}\Delta^H(p,\delta)m_{HH} - \tau^{(0,0)}(H) &\geq \frac{1}{2}\Delta^L(p,\delta)m_{HL} - \tau^{(0,0)}(L),\\
\frac{1}{2}\Delta^H(p,\delta)m_{HL} - \tau^{(0,0)}(H) &\leq \frac{1}{2}\Delta^L(p,\delta)m_{LL} - \tau^{(0,0)}(L),
\end{aligned}$$

for market report (1,1),

$$\begin{aligned}
\frac{1}{2}m_{HH} - \tau^{(1,1)}(H) &\geq \frac{1}{2}m_{HL} - \tau^{(1,1)}(L),\\
\frac{1}{2}m_{HL} - \tau^{(1,1)}(H) &\leq \frac{1}{2}m_{LL} - \tau^{(1,1)}(L),
\end{aligned}$$

for market report (1,0),

$$\begin{aligned}
\frac{1}{2}m_{HH} - \tau^{(1,0)}(H) &\geq \frac{1}{2}\Delta^L(p,\delta)m_{HL} - \tau^{(1,0)}(L),\\
\frac{1}{2}m_{HL} - \tau^{(1,0)}(H) &\leq \frac{1}{2}\Delta^L(p,\delta)m_{LL} - \tau^{(1,0)}(L),
\end{aligned}$$

and for market report (0,1),

$$\begin{aligned}
\frac{1}{2}\Delta^H(p,\delta)m_{HH} - \tau^{(0,1)}(H) &\geq \frac{1}{2}m_{HL} - \tau^{(0,1)}(L),\\
\frac{1}{2}\Delta^H(p,\delta)m_{HL} - \tau^{(0,1)}(H) &\leq \frac{1}{2}m_{LL} - \tau^{(0,1)}(L).
\end{aligned}$$

Combining the incentive constraint of the productive type with the incentive constraint of the unproductive type yields the following conditions on the payment differences:

$$\Delta^H(p,\delta)\frac{m_{HH}}{2} - \Delta^L(p,\delta)\frac{m_{HL}}{2} \geq \tau^{(0,0)}(H) - \tau^{(0,0)}(L) \geq \Delta^H(p,\delta)\frac{m_{HL}}{2} - \Delta^L(p,\delta)\frac{m_{LL}}{2}, \tag{84}$$

$$\frac{1}{2}(m_{HH} - m_{HL}) \geq \tau^{(1,1)}(H) - \tau^{(1,1)}(L) \geq \frac{1}{2}(m_{HL} - m_{LL}), \tag{85}$$

$$\frac{m_{HH}}{2} - \Delta^L(p,\delta)\frac{m_{HL}}{2} \geq \tau^{(1,0)}(H) - \tau^{(1,0)}(L) \geq \frac{m_{HL}}{2} - \Delta^L(p,\delta)\frac{m_{LL}}{2}, \tag{86}$$

$$\Delta^H(p,\delta)\frac{m_{HH}}{2} - \frac{m_{HL}}{2} \geq \tau^{(0,1)}(H) - \tau^{(0,1)}(L) \geq \Delta^H(p,\delta)\frac{m_{HL}}{2} - \frac{m_{LL}}{2}. \tag{87}$$

Thus, an incentive compatible payment difference exists if and only if the following conditions are satisfied:

$$\Delta^H(p,\delta)m_{HH} - \Delta^L(p,\delta)m_{HL} \geq \Delta^H(p,\delta)m_{HL} - \Delta^L(p,\delta)m_{LL}, \tag{88}$$

$$m_{HH} - m_{HL} \geq m_{HL} - m_{LL}, \tag{89}$$

$$m_{HH} - \Delta^L(p,\delta)m_{HL} \geq m_{HL} - \Delta^L(p,\delta)m_{LL}, \tag{90}$$

$$\Delta^H(p,\delta)m_{HH} - m_{HL} \geq \Delta^H(p,\delta)m_{HL} - m_{LL}. \tag{91}$$

Observe that $\Delta^\theta(p,\delta) \leq 1$, for all θ. Hence, (91) implies (88) to (90). To prove that (91) holds whenever the Positive Assortative Policy is optimal, we show that (91) holds if $m_{HL} \leq m^1_{HL}$. Reformulating (91) gives

$$m_{HL} \leq \frac{\Delta^H(p,\delta)}{1 + \Delta^H(p,\delta)}m_{HH} + \frac{1}{1 + \Delta^H(p,\delta)}m_{LL}. \tag{92}$$

We argue that the right-hand side of (92) is larger than m^1_{HL}. To this end, we will show that the multipliers of m_{HH} and m_{LL} in (92) are (weakly) larger than the corresponding factors in m^1_{HL}. First, consider the factor attached to m_{HH}. By (83),

$$\frac{\Delta^H(p,\delta)}{1 + \Delta^H(p,\delta)} = \frac{\delta p}{1 - \delta + 2\delta p}$$

which coincides with the multiplier of m_{HH} in m^1_{HL}. Second, for the factor attached to m_{LL} we obtain

$$\frac{1}{1 + \Delta^H(p,\delta)} = \frac{1 - \delta + \delta p}{1 - \delta + 2\delta p}. \tag{93}$$

(93) is larger than the multiplier of m_{LL} in m^1_{HL} if and only if

$$\frac{1 - \delta + \delta p}{1 - \delta + 2\delta p} \geq \frac{\delta(1-p)}{1 + \delta - 2\delta p} \quad \Leftrightarrow \quad \delta \leq 1.$$

Thus, whenever the Positive Assortative Policy is optimal, we can find an incentive compatible payment pair, for every market report Θ_s.

We construct payments which are positive and individual rational: Set

$$\tau^{(0,0)}(L) = \tau^{(1,0)}(L) = \frac{1}{2}\Delta^L(p,\delta)m_{LL} \geq 0, \tag{94}$$

$$\tau^{(0,1)}(L) = \tau^{(1,1)}(L) = \frac{1}{2}m_{LL} \geq 0. \tag{95}$$

By construction, payments (94) and (95) set the unproductive agent's expected utility to zero and are therefore individual rational. For every market report, choose, given the unproductive type's payment, the maximal payment for the productive type that is consistent with (84) - (87), i.e., such that the payment pair is incentive compatible:

$$\begin{aligned}
\tau^{(0,0)}(H) &= \frac{1}{2}\left(\Delta^L(p,\delta)m_{LL} + \Delta^H(p,\delta)m_{HH} - \Delta^L(p,\delta)m_{HL}\right),\\
\tau^{(1,1)}(H) &= \frac{1}{2}(m_{LL} + m_{HH} - m_{HL}),\\
\tau^{(1,0)}(H) &= \frac{1}{2}\left(\Delta^L(p,\delta)m_{LL} + m_{HH} - \Delta^L(p,\delta)m_{HL}\right),\\
\tau^{(0,1)}(H) &= \frac{1}{2}\left(m_{LL} + \Delta^H(p,\delta)m_{HH} - m_{HL}\right).
\end{aligned}$$

Individual rationality of the payments for the unproductive type and incentive compatibility yield individual rationality for the productive type.
Given that (88) - (91) are satisfied whenever the Positive Assortative Policy is optimal, we can deduce that

$$\begin{aligned}
\tau^{(0,0)}(H) &\geq \frac{1}{2}\left(\Delta^L(p,\delta)m_{LL} + \Delta^H(p,\delta)m_{HL} - \Delta^L(p,\delta)m_{LL}\right) = \frac{1}{2}\Delta^H(p,\delta)m_{HL} \geq 0,\\
\tau^{(1,1)}(H) &\geq \frac{1}{2}(m_{LL} + m_{HL} - m_{LL}) = \frac{1}{2}m_{HL} \geq 0,\\
\tau^{(1,0)}(H) &\geq \frac{1}{2}\left(\Delta^L(p,\delta)m_{LL} + m_{HL} - \Delta^L(p,\delta)m_{LL}\right) = \frac{1}{2}m_{HL} \geq 0,\\
\tau^{(0,1)}(H) &\geq \frac{1}{2}\left(m_{LL} + \Delta^H(p,\delta)m_{HL} - m_{LL}\right) = \frac{1}{2}\Delta^H(p,\delta)m_{HL} \geq 0.
\end{aligned}$$

As all payments are positive, the mechanism runs no deficit. Furthermore, payments support efficient exit because they are charged upon arrival. ■

Proof of Proposition 4. For incentive compatibility, we need to find a single payment pair $(\tau(H), \tau(L))$ such that the difference $\tau(H) - \tau(L)$ satisfies conditions (84) to (87). Observe that (87) yields the lowest upper bound, whereas (86) yields the highest lower bound on the payment difference. Hence, $(\tau(H), \tau(L))$ is incentive compatible if and only if

$$\frac{1}{2}\left(\Delta^H(p,\delta)m_{HH} - m_{HL}\right) \geq \tau(H) - \tau(L) \geq \frac{1}{2}\left(m_{HL} - \Delta^L(p,\delta)m_{LL}\right). \tag{96}$$

Rearranging terms, note that incentive compatible payments exist iff

$$m_{HL} \le \Delta^H(p,\delta)\frac{m_{HH}}{2} + \Delta^L(p,\delta)\frac{m_{LL}}{2}. \tag{97}$$

To see that (97) describes a strict subset of the parameter region in which the Positive Assortative Policy is optimal, we compare it to the boundary of the Positive Assortative Policy m^1_{HL}. We show that (97) is more restrictive than $m_{HL} \le m^1_{HL}$ by separately comparing the factors in front of m_{HH} and m_{LL}. For the factor attached to m_{HH} we note that

$$\frac{1}{2}\frac{\delta p}{(1-\delta(1-p))} < \frac{\delta p}{1-\delta(1-2p)} \quad \Leftrightarrow \quad \delta < 1. \tag{98}$$

Similarly, for the factor in front of m_{LL} observe that

$$\frac{1}{2}\frac{\delta(1-p)}{1-\delta p} < \frac{\delta(1-p)}{1+\delta(1-2p)} \quad \Leftrightarrow \quad \delta < 1. \tag{99}$$

Set $\tau(L) = \frac{1}{2}\Delta^L(p,\delta)m_{LL} \ge 0$. For market reports (1,1) and (0,1), an arriving unproductive type's expected utility from reporting truthfully is

$$\frac{1}{2}m_{LL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL} \ge 0.$$

For market reports (1,0) and (0,0), an arriving unproductive type's expected utility is

$$\frac{1}{2}\Delta^L(p,\delta)m_{LL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL} = 0.$$

Given $\tau(L)$, we choose

$$\tau(H) = \frac{1}{2}\left(\Delta^L(p,\delta)m_{LL} + \Delta^H(p,\delta)m_{HH} - m_{HL}\right)$$

which is consistent with incentive compatibility by (96). Given $\tau(H)$, the productive type's expected utility from truthtelling is

$$\frac{1}{2}(1-\Delta^H(p,\delta))m_{HH} + \frac{1}{2}m_{HL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL} \ge 0,$$

for market reports (1,1) and (1,0), and

$$\frac{1}{2}m_{HL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL} \ge 0$$

for market reports (0,0) and (0,1). Thus, the pair $(\tau(H), \tau(L))$ is individual rational. Furthermore, for the parameter region characterized by (97) it holds that

$$\tau(H) \ge \frac{1}{2}\left(\Delta^L(p,\delta)m_{LL} + m_{HL} - \Delta^L(p,\delta)m_{LL}\right) = \frac{1}{2}m_{HL} \ge 0,$$

therefore, the mechanism runs no deficit. Payments are charged only upon arrival and hence support efficient exit. ∎

Proof of Proposition 5. Fix the share α of the productive agent in the mixed pair. The incentive constraint for the productive and the unproductive type are, for market report (0,0),

$$\frac{1}{2}\Delta^H(p,\delta)m_{HH} \geq \Delta^L(p,\delta)\alpha m_{HL}, \quad \Delta^H(p,\delta)(1-\alpha)m_{HL} \leq \frac{1}{2}\Delta^L(p,\delta)m_{LL},$$

for market report (1,1),

$$\frac{1}{2}m_{HH} \geq \alpha m_{HL}, \quad (1-\alpha)m_{HL} \leq \frac{1}{2}m_{LL},$$

for market report (1,0),

$$\frac{1}{2}m_{HH} \geq \Delta^L(p,\delta)\alpha m_{HL}, \quad (1-\alpha)m_{HL} \leq \frac{1}{2}\Delta^L(p,\delta)m_{LL},$$

and for market report (0,1),

$$\frac{1}{2}\Delta^H(p,\delta)m_{HH} \geq \alpha m_{HL}, \quad \Delta^H(p,\delta)(1-\alpha)m_{HL} \leq \frac{1}{2}m_{LL}.$$

Observe that, for every market report, the incentive constraint of the productive agent provides an upper bound on α, whereas the incentive constraint of the unproductive agent gives a lower bound on α:

$$\frac{1}{2}\frac{\Delta^H(p,\delta)m_{HH}}{\Delta^L(p,\delta)m_{HL}} \geq \alpha \geq \frac{\Delta^H(p,\delta)m_{HL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL}}{\Delta^H(p,\delta)m_{HL}}, \tag{100}$$

$$\frac{1}{2}\frac{m_{HH}}{m_{HL}} \geq \alpha \geq \frac{m_{HL} - \frac{1}{2}m_{LL}}{m_{HL}}, \tag{101}$$

$$\frac{1}{2}\frac{m_{HH}}{\Delta^L(p,\delta)m_{HL}} \geq \alpha \geq \frac{m_{HL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL}}{m_{HL}}, \tag{102}$$

$$\frac{1}{2}\frac{\Delta^H(p,\delta)m_{HH}}{m_{HL}} \geq \alpha \geq \frac{\Delta^H(p,\delta)m_{HL} - \frac{1}{2}m_{LL}}{\Delta^H(p,\delta)m_{HL}}. \tag{103}$$

The incentive constraint of the productive agent given market report (0,1) yields the lowest upper bound, cf. (103), and the incentive constraint of the unproductive agent for market report (1,0) provides the highest lower bound, cf. (102). Thus, any incentive compatible match value split has to satisfy

$$\frac{1}{2}\frac{\Delta^H(p,\delta)m_{HH}}{m_{HL}} \geq \alpha \geq \frac{m_{HL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL}}{m_{HL}}. \tag{104}$$

Note that

$$\frac{1}{2}\frac{\Delta^H(p,\delta)m_{HH}}{m_{HL}} \geq 0 \quad \text{and} \quad 1 \geq \frac{m_{HL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL}}{m_{HL}} \geq \frac{1}{2}.$$

(104) reveals that an incentive compatible match value split exists iff

$$\frac{1}{2}\frac{\Delta^H(p,\delta)m_{HH}}{m_{HL}} \geq \frac{m_{HL} - \frac{1}{2}\Delta^L(p,\delta)m_{LL}}{m_{HL}}.$$

Rearranging terms yields

$$m_{HL} \leq \Delta^H(p,\delta)\frac{m_{HH}}{2} + \Delta^L(p,\delta)\frac{m_{LL}}{2} \tag{105}$$

which coincides with (97).

The Positive Assortative Policy without payments supports efficient exit and provides all agents with (expected) utility of at least zero. Thus, individual rationality is satisfied which concludes the proof. ■

Appendix 2.B Appendix: Supporting Calculations

Calculations for the proof of Theorem 1.

Preliminaries for equations (36)-(39). We derive a couple of useful relationships. Using the definition of $V_{PAP}(0,0)$ and inserting $V_{PAP}(2,0)$, $V_{PAP}(0,2)$, $V_{PAP}(2,1)$, $V_{PAP}(1,2)$, the system (29) - (35) can be reformulated to

$$V_{PAP}(0,0) = \delta[pV_{PAP}(1,0) + (1-p)V_{PAP}(0,1)], \tag{106}$$

$$V_{PAP}(1,0) = \delta[p(m_{HH} + V_{PAP}(0,0)) + (1-p)V_{PAP}(1,1)], \tag{107}$$

$$V_{PAP}(0,1) = \delta[pV_{PAP}(1,1) + (1-p)(m_{LL} + V_{PAP}(0,0))], \tag{108}$$

$$V_{PAP}(1,1) = \delta[p(m_{HH} + V_{PAP}(0,1)) + (1-p)(m_{LL} + V_{PAP}(1,0))]. \tag{109}$$

Firstly, consider $V_{PAP}(1,0) - V_{PAP}(0,1)$. By (107) and (108), we obtain

$$V_{PAP}(1,0) - V_{PAP}(0,1) =$$
$$= \delta[p(m_{HH} + V_{PAP}(0,0) - V_{PAP}(1,1)) + (1-p)(-m_{LL} + V_{PAP}(1,1) - V_{PAP}(0,0))],$$

which yields, inserting (106) and (109), the equation

$$\begin{aligned} V_{PAP}(1,0) - V_{PAP}(0,1) \;=& \delta p\,[m_{HH} + \delta p(V_{PAP}(1,0) - m_{HH} - V_{PAP}(0,1)) \\ & +\delta(1-p)(V_{PAP}(0,1) - m_{LL} + V_{PAP}(1,0))] \\ & +\delta(1-p)\,[-m_{LL} + \delta p(m_{HH} + V_{PAP}(0,1) - V_{PAP}(1,0)) \\ & +\delta(1-p)(m_{LL} + V_{PAP}(1,0) - V_{PAP}(0,1))]\,. \end{aligned}$$

Solving for $V_{PAP}(1,0) - V_{PAP}(0,1)$ gives

$$V_{PAP}(1,0) - V_{PAP}(0,1) \;=\; \frac{\delta[pm_{HH}(1-2p\delta+\delta) - (1-p)m_{LL}(1-\delta+2p\delta)]}{1-\delta^2(1-2p)^2}. \tag{110}$$

Similarly, by (109) and (106), we obtain

$$\begin{aligned} & V_{PAP}(1,1) - V_{PAP}(0,0) \;= \\ & \delta[p(m_{HH} + V_{PAP}(0,1) - V_{PAP}(1,0)) + (1-p)(m_{LL} + V_{PAP}(1,0) - V_{PAP}(0,1))], \end{aligned}$$

which gives, inserting (107) and (108), the equation

$$\begin{aligned} V_{PAP}(1,1) - V_{PAP}(0,0) \;=& \delta p\,[m_{HH} + \delta p(V_{PAP}(1,1) - m_{HH} - V_{PAP}(0,0)) \\ & +\delta(1-p)(m_{LL} + V_{PAP}(0,0) - V_{PAP}(1,1))] \\ & +\delta(1-p)\,[m_{LL} + \delta p(m_{HH} + V_{PAP}(0,0) - V_{PAP}(1,1)) \\ & +\delta(1-p)(V_{PAP}(1,1) - m_{LL} - V_{PAP}(0,0))]\,. \end{aligned}$$

Solving for $V_{PAP}(1,1) - V_{PAP}(0,0)$ yields

$$V_{PAP}(1,1) - V_{PAP}(0,0) \;=\; \frac{\delta[pm_{HH}(1-2p\delta+\delta) + (1-p)m_{LL}(1-\delta+2p\delta)]}{1-\delta^2(1-2p)^2}. \tag{111}$$

Comparing (111) to (110), observe that

$$V_{PAP}(1,1) - V_{PAP}(0,0) \;=\; V_{PAP}(1,0) - V_{PAP}(0,1) + m_{LL}\cdot \underbrace{\frac{2\delta(1-p)}{1+\delta-2p\delta}}_{:=A} \tag{112}$$

with

$$0 \;<\; A \;<\; 1. \tag{113}$$

Exploiting (112), inequalities (36) to (39) can be reformulated to

$$V_{PAP}(1,0) - V_{PAP}(0,1) \;\geq\; m_{HL} - m_{LL}, \tag{114}$$

$$m_{HH} - m_{HL} \;\geq\; V_{PAP}(1,0) - V_{PAP}(0,1), \tag{115}$$

$$m_{HH} + m_{LL} - m_{HL} \;\geq\; V_{PAP}(1,0) - V_{PAP}(0,1) + m_{LL}A, \tag{116}$$

$$V_{PAP}(1,0) - V_{PAP}(0,1) \;\geq\; m_{HL} - m_{LL}A. \tag{117}$$

Deriving equation (36). It is immediate that (117) implies (114), i.e., that (39) implies (36).

Deriving equation (38). Similarly (115) implies (116), i.e., (37) implies (38).

Deriving equation (39). Inserting (110), (117) can be written as

$$m_{HL} \;\leq\; m_{HH}\frac{\delta p}{1-\delta+2p\delta} + m_{LL}\frac{\delta(1-p)}{1+\delta-2p\delta} \tag{118}$$

which corresponds exactly to $m_{HL} \leq m_{HL}^{1}$.

Deriving equation (37). We argue that (39) implies (37), i.e., (118) implies (115). Inserting (110), (115) can be written as

$$m_{HL} \;\leq\; m_{HH}\left(1-\frac{\delta p}{1-\delta+2p\delta}\right) + m_{LL}\frac{\delta(1-p)}{1+\delta-2p\delta}. \tag{119}$$

Therefore (118) implies (115) if

$$\frac{2\delta p}{1-\delta+2p\delta} \;\leq\; 1 \qquad \Leftrightarrow \qquad \delta \;\leq\; 1, \tag{120}$$

which holds.

Deriving equation (51). Rearranging terms to isolate $V_{PIP}(1,0) - V_{PIP}(0,1)$ and inserting the expression for $V_{PIP}(1,0) - V_{PIP}(0,1)$, (51) can be rewritten as

$$m_{HH} - m_{HL} \;>\; \delta[pm_{HH} + (1-2p)m_{HL} - (1-p)m_{LL}]. \tag{121}$$

If (121) holds for $\delta = 1$, it holds for any δ. Setting $\delta = 1$ and rearranging terms yields $m_{HH} + m_{LL} > 2m_{HL}$ which is satisfied by assumption.

Deriving equation (61). We argue that (60) implies (61):

$$\begin{aligned}
&\delta[pV_{MIP}(2,2) + (1-p)V_{MIP}(1,3)] \\
&\qquad = \delta[p(m_{HH} + m_{LL} + V_{MIP}(0,0)) + (1-p)(m_{HL} + m_{LL} + V_{MIP}(0,0))] \\
&\qquad = \delta m_{LL} + V_{MIP}(1,0) \\
&\qquad < m_{LL} + V_{MIP}(1,0) \\
&\qquad \leq m_{HL} + V_{MIP}(0,1),
\end{aligned}$$

where the last inequality follows from (60).

Deriving equation (58). We show (58) exploiting supermodularity and optimality of (0, 1, 0) on state $(1,2)$:

$$\begin{aligned} m_{HL} + \delta[pV_{MIP}(2,2)) + (1-p)V_{MIP}(1,3)] \;&\leq\; m_{HL} + V_{MIP}(1,2) \\ &=\; 2m_{HL} + V_{MIP}(0,1) \\ &\leq\; m_{HH} + m_{LL} + V_{MIP}(0,1). \end{aligned}$$

Deriving equation (57). We show (57) exploiting supermodularity and optimality of (0, 1, 0) on state $(1,1)$:

$$\begin{aligned} m_{HL} + \delta[pV_{MIP}(2,1)) + (1-p)V_{MIP}(1,2)] \;&\leq\; m_{HL} + V_{MIP}(1,1) \\ &=\; 2m_{HL} + V_{MIP}(0,0) \\ &\leq\; m_{HH} + m_{LL} + V_{MIP}(0,0). \end{aligned}$$

Deriving equation (56). (59) and (60) imply (56): (56) is equivalent to

$$\delta V_{MIP}(0,1) - V_{MIP}(0,0) \;\leq\; m_{HL} - \delta p m_{HH} - \delta(1-p)m_{HL},$$

similarly, (59) is equivalent to

$$\delta V_{MIP}(0,1) - V_{MIP}(0,0) \;\leq\; m_{LL} - \delta p m_{HL} - \delta(1-p)m_{LL}.$$

Therefore, a sufficient condition for (56) to hold is

$$m_{LL} - \delta p m_{HL} - \delta(1-p)m_{LL} \;\leq\; m_{HL} - \delta p m_{HH} - \delta(1-p)m_{HL},$$

which is equivalent to $m_{HL} \geq m^2_{HL}$, i.e. (60).

Deriving $m_{HL} \leq m^3_{HL}$. First, note that whenever there exists a m_{HL} such that $m_{HL} > m^3_{HL}$, then $\frac{1}{2}m_{HH} + \frac{1}{2}m_{LL} > m^3_{HL}$. Reformulating $m_{HL} \leq m^3_{HL}$, $\forall p, \delta, m_{HH}, m_{LL}$, gives

$$m_{HL} - m_{LL} \;\leq\; \frac{1-\delta}{\delta p}m_{LL} + \delta[p(m_{HH} - m_{HL}) + (1-p)(m_{HL} - m_{LL})], \quad \forall p, \delta, m_{HH}, m_{LL}.$$

Inserting $m_{HL} = \frac{1}{2}m_{HH} + \frac{1}{2}m_{LL}$ and rearranging terms yields

$$\frac{1}{2}(m_{HH} - m_{LL}) \;\leq\; \frac{m_{LL}}{\delta p}, \quad \forall p, \delta, m_{HH}, m_{LL}.$$

This holds if and if

$$\frac{1}{2}(m_{HH} - m_{LL}) \;\leq\; m_{LL}, \quad \forall m_{HH}, m_{LL},$$

i.e., $m_{HH} \leq 3m_{LL}$, $\forall m_{HH}, m_{LL}$.

Calculations for the proof of Proposition 1.

Preliminaries for equations (69) - (77). It is instructive to rewrite (62) - (68):

$$
\begin{aligned}
V_{P1}(0,0) &= \delta[pV_{P1}(1,0) + (1-p)V_{P1}(0,1)], && (122)\\
V_{P1}(1,0) &= \delta[p(m_{HH} + V_{P1}(0,0)) + (1-p)(m_{HL} + V_{P1}(0,0))], && (123)\\
V_{P1}(0,1) &= \delta[p(m_{HL} + V_{P1}(0,0)) + (1-p)V_{P1}(0,2)], && (124)\\
V_{P1}(0,2) &= \delta[p(m_{HL} + V_{P1}(0,1)) + (1-p)(m_{LL} + V_{P1}(0,1))]. && (125)
\end{aligned}
$$

Plugging (123) into (122) gives

$$V_{P1}(0,0) = \delta p(\delta V_{P1}(0,0) + \delta(pm_{HH} + (1-p)m_{LL})) + \delta(1-p)V_{P1}(0,1). \quad (126)$$

For $V_{P1}(0,1) - \delta V_{P1}(0,2)$ we obtain, inserting (124) and (125),

$$
\begin{aligned}
V_{P1}(0,1) - \delta V_{P1}(0,2) &= \delta p(m_{HL} + V_{P1}(0,0) - \delta m_{HL} - \delta V_{P1}(0,1))\\
&\quad + \delta(1-p)(V_{P1}(0,2) - \delta m_{LL} - \delta V_{P1}(0,1)). \quad (127)
\end{aligned}
$$

Rearranging (125) we see that

$$V_{P1}(0,2) - \delta V_{P1}(0,1) = \delta(pm_{HL} + (1-p)m_{LL}). \quad (128)$$

Inserting (124), (126), and (128) into (127), we can solve for $V_{P1}(0,1) - \delta V_{P1}(0,2)$ which is explicitly given by

$$V_{P1}(0,1) - \delta V_{P1}(0,2) = \frac{\delta p\left[m_{HL} - \delta m_{LL} + m_{HH}p^2\delta^2 - m_{HL}p^2\delta^2 + m_{LL}p\delta - m_{HL}p\delta\right]}{1 + p^2\delta^2 - p\delta^2}. \quad (129)$$

Next, we derive a closed-form expression for $V_{P1}(0,0) - \delta V_{P1}(0,1)$. To this end, note that

$$
\begin{aligned}
V_{P1}(0,0) - \delta V_{P1}(0,1) &= \delta p(V_{P1}(1,0) - \delta m_{HL} - \delta V_{P1}(0,0))\\
&\quad + \delta(1-p)(V_{P1}(0,1) - \delta V_{P1}(0,2)). \quad (130)
\end{aligned}
$$

Furthermore by (123)

$$V_{P1}(1,0) - \delta V_{P1}(0,0) = \delta(pm_{HH} + (1-p)m_{HL}). \quad (131)$$

Inserting (131) and (129) into (130) gives

$$V_{P1}(0,0) - \delta V_{P1}(0,1) = \frac{p\delta^2\left[m_{HH}p + m_{HL}(1 - 2p - p\delta + p^2\delta) + m_{LL}(-\delta + 2p\delta - p^2\delta)\right]}{1 - p\delta^2 + p^2\delta^2}. \quad (132)$$

Deriving equation (72). Rearranging terms in (72) gives

$$V_{P1}(0,1) - \delta V_{P1}(0,2) \geq \delta[pm_{HL} + (1-p)m_{LL}] - m_{LL}. \quad (133)$$

Plugging (129) into (133) and rewriting (133) as a condition on m_{HL}, we obtain

$$m_{HL} \;\leq\; \frac{1}{1-\delta(1-2p)}\left[m_{HH}\delta p + m_{LL}\left(2-\delta(2-p)+\frac{1-\delta+\delta p(1-\delta)^2}{\delta^2 p^2}\right)\right]. \tag{134}$$

The term on the right side of (134) is m_{HL}^4.

Deriving equation (74). We argue that (74) holds if and only if $m_{HL} \geq m_{HL}^3$. Inserting $V_{P1}(1,2)$ and $V_{P1}(0,3)$ in (74) gives

$$\delta\left[p(m_{HL}+V_{P1}(0,1))+(1-p)(m_{LL}+V_{P1}(0,1))\right] \;\geq\; m_{LL}+V_{P1}(0,0).$$

Rearranging terms yields

$$\delta\left[pm_{HL}+(1-p)m_{LL}\right]-m_{LL} \;\geq\; V_{P1}(0,0)-\delta V_{P1}(0,1). \tag{135}$$

Plugging (132) into (135) and some algebra yields

$$m_{HL} \;\geq\; m_{HH}\frac{\delta p}{1-\delta(1-2p)}+m_{LL}\frac{1-\delta(1-p)+\frac{1-\delta}{\delta p}}{1-\delta(1-2p)}. \tag{136}$$

Observe that the right side of (136) coincides with m_{HL}^3.

Deriving equation (76). We show that (72) and (74) imply (76). Reformulating (72) gives

$$m_{LL}-\delta p m_{HL}-\delta(1-p)m_{LL} \;\geq\; \delta V_{P1}(0,2)-V_{P1}(0,1), \tag{137}$$

and (76) gives

$$m_{HL}-\delta p m_{HH}-\delta(1-p)m_{HL} \;\geq\; \delta V_{P1}(0,2)-V_{P1}(0,1). \tag{138}$$

(137) implies (138) if

$$m_{HL}-m_{LL} \;\geq\; \delta p(m_{HH}-m_{HL})+\delta(1-p)(m_{HL}-m_{LL}). \tag{139}$$

Note that (139) coincides with $m_{HL} \geq m_{HL}^2$. As (74) requires $m_{HL} \geq m_{HL}^3$ and $m_{HL}^3 \geq m_{HL}^2$, (139) is implied by (74).

Deriving equation (70). We argue that (74) implies (70). Note that we can rewrite (74) as

$$V_{P1}(0,0)-\delta V_{P1}(0,1) \;\leq\; \delta\left(pm_{HL}+(1-p)m_{LL}\right)-m_{LL}, \tag{140}$$

and (70) as

$$V_{P1}(0,0) - \delta V_{P1}(0,1) \;\geq\; \delta\left(pm_{HH} + (1-p)m_{HL}\right) - m_{HL}. \tag{141}$$

Plugging (132) into (140) and (141) yields, after some algebra, for (140)

$$\delta^2 p^2 (m_{HH} - m_{HL}) + m_{LL}(1-\delta) \;\leq\; \delta^2 p^2 (m_{HL} - m_{LL}) + (1-\delta)p\delta(m_{HL} - m_{LL}), \tag{142}$$

and for (141)

$$\delta p(m_{HH} - m_{HL})(1 - \delta p - \delta^2 p(1-p)) + (1-p)p^2\delta^3(m_{HL} - m_{LL}) \;\leq\; m_{HL}(1-\delta) + (1-p)p\delta^3(m_{HL} - m_{LL}). \tag{143}$$

It is sufficient for (140) to imply (141), i.e. (142) to imply (143), if it holds that

$$(\delta p(m_{HL} - m_{LL}) + (1-\delta)(m_{HL} - m_{LL}))(1 - \delta p - \delta^2 p(1-p)) \;\leq\; m_{HL}(1-\delta) + (1-p)^2 p\delta^3(m_{HL} - m_{LL}). \tag{144}$$

For (144) in turn it is sufficient if

$$(\delta p + (1-\delta))(1 - \delta p - \delta^2 p(1-p)) \;\leq\; (1-\delta) + (1-p)^2 p\delta^3. \tag{145}$$

Simplifying (145) shows that the term on the left side coincides with the term on the right side.

Deriving equation (69). Supermodularity and (74) imply (69). Inserting (123) and (124) into (69) and rearranging terms yields

$$m_{HH} - m_{HL} \;\geq\; \delta\left[p(m_{HH} - m_{HL}) + (1-p)(m_{HL} + V_{P1}(0,0) - V_{P1}(0,2))\right]. \tag{146}$$

By (74) it is sufficient for (146) to hold that

$$m_{HH} - m_{HL} \;\geq\; \delta\left[pm_{HH} + (1-2p)m_{HL} - (1-p)m_{LL}\right]. \tag{147}$$

Reformulating yields

$$\frac{m_{HH} - m_{HL}}{m_{HL} - m_{LL}} \;\geq\; \frac{\delta(1-p)}{1 - \delta p}. \tag{148}$$

By supermodularity the left side of (148) is larger than one, whereas the right side of (148) is smaller than one. Therefore, (74) implies (69).

Deriving equation (77). (77) is implied by (75), (76), and supermodularity: Given (75) and (76), from optimality on state $(1,2)$ we know

$$m_{HL} + \delta[pV_{P1}(2,2) + (1-p)V_{P1}(1,3)] \;\leq\; 2m_{HL} + V_{P1}(0,1), \tag{149}$$

and from supermodularity follows

$$2m_{HL} + V_{P1}(0,1) \ \leq \ m_{HH} + m_{LL} + V_{P1}(0,1). \tag{150}$$

Combining (149) and (150) gives (77).

Deriving equation (71). (71) is implied by (70) and supermodularity: Given (70), from optimality on state $(1,1)$ we know

$$m_{HL} + \delta[pV_{P1}(2,1) + (1-p)V_{P1}(1,2)] \ \leq \ 2m_{HL} + V_{P1}(0,0), \tag{151}$$

and from supermodularity follows

$$2m_{HL} + V_{P1}(0,0) \ \leq \ m_{HH} + m_{LL} + V_{P1}(0,0). \tag{152}$$

Combining (151) and (152) gives (71).

Deriving equation (73). (73) is implied by (74), (70), and supermodularity: Given (74), from optimality on state $(0,2)$ we know

$$m_{HL} + V_{P1}(0,2) \ \geq \ m_{HL} + m_{LL} + V_{P1}(0,0), \tag{153}$$

and given (70), from optimality on state $(1,1)$ follows

$$m_{HH} + m_{LL} + V_{P1}(0,0) \ \leq \ m_{LL} + \delta[pV_{P1}(2,1) + (1-p)V_{P1}(1,2)]. \tag{154}$$

Combining (153) and (154) gives (73).

Deriving equation (75). We argue that (74) implies (75) by showing that "not (75)" implies 'not (74)'. Applying "not (75)" and (74) leads to a contradiction:

$$\begin{aligned} m_{LL} + V_{P1}(0,0) \ &\leq \ \delta[p(m_{HL} + V_{P1}(0,1)) + (1-p)(m_{LL} + V_{P1}(0,1))] \\ &\leq \ \delta[p(m_{LL} + V_{P1}(1,0)) + (1-p)(m_{LL} + V_{P1}(0,1))] \\ &= \ \delta m_{LL} + V_{P1}(0,0). \end{aligned}$$

3

Target Mass or Class? Dynamic Reputation Management with Heterogeneous Consumption Externalities*

This chapter studies a seller whose reputation is determined by the types of her customers. In our model, a monopolist repeatedly sells a good to heterogeneous customers who, depending on their type, increase or decrease the seller's reputation. First, we study a trade-off between realizing current-period profits and building reputation for future periods. Second, we analyze reputation dynamics. Over time, reputation always converges to a stable level. Convergence behavior, however, depends strongly on the good's durability. While the reputation of less durable goods fluctuates around the long-run reputation, the reputation of more durable goods converges monotonically.

3.1 Introduction

The reputation of several goods is substantially influenced by the types of their consumers. This effect is particularly pronounced for luxury and fashion goods, such as exclusive watches and fashionable apparel. When considering buying such a good, consumers not only take into account the utility from the good itself but also whether they want to be associated with its clientele. *"Good-looking people attract other good-looking people"*, as Mike Jeffries, at that time CEO of *Abercrombie & Fitch*, states in a controversial interview in 2006.[1] Economically, past buyers impose an externality on new buyers.

* Based on joint work with Tim Frommeyer (see Frommeyer and Schickner, 2016). We thank Francesc Dilmé, Benny Moldovanu, and Deszö Szalay for guidance throughout the project. Furthermore, we wish to thank Daron Acemoglu, Simon Board, Gonzalo Cisternas, Pauli Murto, Aniko Öry, and participants of the Workshop on Dynamic Incentive Problems 2015 in Bonn for helpful comments and suggestions.

[1] See Denizet-Lewis (January 24, 2006).

This externality can be positive or negative depending on whether the association with the good's clientele is desirable for new buyers. Sellers seem to be aware of their consumers' concerns and strategically try to target a certain clientele. In the same interview, Mike Jeffries emphasizes: *"(...) we want to market to cool, good-looking people. We don't market to anyone other than that."* He goes on to say: *"Are we exclusionary? Absolutely. Those companies that are in trouble are trying to target everybody: young, old, fat, skinny. But then you become totally vanilla. You don't alienate anybody, but you don't excite anybody, either."* These quotes exemplify how sellers manage their reputation by targeting a specific clientele, in Jeffries' words "good-looking people", in order to attract new consumers who wish to be associated with this clientele by possessing the same good. In turn, new consumers influence the good's reputation. How a seller manages reputation depends on the characteristics of the market and the good and may differ substantially across markets and goods.

In this chapter, we present a dynamic model to study how a monopolist optimally manages her reputation by targeting specific clienteles in order to maximize profits. In particular, we are interested in how the underlying characteristics of the market affect the evolution of reputation, demand, and prices over time.

The reputation of a seller has several facets. Traditionally, a good's reputation is linked to its intrinsic properties such as quality. As these properties are usually private information of the seller, we refer to this facet of reputation as *private reputation*. Goods with high private reputation are believed to be of better quality, for example, than goods with low private reputation. However, reputation also has a public facet. Often, buyers do not only derive utility from the good's intrinsic properties but also from what the good symbolizes to others. In itself this *public reputation* has many aspects. For example, a seller forms her public reputation through advertisement. In this chapter, we study the effect of the seller's clientele on her public reputation.

For most goods both facets of reputation are present. However, the importance of one or the other may vary across goods, time, and cultures. In this chapter, we focus solely on the public reputation a seller derives from her clientele. Whenever referring to reputation in the following, we mean reputation in this sense unless explicitly stated otherwise.

In our model, a profit-maximizing monopolist repeatedly offers a good to a continuum of heterogeneous, short-lived buyers. Buyers are ordered according to their type which describes their effect on the seller's reputation. Selling to buyers of higher types improves the seller's reputation, and selling to buyers of lower types decreases the seller's reputation. A buyer's willingness to pay increases in the seller's reputation and in their type.[2] Therefore, the seller can choose her clientele through her pricing strategy and thereby manage her reputation. For any price, buyers purchase the good

[2] Intuitively, the higher a buyer's type, the more reputation improves after he purchases the good. The buyer anticipates his influence and, hence, his willingness to pay increases in his type.

only if their type exceeds a cutoff type.

In the first part of this chapter, we consider a general model in which types are drawn from a general, continuous distribution and a buyer's utility is quasilinear. Reputation tomorrow is a function of reputation today and today's cutoff type, which characterizes the current clientele. We impose, first, that reputation is persistent, i.e., ceteris paribus a higher reputation today yields a higher reputation tomorrow, and, second, that reputation is increasing in today's clientele, i.e., selling to an exclusive group of high-type buyers increases reputation, whereas selling to a broader group of buyers with heterogeneous types decreases reputation.

Within this general framework, we start by establishing existence of a Markov perfect equilibrium and characterize the seller's value in any equilibrium. We then argue that the seller's value is increasing in reputation which implies that reputation is beneficial for the seller. In each period, the seller solves an intertemporal trade-off. On the one hand, selling to a small, exclusive clientele increases her reputation and, hence, her future profits. On the other hand, limiting demand in this way decreases her profits today.

Conventionally, high prices increase the seller's profit by increasing revenue. In our model, in addition, the seller sets high prices to target an exclusive clientele in order to improve her reputation. In other words, high prices prevent buyers with lower types from purchasing the good and, hence, protect the seller's reputation.

To obtain explicit results regarding the reputation dynamics, we specialize our setup in the second part of this chapter. We assume that a buyer's utility is linear and that the distribution of types is uniform. Reputation tomorrow is a convex combination of today's reputation and the cutoff type. We associate the weight of the convex combination on current-period reputation with the good's durability. Intuitively, if the durability of the good is higher, buyers possess the good for a longer period of time. Consequently, the good is longer associated with their type, and the influence of past buyers' types, captured by today's reputation, on tomorrow's reputation is comparatively high. Conversely, if the durability of the good is lower, new buyers constitute a significant fraction of the seller's clientele. Thus, the influence of today's buyers' types, characterized by today's cutoff type, on tomorrow's reputation is relatively high. An example for a market that is characterized by comparatively high durability is the market for watches, whereas, for example, the market for fashionable apparel is characterized by comparatively low durability. With these adjustments, we obtain a linear-quadratic setup. We determine the seller's value and policy function in closed form which makes the setup tractable for a more explicit analysis.

Next, we study optimal reputation dynamics, in particular, their dependence on the good's durability. First, we show that reputation always converges to a long-run reputation. Although it is optimal in the short run for the seller to target different clienteles, this result implies that it is optimal in the long run to target a fixed clientele and maintain a constant reputation. In contrast to the private reputation literature, e.g. Holmström (1999), Cripps et al. (2004), and Cripps et al. (2007), reputation is not a short run phenomenon. Even in the long run, the seller trades off the benefit of increasing her reputation against realizing higher current-period profits. The long-run reputation is increasing in the discount factor and decreasing in the good's durability.

Second, convergence behavior towards the long-run reputation is substantially different for goods with different durability. If the durability of the good is below a threshold, reputation oscillates towards the long-run reputation. A period of high reputation is followed by reputation of low reputation and vice versa. If the durability is above the threshold, reputation and price dynamics are monotone. If the initial reputation is high, reputation decreases monotonically to the long-run reputation, and, in contrast, increases monotonically to the long-run reputation if reputation is initially low.

Despite the substantially different convergence behavior of the cases described above, we identify an underlying monotonicity in the degree of fluctuations across these cases. To this end, we determine an appropriate measure for fluctuations, the normalized distance between two subsequent reputation levels, and show that fluctuations are monotonically decreasing in durability, i.e., the higher the durability of the good the less reputation fluctuates over time.

Our model predicts substantial fluctuations in reputation and prices for goods with low durability and relatively stable, monotone reputation dynamics for goods with high durability. There are many factors that drive price and reputation dynamics. The relative importance of these factors may vary significantly across markets. Nevertheless, our findings seem to be in line with anecdotal evidence. As an example for a durable good, consider the Swiss watchmaker *Rolex*. *Business Insider* documents how prices of *Rolex* watches have steadily increased over the last sixty years, both in absolute terms and as measured as a proportion of average income. At the same time its reputation seems to have improved constantly: *"(...) today's Submariner, the tool-watch of yesterday, has transformed into an internationally recognized status symbol (...)."*[3] As an example for a good with low durability, it is insightful to come back to *Abercrombie & Fitch*. During its reputational high before and at the time of Jeffries' interview, customers were queuing in front its stores. Simultaneously, *Abercrombie & Fitch* was expanding considerably. The German newspaper *FAZ* notes in an article that the brand has "lost its coolness" since then, which some analysts attribute to the fact that it has become too widespread. Further, the authors observe that *Abercrombie & Fitch* is currently cutting back its network of stores.[4] In a similar vein, consider the rapid rise and decline of the fashion label *Ed Hardy*. At its reputational height, many celebrities wore *Ed Hardy* clothes. According to a *CNN* article, the designer himself attributes the subsequent fall of the brand to the fact that " *(...) widespread licensing aspired to make the brand*

[3] See Bredan (January 1, 2015).
[4] See Lindner and Löhr (August 31, 2015).

more accessible to people at every price point." An analyst observes that, as a result, *Ed Hardy* became *"very trailer park"* and states that *"they made it too unexclusive."*[5] Similar upward and downward fluctuations in reputation due to an expansion and contraction of clientele are documented for the fashion labels *Burberry* and *Louis Vuitton*.[6]

Finally, we argue that public reputation provides another rationale for planned obsolescence. Intuitively, if a seller starts with a low reputation, she can improve her reputation by including higher types into her clientele more quickly if the durability of the good is low. This reputational explanation of planned obsolescence complements the traditional demand-driven explanation.

Related Literature. In the private reputation literature, the seller's reputation represents the market's belief about her unknown type, productivity, or quality. In Kreps et al. (1982), Mailath and Samuelson (2001), and Cripps et al. (2004), reputation reflects the market's belief that the seller is a competent type who strategically chooses her effort level instead of a behavioral type who always exerts the same effort. In Holmström (1999), Tadelis (1999), Board and Meyer-ter-Vehn (2013), and Dilmé (2016), reputation is the buyers' belief about the seller's productivity or quality which influences their utility from purchasing the seller's good. Whereas in these models buyers update their beliefs based on openly observable information, it is costly to acquire information about the seller's quality in Liu (2011) as well as in J. Lee and Liu (2013). Overall, the models in the private reputation literature are dynamic and, typically, a monopolistic, long-living seller faces a sequence of myopic, short-lived buyers. For a more detailed overview, see, for example, Bar-Isaac and Tadelis (2008).

To illustrate the difference to the public reputation studied in our work, consider the classical example of a restaurant from the private reputation literature. The restaurant's private reputation describes the customers' belief that food and service are of high quality. Public reputation describes the clientele a customer is associated with when visiting the restaurant. This could be, for example, students or politicians and business men.

The literature on consumption externalities distinguishes the *bandwagon effect*, the *snob effect*, and the *Veblen effect* (Veblen, 1899; Leibenstein, 1950). The bandwagon effect and the snob effect describe the case when demand increases or decreases, respectively, if others consume the same good. The Veblen effect refers to conspicuous consumption, that is, when demand for a good increases in its price. In our model, there are bandwagon effects with respect to high-type buyers and snob effects with respect to low-type buyers. Becker (1991) and Becker and Murphy (1993) consider consumers who care about who else possesses a good and find snob and bandwagon effects. Bagwell and Bernheim (1996) as well as Amaldoss and Jain (2005a) show that Veblen

[5] See Alabi (September 30, 2013).
[6] See also Lindner and Löhr (August 31, 2015).

effects can arise when consumers signal wealth from the consumption of conspicuous goods or when consumers have social needs, such as desire for prestige. Grilo et al. (2001) and Amaldoss and Jain (2005b) examine optimal pricing of conspicuous goods in static duopoly competition.

In our work, we combine the dynamic approach of the private reputation literature with consumption externalities. Pesendorfer (1995) and Hashimoto and Matsubayashi (2014) also consider the problem of a monopolistic seller in a dynamic model with consumption externalities. Pesendorfer (1995) examines fashion cycles in a model where consumers of different types buy a good to distinguish themselves from each other. The good's reputation stems from its ability to signal a buyer's type in a secondary marriage market. Over time, more consumers buy the good such that the signaling effect and, hence, reputation and prices decrease monotonically. When the price has dropped low enough, the seller introduces a new, initially exclusive product line. As opposed to the focus on optimal replacement of a product in Pesendorfer (1995), we examine the effects of changes in the good's characteristics on the seller's reputation management. Moreover, considering the reputation of a good, we obtain richer dynamics. In our model, reputation is not always monotonically decreasing. Depending on initial reputation and the good's durability, reputation can monotonically increase, monotonically decrease, or oscillate. Hashimoto and Matsubayashi (2014) analyze optimal pricing of a monopolistic seller in a dynamic model with either positive or negative consumption externalities. Consumers' utilities depend on past sales which the seller anticipates when solving her pricing problem. For the case where consumption externalities are negative, their results are in line with ours. If the consumers' utilities only depend on recent sales, reputation oscillates downwards, whereas reputation decreases monotonically if consumers discount past sales less strongly. In contrast to their paper, however, in our model some consumers exert positive consumption externalities whereas others exert negative externalities. Furthermore, we provide a full characterization of the seller's optimal pricing strategy, varying the characteristics of the market. Moreover, we model the reputation channel more explicitly through which buyers influence others' utilities from buying the good.

In contrast to social consumption externalities as studied in our and the above works, another branch of literature studies technological consumption externalities. These arise when the utility of a good increases with the size of its user base, such as telephones or social networks. Katz and Shapiro (1985) as well as I. H. Lee and Mason (2001) analyze price competition between firms under positive or negative network effects, and Dhebar and Oren (1985), Bensaid and Lesne (1996), and Gabszewicz and Garcia (2008) consider dynamic pricing strategies of a monopolist.

Further related is Rayo (2013) who studies a monopolist selling a good that signals social status in a screening model as well as the literature on peer groups (see, for example, Board, 2009), and the literature on scarcity (see, for example, Stock and Balachander, 2005).

The rest of this chapter is organized as follows. We introduce the model in Section 3.2. In Section 3.3, we derive the existence result and the trade-off between building reputation and current-period profits in the general model. In Section 3.4, we proceed by studying reputation dynamics in the specified model. In Section 3.5, we discuss the implications on planned obsolescence. Section 3.6 concludes.

3.2 Model

Time is discrete and infinite, $t \in \{0, \ldots\}$. A single long-lived seller (she) repeatedly offers a good to a unit mass of short-lived buyers (he).

Seller. In each period t, the long-lived seller sets a price p_t for the good. She has no production costs and discounts future payoffs with discount factor $\delta \in (0,1)$.

Buyers. Every period-t buyer is characterized by his type $\theta \in [0,1]$ which is distributed according to a continuous distribution function $F(\theta)$ with support $[0,1]$. The types of buyers who purchase the good in periods $0, 1, \ldots, t-1$ determine the *reputation* of the good at time t which we denote by $\lambda_t \in [0,1]$. A period-t buyer with type θ who purchases the good with reputation λ_t at price p_t receives utility

$$u(\theta, \lambda_t) - p_t,$$

where $u : [0,1]^2 \to \mathbb{R}_{\geq 0}$ is strictly increasing in both arguments and continuous. We do not explicitly model buyers as players in the game. Notice that there exists a cutoff type $\theta^\dagger(\lambda_t, p_t) \in [0,1]$ for every reputation λ_t and every price p_t such that $u(\theta^\dagger(\lambda_t, p_t), \lambda_t) = p_t$.[7] We assume that all buyers with type $\theta \geq \theta^\dagger(\lambda_t, p_t)$ purchase the good and all buyers with type $\theta < \theta^\dagger(\lambda_t, p_t)$ do not purchase the good. We justify this assumption with the following reasoning. Assume that each buyer decides whether to buy the good, and normalize buyers' utility from not buying the good to zero. Consequently, a period-t buyer with type θ purchases the good only if

$$u(\theta, \lambda_t) - p_t \;\geq\; 0. \tag{1}$$

Recall that u is strictly increasing in the first argument. If it is optimal for a buyer with type θ to buy the good then it is optimal for any buyer with type $\theta' \geq \theta$ to buy the good as well. Analogously, if it is optimal for a buyer with type θ not to buy the good then it is also optimal for any buyer with type $\theta' \leq \theta$ not to buy the good. The cutoff

[7] For $p_t < u(0, \lambda_t)$ we set $\theta^\dagger(\lambda_t, p_t) = 0$, and for $p_t > u(1, \lambda_t)$ we set $\theta^\dagger(\lambda_t, p_t) = 1$.

is given by the type for which (1) holds with equality.[8] Thus, the seller's demand in period t is $1-F(\theta^\dagger(\lambda_t, p_t))$. We refer to the set of buyers who purchase the good as the seller's *clientele*.

Reputation Transition. Initially, the seller's reputation is $\lambda_0 \in [0,1]$. It evolves dynamically depending on the types of buyers who purchase the good. We assume that reputation satisfies the following two properties. First, reputation is persistent in the sense that, ceteris paribus, a higher reputation today yields a higher reputation tomorrow. Second, the more exclusive the seller's clientele today, the higher her reputation, that is, reputation increases in today's cutoff type. Formally, we assume that reputation evolves according to

$$\lambda_{t+1} \;=\; \phi(\lambda_t, \theta^\dagger(\lambda_t, p_t)), \tag{2}$$

where $\phi : [0,1]^2 \to [0,1]$ is strictly increasing and continuous in both arguments.[9]

Information and Timing. At time t, the seller knows the set of agents who bought the good in periods $0,1,\dots,t-1$ as well as the corresponding prices and reputation levels. The timing of the game in each period t is as follows. First, the seller sets a price p_t. Then, period-t buyers arrive, buy the good or not, and leave the market. Last, reputation updates according to (2).

Histories and Payoffs. Let $p^t = (p_0, \dots, p_{t-1})$ be the history of prices up to time t. For any t and initial reputation λ_0, p^t determines a history of reputation levels $\lambda^t(p^t) = (\lambda_0, \lambda(p^1), \dots \lambda(p^t))$ through (2). The seller's history at the start of period t is given by $h^t = (p^t, \lambda^t(p^t))$ and $h^0 = \lambda_0$. Further, denote the set of all possible histories at the start of period t by $\mathscr{H}^t$.

Fix any history h^t, the seller's continuation payoff from a sequence of prices $(p_s)_{s=t}^\infty$ is

$$\sum_{s=t}^{\infty} \delta^{s-t} p_s (1 - F(\theta^\dagger(\lambda(p^s), p_s)). \tag{3}$$

Strategies and Equilibrium. A (behavioral) pure strategy of the seller is a collection of functions $\rho = (\rho_t)_{t=0}^\infty$, where

$$\begin{aligned} \rho_t \;:\; \mathscr{H}^t &\longrightarrow \mathbb{R}_{\geq 0}, \\ h^t &\longmapsto p_t. \end{aligned}$$

The strategy is Markovian if ρ_t is a function of λ_t only, for all t.

8 In light of this reasoning, our assumption to not model buyers as players is mainly for notational convenience. This is a common assumption in the private reputation literature, see, for example, Board and Meyer-ter-Vehn (2013). Essentially, we only assume that the buyer with the cutoff type purchases the good.

9 We consider a deterministic reputation transition in order to obtain a clean comparison of reputation dynamics.

Definition 1. *A strategy $\rho^* = (\rho_t^*)_{t=0}^{\infty}$ of the seller constitutes*

(i) a Nash equilibrium (NE) if it maximizes (3) at h^0.

(ii) a subgame perfect Nash equilibrium (SPNE) if it maximizes (3) at any h^t, for any t.

(iii) a Markov perfect equilibrium (MPE) if it is Markovian and a subgame perfect equilibrium.

3.3 Equilibrium Existence and Value of Reputation

We start by establishing existence of an equilibrium.

Proposition 1. *There exists a MPE. In any NE the seller's value is $V(\lambda_0)$, where $V(\lambda)$ is the unique solution to*

$$V(\lambda) = \sup_{p \in \mathbb{R}_{\geq 0}} \left\{ p(1 - F(\theta^\dagger(\lambda, p)) + \delta V\left(\phi(\lambda, \theta^\dagger(\lambda, p))\right) \right\}. \tag{4}$$

This and all subsequent proofs are relegated to the Appendix. The seller's problem is to choose an infinite sequence of prices in order to maximize her discounted sum of profits. A strategy of the seller which solves this problem is a Nash equilibrium. In Proposition 1, we show that the value of the seller is characterized by Bellman equation (4). A policy function, corresponding to a solution of the Bellman equation, induces a strategy that is Markovian. Thus, solving the seller's problem through (4) yields a Markov perfect equilibrium. We establish existence of a value function which solves (4) and of a corresponding policy function. As any Markov perfect equilibrium is also a Nash equilibrium, Proposition 1 particularly implies existence of a Nash equilibrium. The proof draws on classical results from the literature on dynamic programming.

We proceed by characterizing the seller's value function in more detail.

Proposition 2. *Reputation is valuable for the seller, that is, $V(\lambda)$ is increasing.*

This result shows that a higher reputation entails a higher value and is thus better for the seller. Specifically, a higher reputation benefits the seller in two ways. First, the buyers' willingness to pay increases in reputation because their utility function is increasing in reputation. Therefore, a seller with a higher reputation can sell to the same group of buyers at a higher price. Second, as reputation is persistent, a higher reputation today implies a higher reputation tomorrow, for a fixed group of buyers. As a result, the first effect carries over to subsequent periods. The proof of Proposition 2 exploits these two effects using a mimicking argument. A high-reputation seller can always induce the same proportion of buyers to purchase the good as a low-reputation seller at higher prices. Consequently, the high-reputation seller can earn higher profits.[10]

[10] The fact that the value function is increasing does not follow from classical sufficient conditions in the literature because the per-period payoff function is not necessarily increasing in reputation.

Moreover, Proposition 2 reveals that there is an intertemporal trade-off between building a reputation for the future and realizing profits in the current period. On the one hand, the seller wants to set a price to increase her next-period reputation in order to get a higher continuation value. Considering only this effect, it is optimal for the seller to set a high price such that only a small exclusive group of high-type buyers purchases the good. This drives up the seller's next-period reputation. On the other hand, the seller wants to set a price that maximizes current-period profits. If the seller sets a high price only a small number of buyers purchase the good. In order to maximize current-period profits, it is better for the seller to set an intermediate price, the static monopoly price, at which a larger set of buyers purchase the good. In turn, however, a larger set of buyers contains a more diverse pool of types. Therefore, setting the static monopoly price comes at the cost of driving down the seller's next-period reputation compared to the case where the seller targets a small, exclusive group of high-type buyers. To maximize overall profits, the seller has to balance these two countervailing effects. Consequently, the seller's optimal price today is higher than the static monopoly price to protect her reputation tomorrow.

In other words, the seller trades off targeting mass, that is a large group of buyers with diverse types, and targeting class, that is a small but exclusive group of buyers with high types.

3.4 Reputation Dynamics

So far, we have shown existence of an equilibrium and that reputation is valuable for the seller in a general model. In this section, we are interested in reputation dynamics. In particular, we study the reputation's long-run behavior and its dynamic evolution over time. To characterize the dynamic properties of the model, we need to impose additional structure which allows us to obtain closed-form solutions for the seller's value and policy function.

Model Adjustments. In the following, we consider the case where types are uniformly distributed, and a buyer's utility from buying the good is linear in λ_t and θ, specifically

$$u(\theta, \lambda_t) \;=\; \theta + \lambda_t. \tag{5}$$

This implies that for a given cutoff $\theta_t^\dagger$ the good's price is $\lambda_t + \theta_t^\dagger$, and demand is $1 - \theta_t^\dagger$. Further, we assume the following reputation transition,

$$\lambda_{t+1} \;=\; \alpha\lambda_t + (1-\alpha)\theta_t^\dagger, \tag{6}$$

where $\alpha \in [0, 1]$. Reputation evolves according to a convex combination of today's reputation and cutoff type. The latter serves as a proxy for the clientele of the current-period.

We interpret α as a measure for the good's physical durability. A good that is more durable is used longer and, hence, also associated longer with its buyers' types. A high value of α corresponds to a good with high durability whose buyers' types are identified with it for a relatively long time. Therefore, new buyers only account for a small fraction of goods in circulation and reputation depends less on their types. In turn, however, their types impact the good's reputation for a long time. Examples for such goods are expensive watches. In contrast, a low value of α represents a good with low durability. It is in circulation for a relatively short period of time and, hence, reputation depends more on the current period's clientele. Fashionable apparel is an example for such a good.

Intuitively, in our model, we keep the size of the group of buyers who determine the good's reputation constant. One could think of a reference group of buyers whose types are associated with the good. This keeps the influence of each buyer on reputation constant over time and simplifies the analysis. Keeping this in mind, we can motivate the specific form of (6) more formally if we imagine that the good becomes unusable with probability $1-\alpha$. The types of buyers in the reference who do not use the good any longer are replaced by those of current-period buyers. Therefore, the fraction $1-\alpha$ of the good's reputation is determined by current-period buyers' types, represented by the proxy $\theta^{\dagger}$, and the fraction α by past-period buyers' types. See also the discussion following Proposition 3.

Value and Policy Function. In this framework, we derive closed-form solutions for the seller's value function and the seller's policy function. We start by establishing a preliminary result.

Lemma 1. *The value function $V(\lambda)$ is continuous in α, δ and λ.*

This result equips us with the properties we need to characterize the value function for the uniform-linear model.

So far, the seller's strategy determines a price for any reputation level. There is a one-to-one relation between the cutoff type and the price. Hence, instead of setting a price, it is convenient to think of the seller's strategy as choosing cutoff types in the following.

Proposition 3.

(i) The value function $V(\lambda)$ is quadratic and continuously differentiable.

(ii) The corresponding unique policy function $\theta^*(\lambda)$ is of the form

$$\theta_1\lambda + \theta_0, \tag{7}$$

for some θ_0 and $\theta_1 \in (-1, 0)$.

This result can be seen from two perspectives. Technically, we derive a closed-form solution for the uniform-linear model which turns out to be linear-quadratic, i.e., the policy function is linear, and the value function is quadratic. Economically, we show that the policy function is decreasing in reputation. In other words a seller with high reputation chooses a lower cutoff type and, thus, sells to a larger group of buyers than a seller with low reputation. To gain some intuition for this result, it is instructive to consider the seller's current-period profits. The effect of $\theta^\dagger$ on current-period profits is twofold. On the one hand, a higher $\theta^\dagger$ corresponds to a higher price. On the other hand, a higher $\theta^\dagger$ corresponds to selling to a smaller group of high-type buyers, that is, to lower demand. Specifically, the price is linearly increasing in $\theta^\dagger$, whereas demand is linearly decreasing in $\theta^\dagger$. At the interior optimum, the seller balances these two opposing effects. Ceteris paribus, a higher reputation increases the price one-to-one but has no effect on demand. Taking the two opposing effects into account, it is profitable for the seller to trade off a fraction of the price increase for an increase in demand by choosing a lower $\theta^\dagger$. This intuition remains valid even when accounting for the effect on the continuation value. The continuation value is increasing in $\theta^\dagger$ because tomorrow's reputation is higher if the seller sells exclusively to a small group of high-type buyers. This strengthens the positive effect that an increase in today's reputation has on today's price through $\theta^\dagger$. Nevertheless, the optimum remains interior and the seller continues to trade off the two opposing effects.

Corollary 1. *There exists a unique SPNE.*

Proposition 1 yields existence of a Markov perfect equilibrium and hence of a Nash equilibrium. By Proposition 3 there exists a unique policy function and therefore a unique Markov perfect equilibrium. Because we only model the seller as a player, the subgame perfect Nash equilibrium is also unique.

In our representation of the reputation transition (6), reputation in the next period depends on the cutoff buyer $\theta^\dagger$ as a summary statistic for demand in the current period. Note that we retain the linear-quadratic structure for the seller's value and policy function if we consider a more general reputation transition where the next period's reputation depends on a linear function of demand or, equivalently, the cutoff buyer. As types are uniformly distributed, this includes, for example, the conditional expectation of the types of buyers who purchase the good. Thus, our qualitative results carry over to these alternative specifications of the reputation transition that depend more generally on demand. For computational convenience, however, we stick to specification (6) in the following.

Long-run Reputation. With the seller's equilibrium strategy at hand which is implicitly determined by her policy function, we analyze in detail how she dynamically manages her reputation over time. We start by showing that reputation converges to some $\hat{\lambda} \in (0, 1)$ in the long run. We refer to $\hat{\lambda}$ also as *long-run reputation* in the following.

Proposition 4. *As $t \to \infty$, reputation converges to a unique $\hat{\lambda}$, for any $\lambda_0 \in [0,1]$.*

(i) For $\alpha < 1$,

$$\hat{\lambda} = \frac{1 + \frac{\delta(1-\alpha)}{1-\delta\alpha}}{3 + \frac{\delta(1-\alpha)}{1-\delta\alpha}}, \tag{8}$$

which is increasing in δ and decreasing in α.

(ii) For $\alpha = 1$, $\hat{\lambda} = \lambda_0$.

For each reputation, the seller trades off the value of reputation tomorrow against demand and, hence, profits today. In order to increase reputation, she has to sell to a more exclusive clientele which decreases her profits today. Targeting a larger group of buyers today decreases her reputation and lowers profits from future periods. From Proposition 4, we learn that there exists a reputation level, $\hat{\lambda}$, where these effects are in balance and the seller does not want to change her reputation. Whereas it is optimal for the seller to dynamically adjust her reputation and clientele in the short run, the seller optimally targets a fixed clientele and has a constant reputation in the long run.

Note that, in contrast to the private reputation literature, reputation is not a short-run phenomenon. Even in the long run, the seller takes into account that her current clientele determines her future reputation.

As second insight, we provide comparative statics results for the long-run reputation $\hat{\lambda}$ with respect to the seller's discount factor δ as well as the good's durability α. Intuitively, as δ increases the seller becomes more patient and values future periods more. In each period, she trades part of her current-period profits off against reputation tomorrow, recall the discussion following Proposition 2. When future periods become more valuable to the seller, she is willing to invest more of her profits today into higher reputation in subsequent periods in order to increase profits from future periods. As a result, the long-run reputation increases.

Furthermore, Proposition 4 reveals that it is optimal for the seller to decrease her reputation over time if she has a high reputation initially. On average, the seller realizes higher profits in early periods at the cost of driving down her reputation over time. Conversely, if the seller starts with a low reputation, it is optimal for her to increase reputation over time at the cost of lower profits in early periods. In the long run, however, reputation is stable. Intuitively, above $\hat{\lambda}$, the benefits of higher profits in early periods outweigh the costs of decreasing reputation in the long run; below $\hat{\lambda}$ the benefits of higher reputation in the long run outweigh the costs of lower profits in early periods. At the long-run reputation, the trade-off between realizing current-period profits and changing reputation is balanced. As α increases, the good becomes more durable and reputation more persistent. From the seller's perspective, she has to give up more demand to increase reputation and loses less reputation when increasing demand. Therefore, the costs of increasing current-period profits go down, whereas the benefits of higher current-period profits are unaffected by the change in α. Consequently, the level where these costs and benefits are in balance is lower, that is, long-run reputation $\hat{\lambda}$ decreases in α.

Although it seems that the conditions imposed on the general model, outlined in Section 3.2, are not sufficient to guarantee existence of a long-run reputation, we can establish existence in frameworks that are more general than our linear-quadratic setup. See, for example, Jose Alexandre Scheinkman (1976), Araujo and Jose A Scheinkman (1977), McKenzie (1986), Stokey et al. (1989) for sufficient conditions.

In the following, we say that the seller has *good reputation* in period t if $\lambda_t > \hat{\lambda}$ and *bad reputation* in period t if $\lambda_t < \hat{\lambda}$. Further, we say that the seller *milks* reputation in period t if she targets a clientele which decreases her next-period reputation, that is $\theta_t^\dagger < \lambda_t$. Conversely, the seller *builds* reputation if she targets a clientele which increases her next-period reputation, that is $\theta_t^\dagger > \lambda_t$.

In the next step, we are interested how the underlying characteristics of the market affect the convergence dynamics towards the long-run reputation. The market is characterized by the good's durability α. Therefore, we analyze the effect of the good's durability on how reputation converges. The following definition turns out to be useful.

Definition 2. *A sequence of reputations* $\{\lambda_t\}_{t=0}^{\infty}$

(i) oscillates towards long-run reputation $\hat{\lambda}$ *if* $\lambda_t \leq (\geq)\hat{\lambda}$ *implies that* $\lambda_{t+1} \geq (\leq)\hat{\lambda}$, *for all* t, *and* $|\lambda_t - \hat{\lambda}| \downarrow 0$.

(ii) stagnates at long-run reputation $\hat{\lambda}$ *if* $\lambda_t = \hat{\lambda}$, *for all* $t \geq 1$.

(iii) converges smoothly to long-run reputation $\hat{\lambda}$ *if* $\lambda_t \leq \lambda_{t+1} \leq \hat{\lambda}$, *for all* t, *or* $\lambda_t \geq \lambda_{t+1} \geq \hat{\lambda}$, *for all* t, *and* $|\lambda_t - \hat{\lambda}| \downarrow 0$.

Intuitively, if reputation oscillates towards $\hat{\lambda}$, a period of good reputation is followed by a period of bad reputation and vice versa. The magnitude of the fluctuation between good and bad reputation decreases over time. If reputation stagnates at $\hat{\lambda}$, it jumps to the long-run reputation $\hat{\lambda}$ in the first period and remains constant from then on. There are two scenarios in which reputation converges smoothly to $\hat{\lambda}$. First, reputation is good in every period and deteriorates over time to $\hat{\lambda}$, i.e., reputation converges to $\hat{\lambda}$ monotonically from above. Second, reputation is bad in all periods and increases over time to $\hat{\lambda}$, i.e., reputation converges to $\hat{\lambda}$ monotonically from below.

For a better understanding of the dynamics, we study two benchmark cases. In the first case, the good is everlasting, i.e. $\alpha = 1$, and in the second case the good is most ephemeral, i.e. $\alpha = 0$.

Benchmark: Everlasting Good ($\alpha = 1$). In the extreme case when the good is everlasting, the good is always associated with the initial buyers' types. Current buyers do not affect reputation. Intuitively, as the good is everlasting, the reference group of buyers who determine the seller's reputation is unchanged over time. Consequently, reputation remains at its initial value λ_0 in each period. From the seller's perspective, the dynamic trade-off disappears. Her reputation tomorrow is not influenced by her clientele today and, hence, she solves the same static problem in each period.

Lemma 2. *Suppose $\alpha = 1$. For any t, it is optimal for the seller to set*

$$\theta_t^{\dagger} \;=\; \frac{1-\lambda_0}{2},$$

and her profits are

$$\frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2.$$

As reputation remains constant at λ_0, this is an extreme case of smooth convergence to $\hat{\lambda}$ in the sense of Definition 2.

Benchmark: Most Ephemeral Good ($\alpha = 0$). Contrary to the previous case, the good is most ephemeral such that buyers use their good for one period only. Therefore, the group of buyers who possess the good consists exclusively of current-period buyers, and, consequently, the good is only associated with their types. The seller's reputation tomorrow is independent of her reputation today. Hence, her choice of today's price not only determines her current-period profits but also fully determines her reputation tomorrow. Consequently, in this benchmark, reputation reacts most sensitively to adjustments in the seller's clientele.

Lemma 3. *Suppose $\alpha = 0$. The seller's policy function is*

$$\theta^*(\lambda) \;=\; -\frac{\lambda}{1+\sqrt{1-\delta}} + \frac{1+\delta}{1+\delta+\sqrt{1-\delta}},$$

and, as $t \to \infty$, her reputation oscillates towards

$$\hat{\lambda} \;=\; \frac{1+\delta}{3+\delta}.$$

See Figure 1 for a graphical illustration of the resulting dynamics of reputation, prices, and demand. We observe that reputation alternates between periods of high reputation and periods of low reputation if the good is most ephemeral. Also, the corresponding sequences of prices and demand fluctuate substantially. Over time, the amplitude of the fluctuations decreases and reputation oscillates towards the long-run reputation.

Intermediate Durability: $\alpha \in (0, 1)$. For goods of intermediate durability, the set of buyers who possess the good, and whose types are thus associated with the good,

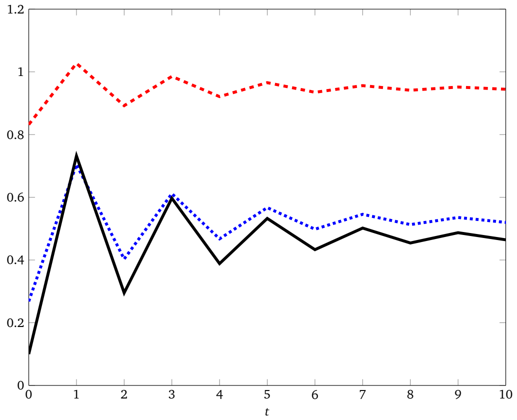

(a)

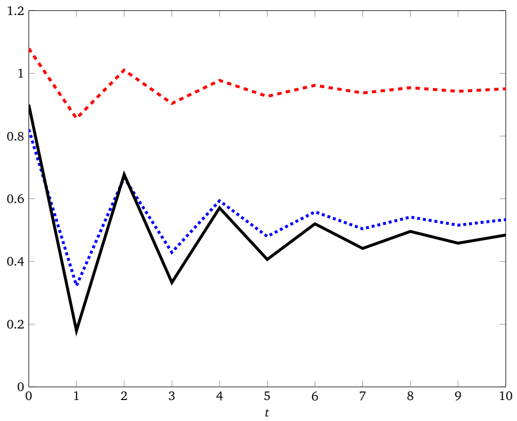

(b)

Figure 1. Reputation (solid black line), prices (dashed red line), and demand (dotted blue line) for $\alpha = 0$ and $\delta = 0.8$ in $t \in \{0, \ldots, 10\}$ with $\lambda_0 = 0.1$ (a), and $\lambda_0 = 0.9$ (b).

consists of current-period buyers ($\alpha < 1$) and buyers from previous periods ($\alpha > 0$). Therefore, both influence the seller's next-period reputation. In this case, we obtain the following result.

Proposition 5. *There exists a cutoff $\bar{\alpha} \in (0, 1)$ such that reputation*

(i) converges smoothly to $\hat{\lambda}$ if $\alpha > \bar{\alpha}$.

(ii) stagnates at $\hat{\lambda}$ if $\alpha = \bar{\alpha}$.

(iii) oscillates towards $\hat{\lambda}$ if $\alpha < \bar{\alpha}$.

Proposition 5 describes how the seller manages her reputation over time, depending on the durability of the good. From Proposition 4 we know that the seller's reputation converges to $\hat{\lambda}$, starting from any initial reputation λ_0. The seller manages her reputation by setting prices and thereby targeting a particular clientele which in turn influences her reputation in subsequent periods. Proposition 5 shows that, depending on the

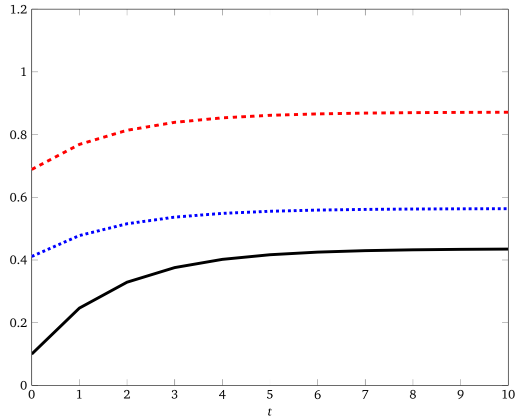

(a)

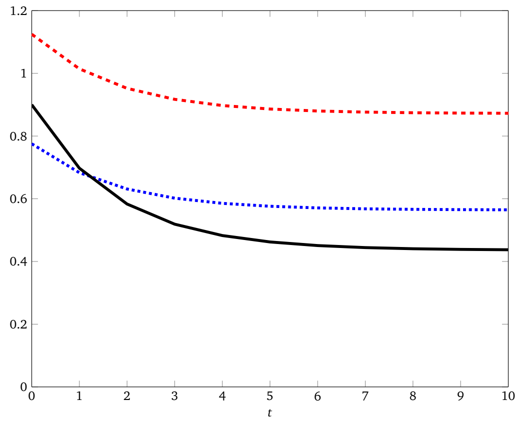

(b)

Figure 2. Reputation (solid black line), prices (dashed red line), and demand (dotted blue line) for $\alpha = 0.7$ and $\delta = 0.8$ in $t \in \{0, \ldots, 10\}$ with $\lambda_0 = 0.1$ (a), and $\lambda_0 = 0.9$ (b).

durability of the good, there are three ways in which the seller's reputation converges to the long-run reputation. These differ substantially in the convergence behavior of reputation and prices.

First, consider the case where the good is relatively durable, $\alpha > \bar{\alpha}$, and the seller's initial reputation is good, $\lambda_0 > \hat{\lambda}$. Initially, the seller sells to a large group of buyers with a diverse set of types by setting a price which is high in absolute terms but comparatively low given her high reputation. As a result, reputation decreases. Because durability is high, it takes time until the low-type buyers constitute a significant proportion of the set of buyers who possess the good and affect the reputation negatively. Therefore, the decline in reputation is relatively small. As time passes, the seller gradually targets less buyers overall, however, she continues to sell to a large group of buyers. Consequently, reputation declines at a decreasing rate. The effect on prices is twofold. On the one hand decreasing demand corresponds to increasing prices. On the other hand, as reputation decreases, the seller has to lower prices because the buyers' willingness to pay decreases.

Overall, we observe that prices decrease smoothly over time which corresponds to demand decreasing slower than reputation. See Figure 2b for a graphical representation. Intuitively, if the good's durability is high and the initial reputation of the seller is good, the seller gradually milks her reputation in every period. That is, she exploits that the good is mostly associated with her initially exclusive clientele so that selling to a broader, less exclusive, clientele only gradually deteriorates her reputation. As a result she realizes high profits in each period at the expense of slowly driving down next period's reputation.

If the good is relatively durable, $\alpha > \hat{\alpha}$, but the seller's reputation is bad initially $\lambda_0 < \hat{\lambda}$, we observe the mirror image of this behavior. See Figure 2a for a graphical representation. At first, the seller targets a small group of high-type buyers, and demand is low. The corresponding price is comparatively high given her low reputation, and, thus, reputation increases. As the good is durable, however, it takes time until the good is associated with these high-type buyers. Therefore, reputation only increases slowly. Over time, the seller gradually expands the set of buyers to which she sells but overall she continues to target a small, exclusive clientele. As a result, reputation increases at a decreasing rate. Analogously to the last case, prices increase over time because demand increases slower than reputation. Intuitively, even though the seller targets a larger group of buyers which would require setting a lower price, she can charge a higher price because the buyer's willingness to pay increases as her reputation improves. In this case, the seller builds reputation over time, i.e, she sacrifices profits in each period by selling only to an exclusive clientele in order to build reputation for future periods.

Altogether, the dynamics are relatively smooth and monotone if the good's durability is high. If the seller's initial reputation is high, reputation deteriorates monotonically to the long-run reputation because the seller milks her reputation in every period. If her initial reputation is low, reputation improves monotonically to the long-run reputation as the seller builds reputation in every period.

Second, we analyze the case when the good has low durability, $\alpha < \hat{\alpha}$. Consider a period in which the seller has good reputation. She sets a comparatively low price given her high reputation such that demand is high, i.e., a large group of buyers with a diverse set of types purchases the good. In other words, the seller milks her high reputation and consequently reputation deteriorates. This is reminiscent of the behavior of a seller with good reputation when the good's durability is high. As the good's durability is low, however, current-period buyers represent a significant proportion of the buyers who possess the good in this case. Thus, the types of recent buyers, in particular of current-period buyers, have a significant impact on the seller's reputation in the next period. Consequently, reputation does not decrease gradually as before but drops below the long-run reputation. As a result, the seller has a bad reputation in the next period. Now consider a period in which the seller has bad reputation. The seller sets a comparatively high price and sells exclusively to a small group of high-type buyers, i.e., she builds reputation. As durability is low, these high-type buyers make up a significant part of the buyers who possess the good. As a result, reputation jumps above the long-run reputation. These

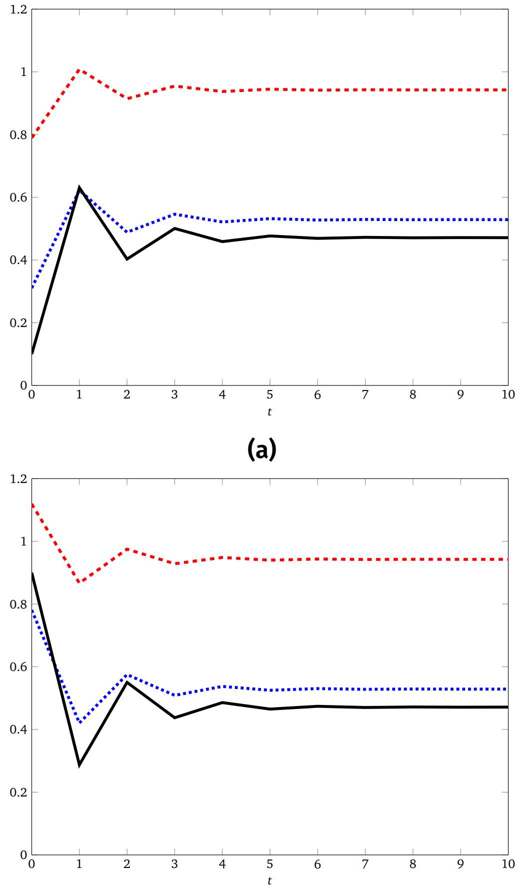

(a)

(b)

Figure 3. Reputation (solid black line), prices (dashed red line), and demand (dotted blue line) for $\alpha = 0.1$ and $\delta = 0.8$ in $t \in \{0, \ldots, 10\}$ with $\lambda_0 = 0.1$ (a), and $\lambda_0 = 0.9$ (b).

observations imply that a period of good reputation is followed by a period of bad reputation and vice versa. As reputation converges to the long-run reputation, the amplitude of the alternations in reputation decreases over time. Analogously, prices alternate with decreasing amplitude. See Figure 3 for a graphical representation.

All in all, if the good's durability is low, reputation, demand, and prices fluctuate substantially. Periods of good reputation in which the seller milks her reputation alternate with periods of bad reputation in which the seller builds reputation. Intuitively, milking good reputation is profitable because the seller increases demand at a comparatively high price. Therefore, her current-period profits are particularly high. Conversely, building reputation in periods of bad reputation is less costly because, in any case, the seller's price is comparatively low due to her bad reputation. Hence, the seller has to give up relatively little in terms of profits today to increase reputation tomorrow.

Third, in the knife-edge case where the durability of the good coincides with $\bar{\alpha}$, reputation jumps immediately to the long-run reputation which balances the trade-off between current profits and the future benefits from a good reputation in the long run. Consequently, reputation remains constant in all subsequent periods.

We further investigate the effect of the good's durability on the reputation dynamics. Proposition 5 suggests that there is a monotone relationship between the durability of the good and the volatility of reputation dynamics. Intuitively, the lower the durability of the good, the more reputation fluctuates around the long-run reputation. This conjecture raises the issue of finding an appropriate measure for volatility to compare the volatility of two reputation sequences. In the following, we propose such a measure and formalize this intuitive conjecture.

It is instructive to make the dependence of reputation on α explicit. Fix a distance $h \in \mathbb{R}$ from long-run reputation $\hat{\lambda}(\alpha)$ and consider the reputation

$$\lambda(\alpha) \ := \ \hat{\lambda}(\alpha) + h.$$

Let $\lambda'(\alpha)$ be the next-period reputation when following the optimal policy if the current-period reputation is $\lambda(\alpha)$, i.e.,

$$\lambda'(\alpha) \ := \ \alpha\lambda(\alpha) + (1-\alpha)\theta^*(\lambda(\alpha)).$$

Intuitively, we define reputation relative to the long-run reputation $\hat{\lambda}(\alpha)$ to account for changes in $\hat{\lambda}(\alpha)$ as the durability of the good α changes.[11] Our measure for volatility is the distance between between $\lambda'(\alpha)$ and $\lambda(\alpha)$. The following result formalizes our conjecture:

Proposition 6. $|\lambda'(\alpha) - \lambda(\alpha)|$ *is decreasing in* α.

We show that the higher the durability of the good, the smaller the volatility of the seller's reputation. In other words, the higher α, the smaller the distance between two subsequent reputation levels. In this sense, there is a continuous transition between the three patterns in which reputation converges to the long-run reputation $\hat{\lambda}$, outlined in Proposition 5. This continuous relationship, gives rise to the question where the substantially different convergence behavior of the three cases comes from. Technically, it is the consequence of the fact that the distance between two subsequent reputation levels is increasing in α combined with the fact that the policy function is decreasing in λ and has a fixed point at $\hat{\lambda}$. For illustration, consider a seller who has a good reputation initially. If α is high, the distance between two subsequent reputation levels is relatively small. Thus, if the current-period reputation is larger than $\hat{\lambda}$, the next-period's reputation is larger than $\hat{\lambda}$ as well. Therefore, the seller milks reputation in both periods. As α decreases the distance between two subsequent reputation levels increases. If α is sufficiently low, the next-period reputation is below $\hat{\lambda}$. Because the policy function

[11] See also the discussion following Proposition 6.

is decreasing and has a fixed point at $\hat{\lambda}$, the seller switches from milking to building reputation in the next period.

To account for the fact that $\hat{\lambda}$ depends on α, we use the distance of two subsequent reputation levels defined relative to $\hat{\lambda}$ as our volatility measure. This is crucial because if we do not measure reputation relative to $\hat{\lambda}$, we ignore the dependence of $\hat{\lambda}$ on α. Then, it can happen for reputations close to $\hat{\lambda}$ that the decrease in $\hat{\lambda}$ outweighs the decrease in fluctuations around $\hat{\lambda}$ such that the distance between two subsequent reputation levels is increasing in α. As we are interested in the way reputation transitions to and fluctuates around $\hat{\lambda}$, however, we measure distances relative to $\hat{\lambda}$ in (the proof of) Proposition 6.

Intuitively, Proposition 5 and 6 predict that for goods with high durability, one should observe less fluctuations in reputation, clientele, and hence prices. The good is longer associated with its buyers' types and the seller manages her reputation such that it only gradually approaches the long-run reputation. In other words, in markets for relatively durable goods in which buyers care about the seller's past clientele, we should rarely observe price discounts. As mentioned in the introduction, this finding is consistent with the observation that the prices and reputation of the Swiss watchmaker *Rolex* have gradually increased over the last sixty years.

In contrast, for producers of goods with low durability our model predicts significant fluctuations in reputation, clientele, and prices. For these goods, reputation is mostly determined by current-period buyers. Hence, if a seller with a good reputation milks reputation by selling to a broader clientele, she incurs a severe negative effect on her next-period reputation. As outlined in the introduction, this provides a rationale for substantial changes in clientele and reputation of fashion labels such as *Abercrombie & Fitch* and *Ed Hardy*.

3.5 Planned Obsolescence

After studying the effect of the good's durability on the seller's reputation dynamics, we analyze how durability affects the profit of the seller. This analysis sheds light on planned obsolescence. Does a good with low durability yield higher profits for the seller than a good with high durability? The classical explanation of planned obsolescence is that it is optimal for a monopolist to produce a good with uneconomically short life to increase demand in future periods, see Bulow (1986). In our model, durability does not affect the size of the market in future periods. Nevertheless, we argue that a seller with low initial reputation can obtain higher profits if the good is less durable. Intuitively, if the seller has a low reputation, she can only charge comparatively low prices and hence, wants to build reputation over time. She improves her reputation to generate higher profits in the future by selling to an exclusive clientele at the cost of current-period profits. If the durability of the good is comparatively low, the seller has to sacrifice less current-period profits to improve reputation because the new high types quickly replace the initial low types in the seller's clientele. As a result, the seller can obtain higher profits if the good is less durable. Put differently, reputation is more sensitive to changes in the seller's clientele if the durability of the good is low. While this makes it more difficult to maintain a high initial reputation, it is desirable for a seller with low initial reputation. Thus, our model reveals additional incentives for planned obsolescence from a reputational perspective which complements the classical explanation.

For the following result, it is convenient to make the dependence of $V(\lambda)$ on α explicit and write $V(\lambda, \alpha)$.

Proposition 7. *For any* $\alpha \in (0,1)$ *there exists a cutoff* $\bar{\lambda}_0(\alpha) \in (0,1)$ *such that*

(i) *for all* $\lambda_0 \geq \bar{\lambda}_0(\alpha)$ *it holds that* $V(\lambda_0, 1) \geq V(\lambda_0, \alpha)$,

(ii) *for all* $\lambda_0 \leq \bar{\lambda}_0(\alpha)$ *it holds that* $V(\lambda_0, 1) \leq V(\lambda_0, \alpha)$.

In Proposition 7 we compare the seller's profits from an everlasting good and a good with durability $\alpha \in (0,1)$. We show that there exists a cutoff $\bar{\lambda}_0(\alpha)$ such that a seller with initial reputation below $\bar{\lambda}_0(\alpha)$ can obtain higher profits if her good is not everlasting. Conversely, an everlasting good yields higher profits than the less durable good for a seller with initial reputation above $\bar{\lambda}_0$.

3.6 Conclusion

In a setting where buyers care about who else consumes the good, we study how a seller dynamically manages reputation which is determined by her clientele. In contrast to the seller's private reputation which is the market's belief about the seller's private type, this public reputation captures the consumption externalities that buyers impose on each other. In a general setup, we show that public reputation is valuable for the seller and that there exists a trade-off between current-period profits and a higher reputation in subsequent periods. In a linear-uniform specification of the model, we find that reputation converges to a stable long-run reputation. In line with anecdotal evidence from various industries, the convergence path depends strongly on the good's characteristics. As the examples of *Ed Hardy* and *Abercrombie & Fitch* illustrate, reputation of less durable goods such as fashionable apparel fluctuates substantially. In contrast, the reputation of goods with higher durability such as luxury watches evolves monotonically as the case of the Swiss watchmaker *Rolex* exemplifies.

In future research, we think it is interesting to analyze reputation dynamics when other market characteristics, in addition to the good's durability, influence the reputation transition. In a similar vein, studying the implications on reputation dynamics of competition between multiple sellers could shed light on how companies with very different reputation coexist in one market. For example, in the automobile industry it seems that the Italian car manufacturer *Ferrari* and the Romanian car manufacturer

Dacia not only differ with respect to quality-related private reputation but also with respect to what their cars symbolize to others.

Although our results about price and reputation dynamics seem to match observations in several markets, our model considers public reputation as the only determinant of prices. Testing empirically how strongly pricing is influenced by public, clientele driven reputation as opposed to private, quality driven reputation would provide interesting insights.

Finally, we would like to highlight that the public reputation analyzed in this chapter is present in several other applications as well, such as in the job market. A professor who considers whether to join a university, for example, takes into account the university's reputation, which depends on the academic standing of its past and current faculty members. In turn, she also influences the university's future reputation. Alternatively, young academics oftentimes work as an intern for a small salary in prestigious companies to improve their resume. Then again, the company's prestige is influenced by the quality of their employees. We believe that analyzing the implications of public reputation in this and further applications are promising avenues for future research.

References

Alabi, Mo (September 30, 2013): "Ed Hardy: From art to infamy and back again." *CNN*, Retrieved from http://edition.cnn.com. [77]

Amaldoss, Wilfred and Sanjay Jain (2005a): "Conspicuous consumption and sophisticated thinking." *Management Science*, 51 (10), 1449–1466. [77]

Amaldoss, Wilfred and Sanjay Jain (2005b): "Pricing of conspicuous goods: A competitive analysis of social effects." *Journal of Marketing Research*, 42 (1), 30–42. [78]

Araujo, Aloisio and Jose A Scheinkman (1977): "Smoothness, comparative dynamics, and the turnpike property." *Econometrica*, 601–620. [86]

Bagwell, Laurie Simon and B Douglas Bernheim (1996): "Veblen effects in a theory of conspicuous consumption." *The American Economic Review*, 349–373. [77]

Bar-Isaac, Heski and Steven Tadelis (2008): "Seller reputation." *Foundations and Trends in Microeconomics*, 4 (4), 273–351. [77]

Becker, Gary S (1991): "A note on restaurant pricing and other examples of social influences on price." *Journal of Political Economy*, 1109–1116. [77]

Becker, Gary S and Kevin M Murphy (1993): "A simple theory of advertising as a good or bad." *The Quarterly Journal of Economics*, 941–964. [77]

Bensaid, Bernard and Jean-Philippe Lesne (1996): "Dynamic monopoly pricing with network externalities." *International Journal of Industrial Organization*, 14 (6), 837–855. [78]

Board, Simon (2009): "Monopolistic group design with peer effects." *Theoretical Economics*, 4 (1), 89–125. [79]

Board, Simon and Moritz Meyer-ter-Vehn (2013): "Reputation for quality." *Econometrica*, 81 (6), 2381–2462. [77, 80]

Bredan, David (January 1, 2015): "How and why Rolex prices have increased over time." *Business Insider*, Retrieved from http://www.businessinsider.com. [76]

Bulow, Jeremy (1986): "An economic theory of planned obsolescence." *The Quarterly Journal of Economics*, 729–750. [93]

Cripps, Martin W, George J Mailath, and Larry Samuelson (2004): "Imperfect monitoring and impermanent reputations." *Econometrica*, 72 (2), 407–432. [76, 77]

Cripps, Martin W, George J Mailath, and Larry Samuelson (2007): "Disappearing private reputations in long-run relationships." *Journal of Economic Theory*, 134 (1), 287–316. [76]

Denizet-Lewis, Benoit (January 24, 2006): "The man behind Abercrombie & Fitch." *Salon*, Retrieved from http://www.salon.com. [73]

Dhebar, Anirudh and Shmuel S Oren (1985): "Optimal dynamic pricing for expanding networks." *Marketing Science*, 4 (4), 336–351. [78]

Dilmé, Francesc (2016): "Reputation building through costly adjustment." *mimeo*. [77]

Frommeyer, Tim and Benjamin Schickner (2016): "Target mass or class? Dynamic reputation management with heterogeneous consumption externalities." *mimeo*. [73]

Gabszewicz, Jean J and Filomena Garcia (2008): "A note on expanding networks and monopoly pricing." *Economics Letters*, 98 (1), 9–15. [78]

Grilo, Isabel, Oz Shy, and Jacques-François Thisse (2001): "Price competition when consumer behavior is characterized by conformity or vanity." *Journal of Public Economics*, 80 (3), 385–408. [78]

Hashimoto, Kaito and Nobuo Matsubayashi (2014): "A note on dynamic monopoly pricing under consumption externalities." *Economics Letters*, 124 (1), 1–8. [78]

Holmström, Bengt (1999): "Managerial incentive problems: A dynamic perspective." *The Review of Economic Studies*, 66 (1), 169–182. [76, 77]

Katz, Michael L and Carl Shapiro (1985): "Network externalities, competition, and compatibility." *The American Economic Review*, 75 (3), 424–440. [78]

Kreps, David M, Paul Milgrom, John Roberts, and Robert Wilson (1982): "Rational cooperation in the finitely-repeated prisoners' dilemma." Tech. rep. DTIC Document. [77]

Lee, In Ho and Robin Mason (2001): "Market structure in congestible markets." *European Economic Review*, 45 (4–6). 15th Annual Congress of the European Economic Association, 809–818. [78]

Lee, Jihong and Qingmin Liu (2013): "Gambling reputation: Repeated bargaining with outside options." *Econometrica*, 81 (4), 1601–1672. [77]

Leibenstein, Harvey (1950): "Bandwagon, snob, and Veblen effects in the theory of consumers' demand." *The Quarterly Journal of Economics*, 183–207. [77]

Lindner, Roland and Julia Löhr (August 31, 2015): "Der Niedergang der amerikanischen Modeketten." *Frankfurter Allgemeine Zeitung*, Retrieved from http://www.faz.net. [76, 77]

Liu, Qingmin (2011): "Information acquisition and reputation dynamics." *The Review of Economic Studies*, 78 (4), 1400–1425. [77]

Mailath, George J and Larry Samuelson (2001): "Who wants a good reputation?" *The Review of Economic Studies*, 68 (2), 415–441. [77]

McKenzie, Lionel W (1986): "Optimal economic growth, turnpike theorems and comparative dynamics." *Handbook of Mathematical Economics*, 3, 1281–1355. [86]

Pesendorfer, Wolfgang (1995): "Design innovation and fashion cycles." *The American Economic Review*, 771–792. [78]

Rayo, Luis (2013): "Monopolistic signal provision." *The BE Journal of Theoretical Economics*, 13 (1), 27–58. [78]

Scheinkman, Jose Alexandre (1976): "On optimal steady states of n-sector growth models when utility is discounted." *Journal of Economic Theory*, 12 (1), 11–30. [86]

Stock, Axel and Subramanian Balachander (2005): "The making of a "hot product": A signaling explanation of marketers' scarcity strategy." *Management Science*, 51 (8), 1181–1192. [79]

Stokey, Nancy L, Robert E Lucas Jr, and EC Prescott (1989): *Recursive Methods in Economic Dynamics*. [86, 97]

Tadelis, Steven (1999): "What's in a name? Reputation as a tradeable asset." *The American Economic Review*, 89 (3), 548–563. [77]

Veblen, Thorstein (1899): *The theory of the leisure class: An economic study of institutions*. London: Unwin Books. [77]

Appendix 3.A Appendix: Proofs

Proof of Proposition 1. We apply Theorems 4.2 - 4.6 from Stokey et al. (1989). Theorems 4.2 - 4.5 establish that we can study the seller's problem of choosing an infinite sequence of prices by means of the Bellman equation (4). Theorem 4.6 yields existence of a value function, $V(\lambda)$, which solves (4) as well as non-emptiness of the corresponding optimal policy correspondence. Taken together, these results yield existence of equilibrium. To employ Theorems 4.2 - 4.6, we verify that Assumptions 4.3 - 4.4 from Stokey et al. (1989) are satisfied in our model. We start by rewriting (4) to match their notational convention:

$$V(\lambda) = \sup_{\lambda' \in [\phi(\lambda,0),\phi(\lambda,1)]} \left\{ \left(1 - F(\phi^{-1}(\lambda,\lambda'))\right) u(\phi^{-1}(\lambda,\lambda'),\lambda) + \delta V(\lambda') \right\}, \tag{9}$$

where $\phi^{-1}(\lambda,\lambda')$ denotes the inverse of ϕ with respect to its second argument with domain $[0,1] \times [\phi(\lambda,0),\phi(\lambda,1)]$. Consider Assumption 4.3. First, note that the set of all possible states, $[0,1]$, is convex. Second, the correspondence which maps the initial state λ into the set of feasible next-period states, $[\phi(\lambda,0),\phi(\lambda,1)]$, is non-empty, compact-valued, and continuous because ϕ is continuous and strictly increasing in its second argument. To verify Assumption 4.4, we consider the per-period return function

$$\left(1 - F(\phi^{-1}(\lambda,\lambda'))\right) u(\phi^{-1}(\lambda,\lambda'),\lambda) \tag{10}$$

with domain $\{(\lambda,\lambda') \mid (\lambda,\lambda') \in [0,1] \times [\phi(\lambda,0),\phi(\lambda,1)]\}$. Because $[0,1]$ is a compact set and F, ϕ, and u are continuous, (10) is continuous as well. Together with the compactness of the domain, this observation also implies that (10) is bounded. Thus, Assumptions 4.3 - 4.4 hold, and we can apply Theorems 4.2 - 4.6. □

Proof of Proposition 2. Consider $\tilde{\lambda}$, $\lambda \in [0,1]$ such that $\tilde{\lambda} \geq \lambda$. We argue that $V(\tilde{\lambda}) \geq V(\lambda)$ by showing that for any strategy of the seller with initial reputation λ (henceforth seller$_\lambda$) we can find a strategy of the seller with initial reputation $\tilde{\lambda}$ (henceforth seller$_{\tilde{\lambda}}$) which gives her weakly higher profits. Any strategy of the seller induces a sequence of prices. Because u is strictly increasing, there is a one-to-one relationship between prices and cutoffs. For this proof, it is instructive to consider sequences of cutoffs rather than sequences of prices. Let $\{\theta_s^\dagger\}_{s=0}^\infty$ be an arbitrary sequence of cutoffs chosen by seller$_\lambda$. Note that seller$_{\tilde{\lambda}}$ can choose prices that induce the same sequence of cutoffs because buyers' utility is quasilinear in the price. If seller$_{\tilde{\lambda}}$ chooses the same sequence of cutoffs $\{\tilde{\theta}_s\}_{s=0}^\infty$ such that $\tilde{\theta}_s^\dagger = \theta_s^\dagger$:

(i) Seller$_{\tilde{\lambda}}$ sells to the same (mass of) buyers as seller$_{\lambda}$ in each period.

(ii) The price that seller$_{\tilde{\lambda}}$ charges in the first period, $u(\theta_0^{\dagger}, \tilde{\lambda})$, is higher than the price that seller$_{\lambda}$ charges in the first period, $u(\theta_0^{\dagger}, \lambda)$, because u is increasing in the second argument.

(iii) As ϕ is increasing in its first argument, we have $\phi(\tilde{\lambda}, \theta_0^{\dagger}) \geq \phi(\lambda, \theta_0^{\dagger})$. Therefore, seller$_{\tilde{\lambda}}$'s next-period reputation is higher than seller$_{\lambda}$'s next-period reputation. Using the same argument as in (ii), we deduce that seller$_{\tilde{\lambda}}$ also sets higher prices in the future.

Taking (i) to (iii) together, we see that $\{\tilde{\theta}_s^{\dagger}\}_{s=0}^{\infty}$ guarantees seller$_{\tilde{\lambda}}$ a higher profit than $\{\theta_s^{\dagger}\}_{s=0}^{\infty}$ guarantees seller$_{\lambda}$. Furthermore, by Proposition 1, there exists a Markovian strategy that yields seller$_{\tilde{\lambda}}$ a weakly higher profit than $\{\tilde{\theta}_s^{\dagger}\}_{s=0}^{\infty}$. We conclude that $V(\tilde{\lambda}) \geq V(\lambda)$. □

Proof of Lemma 1. First, we adjust the Bellman equation, (4), to account for the uniform-linear model specification. Also, as in the proof of Proposition 1, it is convenient to let the seller choose cutoffs directly instead of prices. With linear utility, setting a price p_t at reputation λ_t implies that the cutoff is

$$\theta^{\dagger}(\lambda_t, p_t) \;=\; p_t - \lambda_t.$$

Thus, we can express the price as a function of the cutoff,

$$p_t\left(\lambda_t, \theta^{\dagger}\right) \;=\; \theta^{\dagger} + \lambda_t.$$

For notational convenience, we abbreviate $\theta^{\dagger}(\lambda, p)$ with $\theta^{\dagger}$ in the following. The uniform distribution of types implies that $F\left(\theta^{\dagger}\right) = \theta^{\dagger}$. Therefore, we can rewrite (4) to

$$V(\lambda) \;=\; \max_{\theta^{\dagger}\in[0,1]} \left\{\left(\theta^{\dagger} + \lambda\right)\left(1 - \theta^{\dagger}\right) + \delta V\left(\alpha\lambda + (1-\alpha)\theta^{\dagger}\right)\right\}, \tag{11}$$

where the seller maximizes with respect to $\theta^{\dagger}$ instead of p.

By Proposition 1, $V(\lambda)$ is the fixed point of the Bellman operator. We now prove that the Bellman operator maps the (closed) set of functions which are continuous in α, δ, and λ (henceforth denoted by C) into itself. Thus, the fixed point of the Bellman operator must lie within that closed set and, hence, $V(\lambda)$ is continuous in α, δ, and λ.
To show that the Bellman operator maps C into itself, consider an arbitrary $g \in C$. Applying the Bellman operator, we obtain

$$\max_{\theta^{\dagger}\in[0,1]} \left\{\left(\theta^{\dagger} + \lambda\right)\left(1 - \theta^{\dagger}\right) + \delta g\left(\alpha\lambda + (1-\alpha)\theta^{\dagger}\right)\right\} \tag{12}$$

We use Berge's maximum theorem to deduce that (12) is a continuous function of α, δ, and λ. First, note that for any λ the set of feasible choices is $[0, 1]$ which is compact and continuous as a correspondence of α, δ, and λ because it is independent of α, δ, and λ. Second, note that

$$\left(\theta^\dagger + \lambda\right)\left(1 - \theta^\dagger\right) + \delta g\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right)$$

is continuous in α, δ, λ, and $\theta^\dagger$. Thus, by the maximum theorem, (12) is an element of C. Consequently, the Bellman operator maps C into itself and $V \in C$. □

Proof of Proposition 3. We consider the case where $\alpha < 1$. For $\alpha = 1$ see Lemma 2. Recall that the Bellman equation is given by (11).

(i) *Step 1: We solve the Bellman equation with a guess and verify approach.*

Consider the first-order derivative of the right-hand side of the Bellman equation with respect to $\theta^\dagger$

$$1 - \theta^\dagger - \left(\theta^\dagger + \lambda\right) + \delta(1-\alpha)V'\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right), \tag{13}$$

and the first-order derivative of $V(\lambda)$ with respect to λ

$$V'(\lambda) \;=\; 1 - \theta^\dagger + \delta\alpha V'\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right). \tag{14}$$

We guess that the value function is quadratic, that is,

$$V(\lambda) \;=\; v_2\lambda^2 + v_1\lambda + v_0.$$

This implies for the first-order derivative

$$V'(\lambda) \;=\; 2v_2\lambda + v_1. \tag{15}$$

In order to verify the guess, we consider the first-order condition of the Bellman equation by setting (13) equal to zero,

$$1 - \theta^\dagger - \left(\theta^\dagger + \lambda\right) + \delta(1-\alpha)V'\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right) \;=\; 0. \tag{16}$$

Using (15), we can rewrite (16) to

$$1 - \theta^\dagger - \left(\theta^\dagger + \lambda\right) + \delta(1-\alpha)\left(2v_2\left(\alpha\lambda + (1-\alpha)\theta^\dagger\right) + v_1\right) \;=\; 0.$$

Solving for $\theta^\dagger$ yields

$$\theta^\dagger \;=\; \frac{1 + \delta(1-\alpha)v_1 - \lambda(1 - \delta(1-\alpha)(2v_2\alpha))}{2(1 - \delta(1-\alpha)^2 v_2)}. \tag{17}$$

Now, we consider (14) where we plug in (15), the guess for $V'(\lambda)$, and (17). After some algebra, we obtain

$$\begin{aligned}V'(\lambda) =& \frac{\lambda(1+4\delta\alpha v_2(2\alpha-1))}{2(1-\delta(1-\alpha)^2 v_2)}\\ &+\frac{2(1+\delta\alpha v_1)\big(1-\delta(1-\alpha)^2 v_2\big)-1+\delta(1-\alpha)(2\alpha v_2(1+\delta(1-\alpha)v_1)-v_1)}{2(1-\delta(1-\alpha)^2 v_2)}. \quad (18)\end{aligned}$$

We use the guess for the left-hand side in order to match and thereby solve for coefficients. First, we require that the coefficient of λ on the right-hand side of (18) equals $2v_2$, the coefficient of λ in (15), i.e.,

$$2v_2 = \frac{1+4\delta\alpha v_2(2\alpha-1)}{2(1-\delta(1-\alpha)^2 v_2)}.$$

Solving for v_2 gives

$$v_2 = \frac{1-\delta\alpha(2\alpha-1)\pm\sqrt{(1-\delta\alpha(2\alpha-1))^2-\delta(1-\alpha)^2}}{2\delta(1-\alpha)^2}. \quad (19)$$

We see that there are two possible solutions for v_2. Second, we require that the constant on the right-hand side of (18) equals v_1 which is the constant in the guess (15), i.e.,

$$v_1 = \frac{2(1+\delta\alpha v_1)\big(1-\delta(1-\alpha)^2 v_2\big)-1+\delta(1-\alpha)(2\alpha v_2(1+\delta(1-\alpha)v_1)-v_1)}{2(1-\delta(1-\alpha)^2 v_2)}.$$

Solving for v_1 yields

$$v_1 = \frac{2\big(1-\delta(1-\alpha)^2 v_2\big)-1+2\delta\alpha(1-\alpha)v_2}{2(1-\delta(1-\alpha)^2 v_2)(1-\delta\alpha)+\delta(1-\alpha)(1-2\delta\alpha(1-\alpha)v_2)} \quad (20)$$

which depends on v_2.

Step 2: We argue that only one of the two solutions for v_2 is valid.

Denote the two candidates for v_2 by

$$v_2^+ = \frac{1-\delta\alpha(2\alpha-1)+\sqrt{(1-\delta\alpha(2\alpha-1))^2-\delta(1-\alpha)^2}}{2\delta(1-\alpha)^2}, \quad (21)$$

$$v_2^- = \frac{1-\delta\alpha(2\alpha-1)-\sqrt{(1-\delta\alpha(2\alpha-1))^2-\delta(1-\alpha)^2}}{2\delta(1-\alpha)^2}. \quad (22)$$

First, recall from Lemma 1 that the value function $V(\lambda)$ is continuous in α and δ. Rewriting the square root in the numerator gives

$$\sqrt{(1-\delta\alpha^2)(1-\delta+4\alpha\delta(1-\alpha))}.$$

Observe that both solutions coincide if and only if $\alpha = \delta = 1$. As we consider $\delta \in (0,1)$ we conclude that $v_2^+ \neq v_2^-$. Hence, only one of the two candidates can be the global solution for any pair (α,δ). If both were solutions, $V(\lambda)$ would not be continuous.

Second, we argue that value function induced by v_2^+ is not increasing for all pairs (α,δ) which contradicts Proposition 2. Let v_1^+ be the solution to (20) when we plug in v_2^+ for v_2. Some algebra yields

$$v_1^+ = \frac{\alpha + (1-2\alpha)\left(\delta\alpha(2\alpha-1) - \sqrt{(1-\delta\alpha^2)(1-\delta+4\alpha\delta(1-\alpha))}\right)}{(1-\alpha)\left(1-\delta\alpha+\delta(1-\alpha)(1-2\alpha) - \sqrt{(1-\delta\alpha^2)(1-\delta+4\alpha\delta(1-\alpha))}\right)}. \tag{23}$$

Next, we show that there exists a pair of (α,δ) for which v_1^+ is strictly negative which contradicts that the value function $V(\lambda)$ is strictly increasing in λ for any pair (δ,α). For $\alpha = 0$ we get

$$\frac{-\sqrt{1-\delta}}{1+\delta-\sqrt{1-\delta}}, \tag{24}$$

which is strictly negative for all $\delta \in (0,1)$. By continuity, we conclude that there exists an $\alpha > 0$ such that $v_1^+ < 0$. Consequently, v_2^+ is not the solution. Instead, the solution is given by v_2^- and the corresponding solution for v_1, that is,

$$v_1^- = \frac{\alpha + (1-2\alpha)\left(\delta\alpha(2\alpha-1) + \sqrt{(1-\delta\alpha^2)(1-\delta+4\alpha\delta(1-\alpha))}\right)}{(1-\alpha)\left(1-\delta\alpha+\delta(1-\alpha)(1-2\alpha) + \sqrt{(1-\delta\alpha^2)(1-\delta+4\alpha\delta(1-\alpha))}\right)}. \tag{25}$$

From now on, we identify v_1 with v_1^- and v_2 with v_2^-.

Step 3: We verify that the second-order condition is satisfied.

Consider the second-order derivative of the right-hand side of the Bellman equation with respect to $\theta^\dagger$

$$-2\left(1-\delta(1-\alpha)^2 v_2\right). \tag{26}$$

We start by showing that $v_2 \geq 0$ which is the case if the numerator of v_2^- is positive, i.e.,

$$1-\delta\alpha(2\alpha-1) \;\geq\; \sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}.$$

By squaring both sides of the equation and subtracting $(1-\delta\alpha(2\alpha-1))^2$, the condition simplifies to

$$\delta(1-\alpha)^2 \;\geq\; 0$$

which is always satisfied. Therefore, the second-order condition holds if and only if $\delta(1-\alpha)^2 v_2 < 1$. We show that $2(1-\alpha)v_2 \leq 1$ which implies that $\delta(1-\alpha)^2 v_2 <$

1. After plugging in (22), we can rewrite $2(1-\alpha)v_2 \leq 1$ to

$$1 - \delta\alpha(2\alpha - 1) - \sqrt{(1 - \delta\alpha(2\alpha - 1))^2 - \delta(1 - \alpha)^2} \;\leq\; \delta(1 - \alpha).$$

Rearranging terms and squaring both sides yields

$$(1 - \delta\alpha(2\alpha - 1))^2 - \delta(1 - \alpha)^2 \;\geq\; (1 - \delta + 2\delta\alpha(1 - \alpha))^2 .$$

After some algebra, this condition can be reformulated to

$$-1 + \alpha^2 + \delta\left(1 - 4\alpha + 7\alpha^2 - 4\alpha^3\right) \;\leq\; 0. \tag{27}$$

For $\delta = 0$, it simplifies to

$$\alpha^2 - 1 \;\leq\; 0$$

which is always true, and for $\delta = 1$ we get

$$-4(1 - \alpha)^2 \;\leq\; 0$$

which is satisfied as well. As (27) is linear in δ, the condition holds for all $\delta \in (0, 1)$ and we can conclude that $\delta(1-\alpha)^2 v_2 < 1$. Hence, the second-order condition is satisfied.

(ii) The shape of the policy function follows now immediately from (17), where we plug in v_1 and v_2 in order to determine the solutions for θ_0 and θ_1. We are left to show that $-1 < \theta_1 < 0$. θ_1 is defined as the coefficient of λ in (17), i.e.,

$$\theta_1 \;=\; -\frac{1 - \delta(1 - \alpha)(2v_2\alpha)}{2(1 - \delta(1 - \alpha)^2 v_2)}. \tag{28}$$

Recall from Step 3 of the previous part of the proof that $v_2 \geq 0$ and $2(1-\alpha)v_2 \leq 1$ which implies that $2\delta(1-\alpha)\alpha v_2 < 1$ and $\delta(1-\alpha)^2 v_2 < 1$. Therefore, numerator and denominator are positive and θ_1 is strictly negative. Moreover, $\theta_1 > -1$ if the denominator is larger than the numerator. This is the case if $2\delta(1-\alpha)^2 v_2 < 1$ which follows from $2(1-\alpha)v_2 \leq 1$. Recall also that the second-order derivative of the right-hand side of the Bellman equation, (26), is strictly negative. Hence, the objective is strictly concave and its maximizer is unique. Therefore, the policy function is also unique. This concludes the proof. □

Proof of Proposition 4.

(i) Consider $\alpha < 1$. In the first step, we define

$$T \; : \quad [0,1] \longrightarrow [0,1], \qquad \lambda_t \longmapsto \lambda_{t+1},$$

as the mapping of today's into tomorrow's reputation. Using the seller's policy function (7) and the reputation transition (6), we obtain

$$T(\lambda_t) \;=\; \alpha\lambda_t + (1-\alpha)(\theta_0 + \theta_1\lambda_t).$$

Note that the mapping is linear in λ_t with slope $\alpha + (1-\alpha)\theta_1$. As $\theta_1 \in (-1,0)$, it follows that $\alpha + (1-\alpha)\theta_1 \in (0,1)$ for any $\alpha \in [0,1)$ and, hence, T is a contraction mapping.

By Banach's fixed-point theorem, the contraction mapping T has a unique fixed point, that is, there is a unique $\hat{\lambda}$ such that $T(\hat{\lambda}) = \hat{\lambda}$, and, moreover, the sequence $\lambda_{t+1} = T(\lambda_t)$ converges to $\hat{\lambda}$ as $t \to \infty$.

Second, we explicitly compute $\hat{\lambda}$. To this end, we revisit the seller's problem from the proof of Proposition 3, in particular, the seller's first-order condition, as in (16),

$$1 - 2\theta^{\dagger}(\lambda) - \lambda + \delta(1-\alpha)V'\big(\alpha\lambda + (1-\alpha)\theta^{\dagger}(\lambda)\big) \;=\; 0,$$

and the first-order derivative with respect to λ, as in (18),

$$V'(\lambda) \;=\; 1 - \theta^{\dagger}(\lambda) + \delta\alpha V'\big(\alpha\lambda + (1-\alpha)\theta^{\dagger}(\lambda)\big).$$

Note that $T(\hat{\lambda}) = \hat{\lambda}$ implies $\theta^{\dagger}(\hat{\lambda}) = \hat{\lambda}$. By evaluating both conditions at the fixed point $\lambda = \hat{\lambda}$, we get

$$1 - 3\hat{\lambda} + \delta(1-\alpha)V'(\hat{\lambda}) \;=\; 0, \tag{29}$$

and

$$V'(\hat{\lambda}) \;=\; 1 - \hat{\lambda} + \delta\alpha V'(\hat{\lambda}). \tag{30}$$

Solving (30) for $V'(\hat{\lambda})$ yields

$$V'(\hat{\lambda}) \;=\; \frac{1-\hat{\lambda}}{1-\delta\alpha} \tag{31}$$

which we plug into (29) in order to solve for $\hat{\lambda}$. The solution is

$$\hat{\lambda} = \frac{1 + \frac{\delta(1-\alpha)}{1-\delta\alpha}}{3 + \frac{\delta(1-\alpha)}{1-\delta\alpha}},$$

as claimed.

Finally, we consider the derivatives of $\hat{\lambda}$ with respect to δ and α. For convenience, denote $\frac{\delta(1-\alpha)}{1-\delta\alpha}$ by $x(\alpha,\delta)$. As

$$\frac{\partial\hat{\lambda}}{\partial\delta} = \frac{2\frac{\partial x(\alpha,\delta)}{\partial\delta}}{(3 + x(\alpha,\delta))^2}, \quad \text{and} \quad \frac{\partial\hat{\lambda}}{\partial\alpha} = \frac{2\frac{\partial x(\alpha,\delta)}{\partial\alpha}}{(3 + x(\alpha,\delta))^2},$$

the signs of the derivatives of $\hat{\lambda}$ are determined by the signs of the derivatives of $x(\alpha,\delta)$. The latter are given by

$$\frac{\partial x(\alpha,\delta)}{\partial\delta} = \frac{1-\alpha}{(1-\delta\alpha)^2} > 0, \quad \text{and} \quad \frac{\partial x(\alpha,\delta)}{\partial\alpha} = -\frac{\delta(1-\delta)}{(1-\delta\alpha)^2} < 0.$$

Hence, $\hat{\lambda}$ is increasing in δ and decreasing in α.

(ii) For $\alpha = 1$, (6) implies that $\lambda_{t+1} = \lambda_t$ and, hence, $\lambda_t = \lambda_0$, for all $t \geq 0$, which completes the proof. □

Proof of Lemma 2. Note that for $\alpha = 1$, (6) implies that $\lambda_{t+1} = \lambda_t$. When starting with initial reputation λ_0, it follows that $\lambda_t = \lambda_0$, for all $t \in \{1, \ldots\}$. Thus, the seller's decision about $\theta^\dagger$ in period t does not influence her reputation in the next period. Her problem reduces to maximize stage game payoffs in each period.

Therefore, the seller faces the following problem, in each period,

$$\max_{\theta^\dagger} \{(\lambda_0 + \theta^\dagger)(1 - \theta^\dagger)\}.$$

The objective is strictly concave in $\theta^\dagger$. The first-order condition is given by

$$1 - \lambda_0 - 2\theta^\dagger = 0,$$

and solving for $\theta^\dagger$ yields

$$\theta^\dagger = \frac{1-\lambda_0}{2}.$$

The seller's profits are given by the discounted infinite sum of period profits. In each period, the seller receives

$$\left(\lambda_0 + \frac{1-\lambda_0}{2}\right)\left(1 - \frac{1-\lambda_0}{2}\right) = \left(\frac{1+\lambda_0}{2}\right)^2.$$

Hence, her profits are given by

$$\sum_{t=0}^{\infty} \delta^t \left(\frac{1+\lambda_0}{2}\right)^2 = \frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2.$$

□

Proof of Lemma 3. This result follows from Propositions 3 and 4. First, consider v_1 as in (25) and v_2 as in (22) for $\alpha = 0$,

$$v_1 = \frac{\sqrt{1-\delta}}{1+\delta+\sqrt{1-\delta}}, \tag{32}$$

$$v_2 = \frac{1-\sqrt{1-\delta}}{2\delta}. \tag{33}$$

We determine the seller's policy function by evaluating (13) for $\alpha = 0$,

$$\theta^*(\lambda) = -\frac{\lambda}{2(1-\delta v_2)} + \frac{1+\delta v_1}{2(1-\delta v_2)}. \tag{34}$$

With (32) and (33), we can rewrite (34) after some algebra to

$$\theta^*(\lambda) = -\frac{\lambda}{1+\sqrt{1-\delta}} + \frac{1+\delta}{1+\delta+\sqrt{1-\delta}}. \tag{35}$$

In order to determine the long-run reputation, we evaluate (8) at $\alpha = 0$ which yields

$$\hat{\lambda} = \frac{1+\delta}{3+\delta}.$$

Further, from (6) for $\alpha = 0$ follows $\lambda_{t+1} = \theta_t^\dagger$. From the proof of Proposition 4, we know that $\theta^*(\hat{\lambda}) = \hat{\lambda}$. Moreover, $\theta^*(\lambda)$ is decreasing by Proposition 3. Therefore, $\lambda_t > \hat{\lambda}$ implies $\lambda_{t+1} < \hat{\lambda}$ and vice versa, that is, reputation oscillates towards $\hat{\lambda}$. This completes the proof. □

Proof of Proposition 5. We proceed in several steps.

Step 1: We derive θ_1 in closed form.

It is convenient to define

$$\Delta := (1-\alpha^2\delta)(1-\delta(1-2\alpha)^2).$$

Recall that we obtained θ_1 from the first-order condition as

$$\theta_1 = \frac{\delta(1-\alpha)\alpha 2 v_2 - 1}{2(1-\delta(1-\alpha)^2 v_2)},$$

where

$$v_2 = \frac{1-\delta\alpha(2\alpha-1)-\sqrt{\Delta}}{2\delta(1-\alpha)^2}.$$

Inserting v_2 into θ_1 gives

$$\frac{\frac{\alpha}{1-\alpha}\left(1-\frac{1-\alpha}{\alpha}-\delta\alpha(2\alpha-1)-\sqrt{\Delta}\right)}{1+\delta\alpha(2\alpha-1)+\sqrt{\Delta}}. \tag{36}$$

Next, we multiply the numerator and the denominator of (36) by

$$1-\frac{1-\alpha}{\alpha}-\delta\alpha(2\alpha-1)+\sqrt{\Delta}.$$

Consider first the numerator of (36):

$$\frac{\alpha}{1-\alpha}\left(\left(1-\frac{1-\alpha}{\alpha}-\delta\alpha(2\alpha-1)\right)^2-\Delta\right).$$

Some algebra shows that the numerator can be written as

$$-3+\frac{1}{\alpha}+\alpha\delta(3\alpha-1). \tag{37}$$

Now, consider the denominator of (36),

$$\left(1+\delta\alpha(2\alpha-1)+\sqrt{\Delta}\right)\left(1-\frac{1-\alpha}{\alpha}-\delta\alpha(2\alpha-1)+\sqrt{\Delta}\right).$$

Simplifying the latter expression yields

$$\left(\sqrt{\Delta}+(1-\alpha^2\delta)\right)\left(\frac{3\alpha-1}{\alpha}\right). \tag{38}$$

Putting (37) and (38) together, we obtain for θ_1:

$$\theta_1 = \frac{-3+\frac{1}{\alpha}+\alpha\delta(3\alpha-1)}{\left(\sqrt{\Delta}+(1-\alpha^2\delta)\right)\left(\frac{3\alpha-1}{\alpha}\right)},$$

which we can further simplify to

$$\theta_1 = \frac{-(1-\alpha^2\delta)}{1-\alpha^2\delta+\sqrt{\Delta}}. \tag{39}$$

Step 2: We derive $\frac{\partial \theta_1}{\partial \alpha}$ in closed form.
We start by taking the derivative of θ_1, see (39) from Step 1, with respect to α:

$$\frac{2\alpha\delta(1-\alpha^2\delta+\sqrt{\Delta})}{\left(1-\alpha^2\delta+\sqrt{\Delta}\right)^2} - \frac{(\alpha^2\delta-1)\left(-2\alpha\delta+\frac{1}{2}\Delta^{-\frac{1}{2}}\left(-2\delta\alpha(1-\delta(1-2\alpha)^2)+(1-\alpha^2\delta)(4\delta(1-2\alpha))\right)\right)}{\left(1-\alpha^2\delta+\sqrt{\Delta}\right)^2}.$$

Next, we multiply numerator and denominator by $\sqrt{\Delta}$. The numerator becomes

$$2\alpha\delta\Delta+(1-\alpha^2\delta)\frac{1}{2}\left(-2\delta\alpha(1-\delta(1-2\alpha)^2)+(1-\alpha^2\delta)(4\delta(1-2\alpha))\right),$$

and simplifying yields

$$\delta(1-\alpha^2\delta)(2-3\alpha-\alpha\delta+2\alpha^2\delta).$$

Putting numerator and denominator together gives

$$\frac{\partial\theta_1}{\partial\alpha} = \frac{\delta(1-\alpha^2\delta)(2+2\alpha^2\delta-3\alpha-\alpha\delta)}{(1-\alpha^2\delta+\sqrt{\Delta})^2\sqrt{\Delta}}.$$

Inserting the definition of Δ yields

$$\frac{\partial\theta_1}{\partial\alpha} = \frac{\delta(1-\alpha^2\delta)(2+2\alpha^2\delta-3\alpha-\alpha\delta)}{(1-\alpha^2\delta+\sqrt{(1-\alpha^2\delta)(1-\delta(1-2\alpha)^2)})^2\sqrt{(1-\alpha^2\delta)(1-\delta(1-2\alpha)^2)}}. \tag{40}$$

Step 3: $(1-\alpha)\theta_1$ is increasing in α.
The statement is equivalent to showing that

$$-\theta_1+(1-\alpha)\frac{\partial\theta_1}{\partial\alpha} \geq 0. \tag{41}$$

We plug in the explicit formulas for θ_1 and $\frac{\partial\theta_1}{\partial\alpha}$ which we derived in Step 1 and 2 and rewrite the left side of inequality (41) as a single fraction over the denominator

$$(1-\alpha^2\delta+\sqrt{\Delta})^2\sqrt{\Delta}.$$

As the denominator is positive, it suffices to argue that the numerator

$$(1-\alpha)\delta(1-\alpha^2\delta)(2+2\alpha^2\delta-3\alpha-\alpha\delta)+(1-\alpha^2\delta)(1-\alpha^2\delta+\sqrt{\Delta})(\sqrt{\Delta})$$

is also positive. The latter expression is bounded from below by

$$(1-\alpha)\delta(1-\alpha^2\delta)(2+2\alpha^2\delta-3\alpha-\alpha\delta)+(1-\alpha^2\delta)\Delta.$$

The definition of Δ and some algebra yield

$$(1-\alpha^2\delta)\big((1-\delta)(1-\alpha^2\delta)+\delta(1-\alpha)(2+2\alpha^2\delta+\alpha-\alpha\delta-4\alpha^3\delta)\big).$$

It remains to be shown that

$$2+2\alpha^2\delta+\alpha-\alpha\delta-4\alpha^3\delta \tag{42}$$

is positive. Observe that (42) is linear in δ. Therefore, it attains its minimum either at $\delta=1$ or $\delta=0$. For $\delta=0$ we obtain

$$2+\alpha \;>\; 0,$$

and for $\delta=1$, (42) becomes

$$2+2\alpha^2-4\alpha^3 \;\geq\; 0.$$

Consequently, (42) is positive for all $\delta\in(0,1)$ which implies that (41) holds.

Step 4: We prove (i)-(iii).

Note that $\alpha+(1-\alpha)\theta_1$ is continuous in α, strictly increasing in α by Step 3, less than zero at $\alpha=0$, and larger than zero at $\alpha=1$. Therefore, we can uniquely define $\bar{\alpha}$ through

$$\bar{\alpha}+(1-\bar{\alpha})\theta_1 \;=\; 0.$$

(i) Consider $\alpha>\bar{\alpha}$. Fix $\lambda_0\geq\hat{\lambda}$. Denote by λ the current-period reputation and by λ' the next-period reputation. First, we argue that $\lambda\geq\hat{\lambda}$ implies $\lambda'\geq\hat{\lambda}$:

$$\begin{aligned}\lambda' &= \alpha\lambda+(1-\alpha)\theta^*(\lambda)\\ &= \alpha\lambda+(1-\alpha)(\theta_1\lambda+\theta_0)\\ &= \lambda\underbrace{(\alpha+(1-\alpha)\theta_1)}_{>0 \text{ as } \alpha>\bar{\alpha}}+(1-\alpha)\theta_0\\ &\geq \hat{\lambda}(\alpha+(1-\alpha)\theta_1)+(1-\alpha)\theta_0\\ &= \alpha\hat{\lambda}+(1-\alpha)\theta^*(\hat{\lambda})\\ &= \hat{\lambda}.\end{aligned} \tag{43}$$

Thus, if $\lambda_0\geq\hat{\lambda}$ then $\lambda_t\geq\hat{\lambda}$, for all t. Second, we show that $\lambda\geq\hat{\lambda}$ implies $\lambda\geq\lambda'$. Recall that $\theta^*(\hat{\lambda})=\hat{\lambda}$, and that θ^* is decreasing in λ. Thus, we have

$$\theta^*(\lambda) \;<\; \hat{\lambda} \;<\; \lambda,$$

and consequently

$$\lambda' \;=\; \alpha\lambda+(1-\alpha)\theta^*(\lambda) \;\leq\; \alpha\lambda+(1-\alpha)\lambda \;=\; \lambda.$$

We deduce that if $\lambda_0 \geq \hat{\lambda}$ then $\lambda_t \geq \lambda_{t+1}$, $\forall t$. From Proposition 4 we know that $\lim_{t\to\infty} \lambda_t = \hat{\lambda}$, hence, $|\lambda_t - \hat{\lambda}| \downarrow 0$. The proof for $\lambda_0 \leq \hat{\lambda}$ is analogous.

(ii) Let $\alpha = \bar{\alpha}$. Fix any $\lambda_0 \in [0,1]$. Carefully inspecting the arguments from (i) shows that (43) holds with equality if $\alpha = \hat{\alpha}$. Thus, we have $\lambda_t = \hat{\lambda}$, for all $t \geq 1$.

(iii) Consider $\alpha < \bar{\alpha}$. Fix any $\lambda_0 \in [0,1]$. In this case the inequality in (43) is reversed, i.e., if $\lambda \geq \hat{\lambda}$ then $\lambda' \leq \hat{\lambda}$. The proof that if $\lambda \leq \hat{\lambda}$ then $\lambda' \geq \hat{\lambda}$ follows from an analogous computation. Next, we show that

$$|\lambda - \hat{\lambda}| \;\geq\; |\lambda' - \hat{\lambda}|. \tag{44}$$

Again, we focus on $\lambda \geq \hat{\lambda}$, the proof for $\lambda \leq \hat{\lambda}$ is analogous. If $\lambda \geq \hat{\lambda}$, (44) becomes

$$\lambda - \hat{\lambda} \;\geq\; \hat{\lambda} - \lambda' \quad \Leftrightarrow \quad \lambda + \lambda' \;\geq\; 2\hat{\lambda}. \tag{45}$$

Observe that

$$\begin{aligned} \lambda + \lambda' \;&=\; \lambda \underbrace{(1 + \alpha + (1-\alpha)\theta_1)}_{>0} + (1-\alpha)\theta_0 \\ &\geq\; \hat{\lambda} + \hat{\lambda}(\alpha + (1-\alpha)\theta_1) + (1-\alpha)\theta_0 \\ &=\; 2\hat{\lambda}, \end{aligned}$$

which yields (45). Together with the convergence established in Proposition 4, we obtain $|\lambda_t - \hat{\lambda}| \downarrow 0$. This concludes the proof. □

Proof of Proposition 6. We argue that

$$|\lambda'(\alpha) - \lambda(\alpha)| \tag{46}$$

is decreasing in α. Inserting the definition of $\lambda'(\alpha)$ into (46), we obtain

$$(1-\alpha)|\theta^*(\lambda(\alpha)) - \lambda(\alpha)|.$$

For $h < 0$, we have $\lambda(\alpha) < \hat{\lambda}(\alpha)$ and thus $\theta^*(\lambda(\alpha)) > \hat{\lambda}(\alpha) > \lambda(\alpha)$. Therefore, we obtain

$$(1-\alpha)(\theta^*(\lambda(\alpha)) - \lambda(\alpha)). \tag{47}$$

Recall that

$$\theta^*(\lambda(\alpha)) \;=\; \theta_1\lambda(\alpha) + \theta_0. \tag{48}$$

Taking the derivative of (47) with respect to α we obtain

$$-\big(\theta_1(\hat{\lambda}(\alpha) + h) + \theta_0 - (\hat{\lambda}(\alpha) + h)\big) + (1-\alpha)\left(\frac{\partial\theta_1}{\partial\alpha}(\hat{\lambda}(\alpha) + h) + \theta_1\frac{\partial\hat{\lambda}}{\partial\alpha} + \frac{\partial\theta_0}{\partial\alpha} - \frac{\partial\hat{\lambda}}{\partial\alpha}\right), \tag{49}$$

where we used (48) and the definition of $\lambda(\alpha)$. By definition of the long-run reputation $\hat{\lambda}(\alpha)$, it holds

$$\hat{\lambda}(\alpha) = \theta_1 \hat{\lambda}(\alpha) + \theta_0. \tag{50}$$

Taking the derivative of the fixed-point condition with respect to α, we get

$$\frac{\partial \hat{\lambda}}{\partial \alpha} = \frac{\partial \theta_1}{\partial \alpha}\hat{\lambda}(\alpha) + \theta_1 \frac{\partial \hat{\lambda}}{\partial \alpha} + \frac{\partial \theta_0}{\partial \alpha}. \tag{51}$$

Substituting (50) and (51), simplifies (49) to

$$h\left(1 - \theta_1 + (1-\alpha)\frac{\partial \theta_1}{\partial \alpha}\right). \tag{52}$$

We want to argue that (52) is negative. This is equivalent to showing that the term in brackets is positive as $h < 0$. From Step 3 of the proof of Proposition 5 we know that $(1-\alpha)\theta_1$ is increasing in α, i.e.,

$$-\theta_1 + (1-\alpha)\frac{\partial \theta_1}{\partial \alpha} \geq 0. \tag{53}$$

Thus (52) is negative and the distance between $\lambda'(\alpha)$ and $\lambda(\alpha)$ is decreasing in α for $h < 0$. The proof for $h > 0$ is analogous. □

Proof of Proposition 7. We know from Lemma 2 that

$$V(\lambda_0, 1) = \frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2,$$

and from Proposition 3 that

$$V(\lambda_0, \alpha) = v_2 \lambda_0^2 + v_1 \lambda_0 + v_0,$$

for $\alpha < 1$, where v_2 is given by (22) and v_1 by (25). Hence, $V(\lambda_0, 1) - V(\lambda_0, \alpha)$ is given by

$$\frac{1}{1-\delta}\left(\frac{1+\lambda_0}{2}\right)^2 - \left(v_2 \lambda_0^2 + v_1 \lambda_0 + v_0\right). \tag{54}$$

From here, we proceed in four steps.

Step 1: The difference $V(\lambda_0, 1) - V(\lambda_0, \alpha)$ is strictly convex, for all $\alpha < 1$.

In order to determine the curvature of (54), we consider the coefficient of λ_0^2, which is

$$\frac{1}{4(1-\delta)} - \frac{1 - \delta\alpha(2\alpha-1) - \sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}}{2\delta(1-\alpha)^2}.$$

Rewriting fractions to have the same denominator yields

$$\frac{(1-\alpha)^2\delta - 2(1-\delta)(1-\delta\alpha(2\alpha-1)) + 2(1-\delta)\sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}}{4\delta(1-\delta)(1-\alpha)^2}. \tag{55}$$

Note that the denominator of (55) is strictly positive for any $\alpha < 1$. Let

$$\begin{aligned} a &:= (1-\alpha)^2\delta - 2(1-\delta)(1-\delta\alpha(2\alpha-1)), \quad \text{and} \\ b &:= 2(1-\delta)\sqrt{(1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2}. \end{aligned}$$

The difference, (54), is strictly convex if and only if the numerator of (55) is strictly positive, that is, if $a + b > 0$. We start by considering the product $(a+b)(a-b) = a^2 - b^2$, that is

$$\left((1-\alpha)^2\delta - 2(1-\delta)(1-\delta\alpha(2\alpha-1))\right)^2 - 4(1-\delta)^2\left((1-\delta\alpha(2\alpha-1))^2 - \delta(1-\alpha)^2\right).$$

After some algebra, we obtain

$$\delta^2(1-\alpha)^2\left[(1-\alpha)^2 - (1-\delta)\left(4+4\alpha-8\alpha^2\right)\right]. \tag{56}$$

The sign of (56) is determined by the sign of

$$(1-\alpha)^2 - (1-\delta)\left(4+4\alpha-8\alpha^2\right). \tag{57}$$

As $4+4\alpha-8\alpha^2 = 4(1-\alpha^2)+4\alpha(1-\alpha) > 0$, (57) is strictly increasing in δ. Consider (57) for $\delta = 0$,

$$(1-\alpha)^2 - \left(4+4\alpha-8\alpha^2\right) = -\left[3(1-\alpha^2) + 6\alpha(1-\alpha)\right] < 0.$$

Showing that (57) is strictly negative for $\delta = 0$ implies that $(a+b)(a-b) = a^2 - b^2$ is strictly negative for all $\delta \in (0,1)$. Hence, either $a-b$ or $a+b$ must be negative but not both. Since $a+b > a-b$, we know that only $a-b$ can be negative and, hence, $a+b$ is positive which implies that (54) is strictly convex, as claimed.

Step 2: $V(1,1)$ is an upper bound for $V(\lambda_0,\alpha)$.

In period t, the seller's period profits from choosing $\theta_t^\dagger$ are given by

$$\left(1-\theta_t^\dagger\right)\left(\lambda_t + \theta_t^\dagger\right). \tag{58}$$

Note that (58) is strictly increasing in λ_t. Hence, obtaining the maximum of (58) for $\lambda_t = 1$ in each period is an upper bound for $V(\lambda_0,\alpha)$. For the case where $\alpha = 1$, we have $\lambda_t = \lambda_0 = 1$ for all $t \in \{0,1,\dots\}$, and, therefore, the seller attains this upper bound by Lemma 2. For $\alpha < 1$, the seller's value $V(\lambda_0,\alpha)$ is smaller than the upper bound which implies that (54) is positive for $\lambda_0 = 1$.

Step 3: For any $\alpha < 1$, $V(0,1)$ is strictly smaller than $V(0,\alpha)$.

From Lemma 2, we know that $\theta_t^\dagger = 1/2$, for all $t \in \{0,1,\dots\}$, if $\alpha = 1$ and $\lambda_0 = 0$. In each period, the seller thus realizes demand of $1/2$ at a constant price of $1/2$ which yields a value of

$$\frac{1}{4(1-\delta)}.$$

Now, consider the seller's profit for the case where $\alpha < 1$ when she equally sets $\theta_t^{\dagger} = 1/2$ for all $t \in \{0, 1, \ldots\}$. Again, she realizes demand of $1/2$ in each period but at higher prices for $t \geq 1$ because $\alpha < 1$ implies that $\lambda_t > 0$ for $t \geq 1$. Consequently, $V(0, 1)$ is strictly smaller than $V(0, \alpha)$ which implies that (54) is negative for $\lambda_0 = 0$.
Step 4: Existence of $\bar{\lambda}_0(\alpha) \in (0, 1)$.
We know from the previous steps that (54) is strictly convex, strictly positive at $\lambda_0 = 1$, and strictly negative at $\lambda_0 = 0$. This implies that for any $\alpha < 1$, there exists a unique cut-off in $(0, 1)$. For any $\lambda_0 > \bar{\lambda}_0(\alpha)$, (54) is strictly positive, and hence $V(\lambda_0, 1) > V(\lambda_0, \alpha)$. Conversely, for any $\lambda_0 < \bar{\lambda}_0(\alpha)$, (54) is strictly negative, and hence $V(\lambda_0, 1) < V(\lambda_0, \alpha)$. For $\lambda_0 = \bar{\lambda}_0(\alpha)$, (54) equals zero which concludes the proof. □

4

Redistributional Effects of Health Insurance in Germany: Private and Public Insurance, Premia and Contribution Rates*

This chapter develops a model of the German health insurance system for identifying redistribution streams and evaluating proposals to change the system. A population, characterized by health and income, obtains health insurance either from a budget-balancing public insurer or a more flexible, revenue-maximizing private insurer. Redistribution occurs across health and income, and the private insurer extracts surplus, by attracting profitable customers, which cannot be used for redistribution. We analyze changes in redistribution when switching from the current contribution-based system to a premium-based system with only one type of insurer. Furthermore, we study the properties of welfare-maximizing fee schedules.

4.1 Introduction

In many industrialized countries health costs constitute a significant portion of GDP. As health costs continue to increase due to technological progress, higher life expectancy, and an aging population, and at the same time the prospects for the wider economy decline with periods of low economic growth and high rates of unemployment, the question arises of how the burden of financing health costs is currently distributed and how it could be distributed in the future. Which parts of the population (will) subsidize which other parts of the population through the health insurance system?

---

* Based on joint work with Linda Schilling (see Schickner and Schilling, 2016). We would like to thank Deniz Dizdar, Eugen Kovac, Daniel Krähmer, Benny Moldovanu, and participants of the BGSE Micro Workshop, the BGSE Brown Bag Seminar, and the MPI Econ Workshop in Bonn for helpful comments and suggestions.

Understanding the organizational structure of the health insurance market is key to understanding how the health insurance system redistributes income between different parts of the population. As in many other countries, the question of how to organize the market for health insurance is a topic of intense public debate in Germany. Several competing policy proposals for changing the organizational structure of the health insurance market are on the table, with commentators varying in their praise of the advantages in terms of efficiency and equity of one or the other.

While economic theory has extensively studied abstract health insurance markets, in particular the implications of private information and the resulting problems of adverse selection and moral hazard, it comes at a surprise that relatively little work has been done on explicitly modeling specific health insurance systems in order to substantiate descriptive analysis and evaluate policy proposals.

In this chapter we provide an analytic framework for the German health insurance system for assessing redistribution streams in the current health insurance system and evaluating policy proposals aimed at changing the organizational structure of the health insurance market.

We consider a model of the German health insurance market with three kinds of agents.[1] There is a continuum population of customers. Each customer is characterized by his income and his health, which are positively correlated. Corresponding to the customer's health are his health costs. Health insurance is provided by two insurers, a public insurer and a private insurer. Health insurance is obligatory. There exists an income threshold, the opt-out threshold, below which customers must insure with the public insurer; customers with income above the opt-out threshold can choose between the public and the private insurer. Health insurance provides customers with a monetary benefit; it reimburses their health costs up to a fixed level, the benefit level. We refer to this monetary benefit as health benefit. Health insurers charge their customers a fee to finance their expenditures. The public insurer charges a fixed percentage, the contribution rate, of its customers' income up to a contribution cap. The private insurer is more flexible in that its fee may additionally condition on the health type. We assume that all market participants can observe customers' characteristics. Also, the private insurer is obliged to charge a fee which is lower than the contribution cap. Whereas the private insurer maximizes profits, the public insurer aims at running a balanced budget, i.e., it balances the health benefits to its customers with the fees it collects from its customers. We study the health insurance market in two periods. First, insurers set their fees for every eligible customer. Then, customers choose their insurance.[2]

---

[1] See Breyer (2004), Wörz and Busse (2005), or Grunow and Nuscheler (2014) for exhaustive descriptions of the German health insurance system.

[2] Although somewhat peculiar, the German health insurance system bears some similarities with the health insurance systems in Argentina, Chile, Colombia, Peru, and the Netherlands, cf. Thomson and Mossialos (2006) and Panthöfer (2016). Also, other European countries have considered emulating the German system, see Thomson and Mossialos (2006).

To gain some intuition for our model, we start by analyzing the case when health insurance is voluntary. In this case, there exists no equilibrium in the health insurance market, i.e., the health insurance market collapses. Call a customer profitable if the difference between his fee at the public insurer and his health benefit is positive and unprofitable otherwise. Intuitively, if health insurance is voluntary, profitable customers prefer to bear their health costs themselves. Thus, health insurers only attract unprofitable customers. In particular, there is a group of low-income customers who are only eligible for public insurance and whose fees are strictly less than their health benefit. This group insures with the public insurer, which makes it impossible for the public insurer to run a balanced budget. In contrast to Rothschild and Stiglitz (1976), this unraveling of the market is not a result of asymmetric information but a result of the way the market is organized, in particular the restriction that the public insurer's fee only depends on the income type and not on the health type. As a result, the public insurer cannot fully account for customers' health costs. In addition to potential issues of asymmetric information, this result provides a rationale for why health insurance is obligatory in Germany.

Turning to obligatory insurance, we derive a condition under which there exists an equilibrium in the health insurance market. Roughly, the condition states that the population's health costs are not larger than the income of customers who must insure with the public insurer. It is instructive to look at the strategy of the private insurer. Because its fee can also depend on the customer's health, the private insurer is more flexible than the public insurer.[3] In equilibrium, the private insurer distinguishes between profitable and unprofitable customers. It tries to attract profitable customers by slightly undercutting the public fee and to deter unprofitable customers by setting the highest possible fee, i.e., the private insurer tries to cream skim customers. In equilibrium, all customers, profitable and unprofitable, with income below the opt-out threshold choose public insurance. Profitable customers with income above the opt-out threshold join the private insurer. Unprofitable customers with high income are indifferent between public and private insurance because both insurers charge the highest possible fee.

The redistribution streams are as follows. Profitable customers with income below the opt-out threshold subsidize all unprofitable customers which join the public insurer. The private insurer uses its flexibility in setting fees to extract the surplus of profitable, high-income customers and turns them into profits. These customers' surplus is lost for redistribution. Note that redistribution occurs across income and health.

It is interesting to observe that high-income customers would pay less if they were forced to insure with the public insurer. In equilibrium, they join the private insurer at a cost that is essentially equal to the fee of the public insurer. If they would join the public insurer, their surplus could be used for redistribution which would allow for a decrease in the contribution rate and thus a decrease in the fee they pay. Because the

[3] It is convenient to assume that all insurers provide the same benefit level which allows us to focus on fees. See Section 4.5 for the extension to heterogeneous benefit levels.

effect of a single customer on the contribution rate is negligible, the private insurer can prevent this by marginally undercutting the public fee.

The opportunity to opt-out of the redistributive public system was originally granted to enhance consumer choice and stimulate competition. However, our analysis indicates that because the private insurer has more flexibility in setting its fee it can favorably select its customers which leads to a higher fee for all customers. Surplus from profitable high-income customers is converted into profits for the private insurer and lost for redistribution through public insurance.

Our findings are consistent with the result in Pauly (1984) that cream skimming may occur as a result of regulation of the fee structure. Cream skimming in the German health insurance market is in line with the empirical findings of Grunow and Nuscheler (2014) and Bünnings and Tauchmann (2015). Grunow and Nuscheler (2014) find that private insurance customers who experienced a negative health shock are more likely to switch to the public system, thereby imposing additional costs on the public system. Bünnings and Tauchmann (2015) find that young, healthy, high-earning individuals are more likely to leave the public sector.[4]

We then proceed by analyzing how changes in policy parameters affect the equilibrium in the health insurance market. We find that increasing the opt-out threshold decreases the contribution rate and thereby the fee of the entire population. Consequently, all customers have higher utility. This is due to two effects: First, increasing the opt-out threshold increases the part of the population which must insure with the public insurer, thereby decreasing the range in which the private insurer may cream skim. Second, the customers which join the public insurer as a result of the increase in the opt-out threshold are on average more profitable than existing customers of the public insurer because they have comparatively high income, and income and health are correlated.

Next, we study structural changes in the population's health and income. First, we show that, surprisingly, systematic improvements of the population's health and income do not necessarily lead to a decrease in the contribution rate. For this purpose, we identify subgroups of the population which determine whether a specific improvement increases or decreases the contribution rate. The contribution rate may increase if many healthy, medium-income customers of the public insurer become eligible to choose between public and private insurance as a result of their income increase. These profitable customers leave the public insurer and join the private insurer. Thus, government programs aimed at increasing health or income may increase the contribution rate because of the organization of the health insurance market. Second, inspired by the increasing correlation between health and income (see, for example, Deaton and

[4] For an empirical assessment of selection between public and private plans in the US health insurance system see Brown et al. (2014), Newhouse et al. (2012), and Newhouse et al. (2015).

Paxson, 1998), we study the effect of an increase in correlation between health and income on the health insurance market.[5] We specify conditions under which an increase in correlation leads to an increase in the contribution rate. Intuitively, low-income customers choose public insurance; high-income customers choose private insurance. If the correlation between health and income increases, low-income customers also have worse health. Consequently, to balance its budget, the public insurer has to increase the contribution rate.

Having analyzed redistribution in the current German health insurance system, we evaluate how proposals to reorganize the health insurance market, currently being put forward by policy makers, would change redistribution. Two ideas have been at the center of the discussion (see, e.g., Breyer, 2004; Wissenschaftlicher Beirat beim Bundesministerium der Finanzen, 2004; Wörz and Busse, 2005). First, to abandon the difference in regulation between private and public insurers. Second, to replace the income-dependent contribution by a flat fee which is independent of income and health, a so called health premium. A comparable system is currently in place in Switzerland.

We adjust our model to incorporate these organizational changes and prove existence of equilibrium in the adjusted model. We find that in the premium-based system redistribution occurs only across the health dimension not across the income dimension. Consequently, compared to the contribution-based system, low-income customers suffer, whereas high-income customers gain. We further analyze these findings by comparing the (utilitarian) welfare of the population in the two systems. We identify two effects that are important. Because the customers' utility function is concave, redistributing income from high-income customers to low-income customers increases welfare. This redistribution effect speaks in favor of the contribution-based system in which the fee depends on income. On the other hand, in the contribution-based system, the private insurer extracts surplus from its customers which cannot be used for redistribution via the public insurance. This competition effect speaks in favor of the premium-based system in which insurers compete on equal grounds for customers from the entire population. Consequently, we obtain the following two results: If the opt-out threshold is sufficiently high, i.e., the influence of private insurance on the contribution-based system sufficiently low, the current contribution-based system yields higher welfare than the premium-based system. However, introducing a simple, budget-balanced income redistribution scheme (income tax) into the premium-based system permits to obtain the same level of welfare as one could obtain in the current German system, given the most favorable specification of the policy parameters. This implies that a change of the current system to a premium-based system ought to be accompanied by an appropriate increase in redistribution through, e.g., income taxation.

---

[5] For an analysis of the causal effect of income on health, using the German reunification as a natural experiment, see Frijters et al. (2005).

Motivated by the preceding analysis, we then study properties of welfare-maximizing fee schedules if health insurance is provided exclusively by a benevolent authority who has to finance the aggregate health expenditures. This gives practical insights into how to adjust the regulation of fee schedules to improve welfare. Equivalently, it sheds light on the optimal adjustment of income taxation when introducing a premium-based system.

First, we consider the case when the fee may condition on health and income. In this case, we derive the welfare-maximizing fee schedule in closed form. Let a customer's net income be his income minus his health costs. The optimal scheme exhibits strong redistributional features. First, it charges each customer the health costs he imposes on the system. Second, it charges/reimburses each member of the population the difference between his net income and the average net income in the population.

Second, inspired by the fee of the public insurer which only depends on income, we study properties of the welfare-maximizing fee schedule that only depends on income. We show that the optimal fee schedule increases quicker than income itself. This stands in stark contrast to the fee of the public insurer which is a percentage of income and therefore increases slower than income. Intuitively, the fee schedule tries to mimic the strong redistributional features of the welfare-maximizing fee schedule that conditions on both health and income. As income increases by one unit, net income increases by more than one unit because higher income entails lower health costs due to the correlation between health and income. The fee schedule aims at compensating this increase and therefore increases by more than one as well.

Related Literature. The related literature can be roughly grouped into two strands: on the one hand, the theoretical literature on insurance; on the other hand, policy-oriented papers addressing specifically the German health insurance system.

The theoretical literature has mostly analyzed abstract insurance markets, focusing on the implications of private information. The seminal paper by Rothschild and Stiglitz (1976) demonstrates that competitive insurance markets can fail as a consequence of asymmetric information. Combining ideas from Rothschild-Stiglitz and optimal income taxation (see Mirrlees, 1971), Rochet (1991) and Cremer and Pestieau (1996) argue that social insurance is desirable to supplement an optimal income tax if productivity and health risk are negatively correlated, that is, a more productive individual has a lower health risk. See also Blomqvist and Horn (1984). Nishimura (2009) and Boadway et al. (2006) extend this idea to include adverse selection and moral hazard in the health insurance market. Neudeck and Podczeck (1996) study different insurance systems within the Rothschild-Stiglitz model. They show how various combinations of private and public insurance, among them a German-style system, can implement second-best. Employing the Rawlsian welfare criterion, Franc and Abadie (2004) study when opting-out of public insurance is optimal in the Rochet model.

Besley (1989) and Blomqvist and Johansson (1997) discuss the welfare implications of moral hazard for insurance systems in which private insurance is supplementary to public insurance. Petretto (1999) studies the optimal design of such a system under efficiency and equity considerations.

In contrast to this strand of literature, we explicitly model the organizational structure of a specific health insurance system and study its redistributional implications, abstracting from issues of private information.

Cream skimming of private insurers due to the organizational structure of the German health insurance system has been described in, e.g., Breyer (2004), Jacobs and Schulze (2004), Wasem et al. (2004), Wissenschaftlicher Beirat beim Bundesministerium der Finanzen (2004), and Thomson and Mossialos (2006). Our work complements these papers by providing an analytical framework to corroborate their observations.

In contemporary, independent work Panthöfer (2016) studies selection in the German health insurance system by augmenting a Rothschild-Stiglitz model to fit the German health insurance system. In a subsequent empirical analysis, he finds evidence for risk selection: Good risks leave the public system and join the private sector.

Our comparison of the redistributional effects of financing the insurance system through income-dependent contributions and income-independent premiums relates to the seminal work of Atkinson and Stiglitz (1976). They show that in the presence of private information any Pareto-efficient allocation can be implemented solely relying on income taxation. No additional income redistribution, e.g. through contributions in the health insurance system, is necessary. In a policy-oriented approach, Breyer and Haufler (2000) put forward two arguments against income redistribution through health contributions: first, the facilitated adoption of incentive-compatible health insurance contracts, and, second, a reduction of the excess burden from implicit income taxation. Abstracting from informational considerations, Buchholz (2005) and Haufler (2004) provide equivalence results between the redistributional effect of contributions and the redistributional effect of premiums combined with an appropriate adjustment of income taxation.

By deriving an analogous equivalence result in an equilibrium model, our analysis underpins their findings.

Fehr and Jess (2006) and Schubert and Schnabel (2009) simulate computable general equilibrium models to assess the economic implications of introducing various variants of a premium system. Kemnitz (2013) compares the effect of a contribution-based financing scheme and a premium-based financing scheme on competition of health insurers in a duopoly model. Gouveia (1997) studies an abstract model in which the population votes to determine the level of state provision of health services. Breyer (2001) applies this model to compare the premium-based Swiss health insurance system to the contribution-based German health insurance system. His analysis predicts higher health benefits and consequently higher health costs in a contribution-based system. Kifmann (2005) analyzes a model in which the population democratically votes on its health insurance system. He shows that a contribution-based system is likely to arise in a democracy. In a similar vein, Hansen and Keiding (2002) analyze welfare differences between provision of health insurance through a competitive market and a system with compulsory insurance but democratic choice of the level of health benefits. They show that under some conditions the competitive market may outperform the compulsory, democratic system.

The rest of this chapter is organized as follows. In Section 4.2 we provide a formal description of the health insurance market. In Section 4.3 we prove existence of equilibrium and analyze changes in policy parameters and the distribution of health and income. In Section 4.4 we compare the current health insurance system to a premium-based health insurance system and study welfare-maximizing fee schedules. Section 4.5 discusses extensions of the model, and Section 4.6 concludes.

4.2 Model

In the health insurance market a population of customers purchases health insurance contracts from either of two insurers: a private health insurer and a public health insurer.

4.2.1 Population

The population consists of a unit mass continuum of customers. Every customer is characterized by his health type h and his income type e; a high value of h corresponds to a good state of health. Types are distributed according to a distribution with distribution function $F(h,e)$ with compact and connected support $\mathcal{H} \times \mathcal{E} = [0,\bar{h}] \times [0,\bar{e}]$. The distribution has a strictly positive, twice differentiable density $f(h,e)$. We assume that health type and income type are affiliated.[6] Associated with a customer's health type h are health costs $c(h)$ where $c(\cdot)$ is a continuous, positive, and strictly decreasing function. Every customer is endowed with a strictly increasing, strictly concave, and twice continuously differentiable utility function $u(w)$ over wealth w which is, for an uninsured customer, the difference between income and health costs

$$w = e - c(h).$$

Customers' health and income type are observable by all agents in the market.

[6] Affiliation is a strong form of positive correlation and widely used in Auction Theory, cf. Milgrom and Weber (1982).

4.2.1.1 Choice of Insurance

Health insurance is compulsory, i.e., every customer must choose a contract. A customer has to insure with the public insurer if his income is less than the *opt-out threshold* K_1; otherwise, he is eligible to choose between private and public insurance. Neither insurer is allowed to exclude customers from its service, that is, both insurers have to offer a contract to every eligible customer.

4.2.1.2 Contracts

A health insurance contract is defined by its fee $p(h,e)$ and the maximum monetary reimbursement of health costs, *benefit level* L. We assume that both insurers offer the same fixed benefit level L.[7] Denote by $\mathscr{C}(h,e)$ the set of contracts available to a type-(h,e) customer. Upon signing a contract, a customer pays $p(h,e)$ and receives a monetary reimbursement of $\min(c(h),L)$, the *health benefit*.[8] For a fixed contract, we refer to $\min(c(h),L) - p(h,e)$ as the customer's *net benefit*. Formally type-(h,e)'s problem is to choose $(L, p(h,e))$ such that

$$(L,p(h,e)) \in \underset{(L,p)\in\mathscr{C}(h,e)}{\arg\max}\ u(e + \min(L - c(h), 0) - p(h,e)). \tag{1}$$

4.2.2 Public Health Insurance

The public health insurer (PU) charges its customers a fixed percentage α, the *contribution rate*, of their income. Above income K_2, however, the fee remains constant; we refer to K_2 as the *contribution cap*. Therefore, the public fee is

$$p_{\text{PU}}(e) = \alpha \min(K_2, e). \tag{2}$$

PU's fee does not depend explicitly on the customer's health type. PU commits to operate with a balanced budget, i.e., it equates revenues and expenditures. Formally, PU's objective is to set $\alpha \in [0,1]$ such that

$$\alpha\ \mathbb{E}[\min(K_2,e)\mathbf{1}_{\{PU(\alpha)\}}] = \mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU(\alpha)\}}], \tag{3}$$

where $\{PU(\alpha)\}$ denotes the set of PU customers, and the expectation is taken with respect to $F(\cdot)$.

---

[7] Thomson et al. (2002) note that the difference in the benefit level between public and private insurance is small. Whereas private insurance is more generous with respect to some treatments such as dental care, public insurance is more generous with respect to others, for example, psychotherapy. See also Section 4.5 for a discussion what happens if we relax this assumption.

[8] The minimum function allows us to model overinsurance, i.e., the case when $L > c(h)$.

4.2.3 Private Health Insurance

The private health insurer (PR) charges each customer a fee $p_{\text{PR}}(h,e)$ that may depend on the customer's health type and income. We call the function $p_{\text{PR}}(\cdot)$ PR's *fee schedule*. PR is obliged to set a fee below αK_2, i.e., $p_{\text{PR}}(h,e) \leq \alpha K_2$ for all health and income types. We refer to a fee $p_{\text{PR}}(h,e)$ satisfying this requirement as *feasible fee* and to $p_{\text{PR}}(\cdot)$ as *feasible fee schedule*. PR aims at maximizing profit, i.e., fees collected minus health benefits. Thus, PR's objective is to choose $p_{\text{PR}}(\cdot)$ such that

$$p_{\text{PR}}(\cdot) \in \underset{p(\cdot)\text{ feasible}}{\arg\max} \; \mathbb{E}[(p(h,e) - \min(L, c(h))) \; \mathbf{1}_{\{PR(\alpha)\}} \;], \tag{4}$$

where $\{PR(\alpha)\}$ denotes the set of PR customers.[9]

4.2.4 Remarks

In contrast to PR, PU does not maximize profits but aims at balancing its budget. Legally, public insurers in Germany are not private companies but "public bodies", see Jacobs and Schulze (2004), or as Bauhoff (2012) notes, "*(...) sickness funds are non-profit institutions (...)*". Thus, one interpretation of a budget-balancing PU is that it is an institution set up by a benevolent regulator to cover the population's health benefits at minimal cost. Alternatively, we can interpret PU as a proxy for an entire competitive market of public insurers.

To single out the effect of the organizational structure of the health insurance market, we assume that insurers observe customers' characteristics, in particular their health type, thereby shutting down confounding channels such as adverse selection as a result of private information. We believe that this assumption is reasonable for two reasons. First, (private) insurers require potential customers to fill in binding questionnaires about their medical history. Second, insurers can draw on internal statistics to estimate precisely the likelihood that a customer falls sick with a certain disease.[10] For a somewhat critical view on the practical importance of asymmetric information in health insurance markets see Pauly (1984) and Siegelman (2004). Both authors also emphasize the importance of studying the organizational and institutional setting of health insurance markets over the implications of asymmetric information. See Breyer (2004) for a similar conclusion in the realm of European health insurance markets. On the empirical side, the finding of asymmetric information in insurance markets is mixed. For a literature overview see Siegelman (2004), Cohen and Siegelman (2010), and Einav et al. (2010). These studies conclude that the importance of asymmetric information varies significantly across markets. In a similar spirit, see also Finkelstein and Poterba (2004).

---

[9] Note that PU's and PR's objective are only well-defined if $\{PU(\alpha)\}$ and $\{PR(\alpha)\}$ are measurable. In the following, we abstract from these measurability issues and restrict attention to cases where $\{PU(\alpha)\}$ and $\{PR(\alpha)\}$ are measurable.

[10] In Section 4.5, we study the case when insurers observe only a health signal which is correlated to the true health type.

As customers decide on their insurance contract after observing their type, customers face no uncertainty about future income and health costs. Thus, the motive to insure is imposed by regulation, more precisely, the fact that health insurance is obligatory. This assumption simplifies the analysis and allows for a clean investigation of redistribution streams.

4.3 Equilibrium

4.3.1 Timing and Equilibrium Concept

Agents interact on the health insurance market in two periods. In the first period, PU and PR simultaneously devise fees for all customer types; in the second period, every customer chooses from his set of contracts. We are now in the position to define an equilibrium of the health insurance market.[11]

Definition 1. *An equilibrium of the health insurance market is a tuple* $\left(\alpha^*, p^*_{PR}(\cdot), (L^*, p^*(h,e))\right)$ *such that*

(i) α^* *satisfies* (3) *given* $p^*_{PR}(\cdot)$ *and* $(L^*, p^*(h,e))$,

(ii) $p^*_{PR}(\cdot)$ *solves* (4) *given* α^* *and* $(L^*, p^*(h,e))$,

(iii) $(L^*, p^*(h,e))$ *solves* (1), *for all* $\alpha \in [0,1]$, *feasible* $p_{PR}(\cdot)$, *and* h, e.

4.3.2 Voluntary Health Insurance

Before establishing equilibrium existence, we study the health insurance market when relaxing the assumption that health insurance is obligatory. In this case, instead of purchasing contracts from PU and PR, every customer may choose to be uninsured and bear health costs himself. As is not uncommon in models of insurance, voluntary insurance leads to a complete unraveling of the health insurance market:

Proposition 1. *For any positive benefit level there exists no equilibrium in the health insurance market if health insurance is voluntary.*

We provide intuition for Proposition 1; formal details of this and all following proofs can be found in the Appendix. If health insurance is voluntary, a customer is willing to insure only if a contract offers him a net benefit, that is, if the difference between health benefits and fee is positive. All customers whose contract set contains only contracts with negative net benefit decide to remain uninsured. A net benefit for the customer translates one-to-one into a loss in profits for the insurer. Hence, only customers who inflict a (weak) loss on the health insurer choose to be insured. PR can avoid the loss on most parts of the population as it can finetune its contract to customer's health and

---

[11] Besides the measurability issue mentioned above, our equilibrium concept coincides with the definition of pure-strategy subgame-perfect Nash equilibrium.

income. PR may only incur a loss if the upper bound on its fee binds, restricting PR's ability to deter unprofitable customers. This might occur for rich and unhealthy customers. PU is less flexible because it only discriminates along the income dimension. As a consequence, for positive benefit level, PU cannot avoid a loss on comparatively poor and unhealthy customers. These customer want to insure; since their income is sufficiently low, PR can deter them from insuring privately. Thus, they insure with PU. As a result, PU is unable to run a balanced budget, and, hence, there exists no equilibrium.

To sum up, if health insurance is voluntary, customers who are attractive from the insurers' perspective remain uninsured, leaving insurers with unprofitable customers. This makes health insurance non-viable. Note that in contrast to Rothschild and Stiglitz (1976), market breakdown is not a result of asymmetric information since customer types are observable but purely a consequence of the organizational structure of the health insurance market. Proposition 1 provides a rationale for why health insurance is obligatory in Germany and more generally in many health insurance markets.

4.3.3 Equilibrium Existence

Retaining obligatory health insurance, we prove existence of equilibrium in the health insurance market. We proceed backward from the second period, first studying customers' optimal insurance choice.

Lemma 1. *Given any contribution rate set by PU and any feasible fee schedule of PR, it is optimal for customers to choose the insurance which offers the contract with the lowest fee.*

Customers whose income is below the opt-out threshold can only choose PU's contract. All other customers have the choice between PU and PR. As the utility function is strictly increasing, every customer chooses the contract which offers him the largest net benefit. Since the benefit level is fixed and equal for PU and PR, it is the contract's fee that determines the net benefit and, thereby, its attractiveness for customers.

In the following, we assume that customers choose PR when they are indifferent, i.e., if both insurers charge the same fee.[12]

Having determined the population's optimal insurance choice, we analyze PR's optimal fee schedule. We call a customer profitable if, for a given contribution rate, the fee PU charges exceeds his health benefits; otherwise, we call the customer unprofitable.

Lemma 2. *Given customers' optimal choice and any contribution rate set by PU, it is optimal for PR to set its fee equal to PU's fee if a customer is profitable and to set the highest possible fee if a customer is unprofitable.*

For a profitable customer, PR faces the trade-off between attracting the customer and charging a high fee. If PR's fee exceeds PU's fee, the customer rejects PR's contract and chooses PU. If PR's fee is strictly less than PU's fee, PR can increase profits by increasing

[12] See also the remarks following Theorem 1.

its fee slightly without losing the customer. Hence, it is optimal for PR to set its fee exactly equal to PU's fee for all profitable customers.

If a customer is unprofitable and PR sets a fee below PU's fee, PR incurs a loss. Because PR may not reject customers, PR tries to deter unprofitable customers by setting its fee as high as possible. Note that PR may not deter all unprofitable customers because of the upper bound on its fee.

In contrast to PU, PR sets a flexible fee and discriminates based on both health and income. The above argument shows that PR exploits its greater flexibility to cream skim all profitable customers with sufficiently high income, i.e., with income exceeding the opt-out threshold. In fact, if it would not be for the opt-out threshold, PR could cream skim all profitable customers in the population which would make it impossible for PU to run a balanced budget. Hence, the opt-out threshold is essential for the existence of equilibrium in the health insurance market.

This observation motivates the following assumption which we maintain throughout this chapter.

Assumption 1. *(Viable health insurance market.) The aggregate income of customers with income below the opt-out threshold and below the contribution cap exceeds the entire population's health benefits:*

$$\mathbb{E}[\min(L, c(h))] < \mathbb{E}[e\mathbf{1}_{\{e<\min(K_1,K_2)\}}].$$

Roughly, Assumption 1 says that the total income of all customer who must insure with PU covers the health costs of the whole population. It guarantees that the population structure is such that, at least potentially, PU can run a balanced budget. There are several reasons why Assumption 1 may be satisfied; some of which of may lie under the direct control of an exogenous regulator (benefit level, opt-out threshold, contribution cap) some of which may not (health costs). In particular, Assumption 1 holds if the benefit level is sufficiently low or the opt-out threshold is sufficiently high. With this assumption in place, we obtain the following theorem:

Theorem 1. *Assume that the health insurance market is viable. Then the health insurance market has an equilibrium. Furthermore, the equilibrium contribution rate α^* is unique if customers and PR behave as described in Lemma 1 and Lemma 2.*

The proof of Theorem 1 relies on the intermediate value theorem. The key step is to establish continuity of PU's objective in the contribution rate. See the Appendix for details.

In equilibrium, customers with income below the opt-out threshold choose PU. Above the opt-out threshold customers who are profitable insure with PR; unprofitable customers insure with PU. However, all customers, profitable or unprofitable, with in-

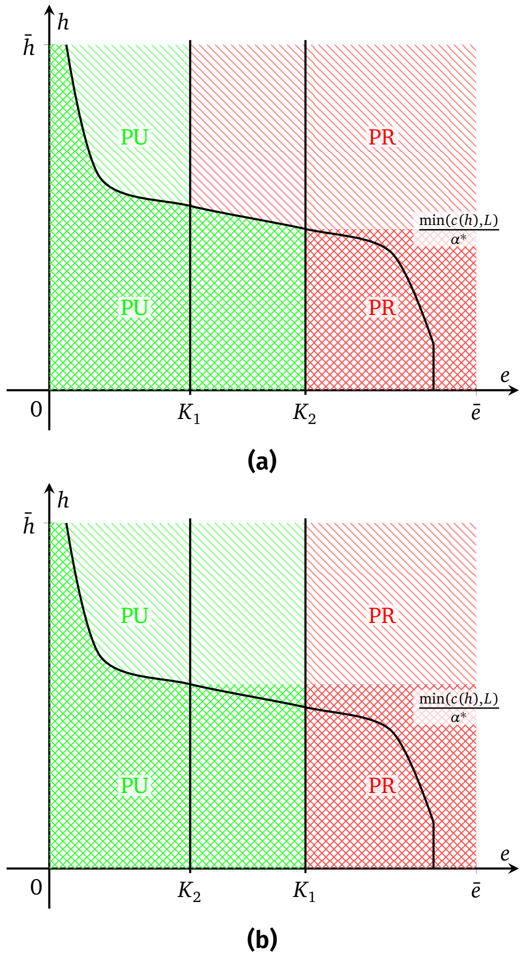

Figure 1. Customers' insurance choice by customer type for the cases (a) $K_1 < K_2$ and (b) $K_2 < K_1$. The green area is composed of PU customers, the red area of PR customers. The cross-hatched area represents unprofitable customers, and the area with diagonal stripes profitable customers.

come above the contribution cap and above the opt-out threshold choose PR. See Figure 1 for a graphical illustration.

Interestingly, independent of their insurance choice, all customers pay the amount they would pay if they insured with PU. That is, PR's fee is coupled to PU's fee in equilibrium. Intuitively, its monopolistic power in the private sector allows PR to charge customers a fee that makes them indifferent to choosing their outside option which is insuring with PU.[13] Attracting profitable customers entails two positive effects for PR: First, there is an immediate gain in profits. Second, if PU loses profitable customers to PR, PU has to increase the contribution rate which leads to a higher fee for all customers. This in turn allows PR to increase fees for all its customers as their outside option has

[13] We study the case of competing private insurers in Section 4.5.

become less attractive. In fact, note that if profitable customers with income above the opt-out threshold would collectively choose to insure with PU instead of PR, PU could adjust the contribution rate downward which would lead to a lower fee for the entire population. Intuitively, PR prevents this by slightly undercutting PU's fee.

What are the redistributional effects of the health insurance market? Profitable customers with income below the opt-out threshold subsidize all unprofitable customers of PU. Furthermore, the relative profitability of these two customer groups determines the fee for the entire population through their effect on the contribution rate. The surplus of profitable customers with high income above the opt-out threshold is transformed one-to-one into a profit for PR and is lost for the population. PR may incur a loss on unprofitable customers with income above the contribution cap and above the opt-out threshold. However, as a consequence of the organizational structure of the health insurance market, PR obtains an overall profit: PU runs a balanced budget; relative to PU, PR attracts customers with higher income. As health and income are positively correlated, a higher income entails also a better health type. Thus, PR draws upon a more lucrative part of the population and earns positive profits. See the proof of Theorem 1 for details.

A couple of technical remarks are in order. First, as can be seen from the proof of Theorem 1, the assumption that health and income are affiliated is not required for the existence of equilibrium.

Second, note that PU's contribution rate is only unique given the behavior of customers and PR. However, customers indifference behavior is not unique. Our specification that customers choose PR if they are indifferent resolves existence issues which stem from profitable customers with income exceeding the opt-out threshold: If these customers would choose PU when they are indifferent, PR would like to set a fee arbitrarily close but not equal to PU's fee. However, one could imagine other, plausible specifications for unprofitable customers with income above the contribution cap and above the opt-out threshold. For example, these customers could join PU. This could be justified as follows. Bauhoff (2012) shows that insurers do not not only cream skim explicitly based on prices but also implicitly, for example, by signaling poor service quality through delayed responses. A general theme of our model is that PR is more flexible than PU, thus, it is plausible to assume that PR is also more successful in deterring customers implicitly. Even if PR is only slightly better at deterring customers implicitly, all profitable customers with income above the opt-out threshold join PR, and all unprofitable customers with income above the opt-out threshold join PU. For this and other specifications an analogous analysis applies.

Relatedly, PR's optimal fee schedule may not be unique (even on a set with positive measure): In order to deter unprofitable customers with income between the opt-out threshold and the contribution cap, PR can set any fee that exceeds PU's fee. Note, however, that this does not change customers decisions and thus the equilibrium contribution rate is the same as under Lemma 2. Furthermore, our specification is particularly robust to tremble-like errors in the behavior of customers.

4.3.4 Comparative Statics in Policy Parameters

Having established existence of equilibrium, we analyze how changes in the opt-out threshold, the contribution cap, and the benefit level affect the equilibrium in the health insurance market. As these three parameters might be controlled by an exogenous regulator, we refer to them as “policy” parameters.

Proposition 2. *An increase of the opt-out threshold decreases the contribution rate.*

First, consider the case where both the old opt-out threshold and the new opt-out threshold are below the contribution cap. Recall that PR cream skims customers in the population with income between the opt-out threshold and the contribution cap, i.e., profitable customers in this range insure with PR, whereas unprofitable customers in this range insure with PU. An increase in the opt-out threshold limits PR’s possibility to cream skim profitable customers. Some of the profitable customers are forced to insure with PU under the new opt-out threshold. No additional unprofitable customers join PU because they insured with PU already under the old opt-out threshold. Thus, all new PU customers are profitable which allows PU to adjust the contribution rate downward.

Next, consider the case where the old and the new opt-out threshold lie above the contribution cap. In this case, all customers with income below the the opt-out threshold insure with PU; customers with income above the opt-out threshold insure with PR. An increase in the opt-out threshold forces additional customers to insure with PU. In contrast to the first case, some of these customers might be unprofitable. However, compared to existing PU customers, the new customers have higher income. As income and health are correlated, a higher income entails (on average) also a better health type. These two factors allow PU to decrease the contribution rate.

Surprisingly, limiting their choice benefits customers in that it decreases the contribution rate and thus their fee. As PR’s fee is coupled to PU’s fee in equilibrium not only PU customers benefit from the lower contribution rate but all customers do. Intuitively, with a higher opt-out threshold, more profitable customers are forced to insure with PU rather than being cream skimmed by PR. These profitable customers’ surplus is redistributed to all other customers in the population (including themselves) rather than translated into a profit for PR. PR’s profits decrease because PR loses profitable parts of the population to PU and has to charge a lower fee to attract customers. Observe that from the customers’ perspective, it would be desirable to set the opt-out threshold as high as possible, essentially eliminating PR from the market. This provides a theoretic foundation for the arguments in Breyer (2004) to increase compulsory membership in public insurance to the entire population.

Proposition 3.

(i) *If the contribution cap is above the opt-out threshold, a decrease of the contribution cap to a level that is still above the opt-out threshold decreases the contribution rate.*

(ii) *If the contribution cap is below the opt-out threshold, a decrease of the contribution cap increases the contribution rate.*

Consider first the case where the old and the new contribution cap lie above the opt-out threshold. In this case, lowering the contribution cap decreases the range of income in which PR can cream skim customers because PR faces a lower upper bound on its fee. As a result, PR can deter fewer unprofitable customers from its service. Furthermore, PR does not attract any new profitable customers. PU loses unprofitable customers while retaining all profitable customers. Also, note that the fees that all remaining PU customers pay are not reduced by the change since these customers have income below the contribution cap. Therefore, PU can adjust the contribution rate downward.

Intuitively, after the decrease, profitable customers with income below the opt-out threshold subsidize a smaller number of unprofitable customers which allows for a decrease in the contribution rate. Consequently, all customers pay less, and PR's profits decrease.

Now consider the case where both the old and the new contribution cap lie below the opt-out threshold. In this case, customer sets do not change through the decrease of the contribution cap. However, PU is forced to reduce its fee for those customers for whom the old contribution cap was binding. To compensate this loss, PU has to adjust the contribution rate upward. Thus, customers' fees increases, and PR's profits increase.

Proposition 4. *An increase of the benefit level increases the contribution rate.*

An increase in the benefit level affects PU in two ways. First, existing PU customers for whom the old benefit level was binding become less profitable. Additionally, if the opt-out threshold is below the contribution cap, there is an income range where PR cream skims. Some of the customers with income in this range are profitable under the old benefit level but become unprofitable under the new benefit level. Under the new benefit level, PR deters these customers who thus join PU. As a result of these two effects, PU has to adjust the contribution rate upward to cover the increased health benefits of its customers. The effect on customers' utility is twofold. On the one hand, all customers face a higher fee; on the other hand, some customers enjoy more health benefits. Accordingly, PR can charge a higher fee but also needs to cover higher health benefits.

4.3.5 Structural Population Changes

In this section we study how structural changes in the population's health and income affect the health insurance market. To this end, we analyze two different changes in the underlying distribution of health and income: a systematic improvement of health and income and an increased correlation between health and income.

4.3.5.1 Systematic Improvement of Health and Income

First, we investigate the effect of a systematic improvement of the population's income and health on the contribution rate and PR's profit.[14]

Intuitively, as the population's income and health improve, the population should spend a lower percentage of its income on health insurance given that the benefit level stays constant. Indeed, if the entire population would insure with the budget-balancing PU, the contribution rate could be adjusted downward. However, to account for the precise organizational structure of the market a more thorough analysis is needed.

We start by analyzing how customer sets change as the distribution changes. It is instructive to divide the population into four subgroups and study the effect of a systematic improvement on each of these subgroups separately.

Let $\{PU^+(\alpha^*)\}$ be the set of profitable PU customers and $\{PU^-(\alpha^*)\}$ the set of unprofitable PU customers. Analogously, let $\{PR^+(\alpha^*)\}$ be the set of profitable PR customers and $\{PR^-(\alpha^*)\}$ the set of unprofitable PR customers.

Profitability and unprofitability are defined relative to the original distribution and the corresponding contribution rate α^*.

First, consider the effect on the subgroup of unprofitable PU customers, $\{PU^-(\alpha^*)\}$. As health and income improve, PU's profits from this subgroup unambiguously increase: customers who remain in the group even after the improvement are less unprofitable than before; additionally, some unprofitable customers leave $\{PU^-(\alpha^*)\}$ to join one of the other subgroups.

Second, consider how PR's profit is affected on the set of its unprofitable customers, $\{PR^-(\alpha^*)\}$. Customers remaining in the group even after the improvement are less unprofitable than before, and some unprofitable customers join $\{PR^+(\alpha^*)\}$. This effect increases PR's profit. On the other hand, there is an inflow of new unprofitable customers from $\{PU^-(\alpha^*)\}$. These customers are unprofitable before and after the change of distribution but had income lower than the contribution cap before the change and income exceeding the contribution cap after the change. This effect decreases PR's profit. Which of the two effects dominates depends on the precise change in health and income.

Third, we analyze the effect on PR's profit generated from $\{PR^+(\alpha^*)\}$: Customers remaining in the group are more profitable than before. Additionally, there is an inflow of new profitable customers from all other subgroups. Thus, PR's profit from this subgroup increases.

[14] Technically, we consider a stochastic improvement of $f(h,e)$ to a distribution with density function $\tilde{f}(h,e)$ in the sense of (multivariate) first-order stochastic dominance. See the Appendix for a definition and technical details.

Finally, consider $\{PU^{+}(\alpha^{*})\}$. Again, customers remaining in this group are more profitable than before. Also, there is an inflow of new profitable customers from $\{PU^{-}(\alpha^{*})\}$. These two effects suggest that PU's profit should increase. However, there is a countervailing effect. Profitable PU customers whose income exceeds the opt-out threshold after the change are attracted by PR which decreases PU's profit on $\{PU^{+}(\alpha^{*})\}$. Therefore, the overall change in profit depends on the exact changes in health and income.

Even though we cannot determine the sign of the change in profits on $\{PR^{-}(\alpha^{*})\}$ and $\{PU^{+}(\alpha^{*})\}$ in general, we can derive an upper bound on a potential loss. In fact, a negative change in profits on $\{PR^{-}(\alpha^{*})\}$ never outweighs the increase in profits on $\{PU^{-}(\alpha^{*})\}$. To see this, observe that the loss in profit on $\{PR^{-}(\alpha^{*})\}$ is caused by unprofitable customers switching from $\{PU^{-}(\alpha^{*})\}$ to $\{PR^{-}(\alpha^{*})\}$. Thus, the loss on $\{PR^{-}(\alpha^{*})\}$ corresponds one-to-one to a gain in profit on $\{PU^{-}(\alpha^{*})\}$. All other effects increase profit. Put differently, profit on $\{PR^{-}(\alpha^{*})\} \cup \{PU^{-}(\alpha^{*})\}$ increases. An analogous argument applies to the change in profit on $\{PU^{+}(\alpha^{*})\}$.

Our analysis reveals that the systematic improvement of health and income may affect the overall profit of PU and PR positively or negatively, depending on the precise inflow and outflow into and from $\{PU^{+}(\alpha^{*})\}$ and $\{PR^{-}(\alpha^{*})\}$. However, as noted above, the overall effect is positive. Thus, it cannot be that PU's *and* PR's profits decrease. The change in PU's profit determines whether PU adjusts the contribution rate upward or downward in response to the systematic improvement. As PR's profit is increasing in the contribution rate, this effect may reinforce or counteract the initial change in PR's profit. The following proposition summarizes our findings.

Proposition 5. *Consider a systematic improvement of the population's health and income. Then always exactly one of the following three scenarios arises:*

- *(i) If the loss in PU's profit on $\{PU^{+}(\alpha^{*})\}$ outweighs the gain in profit on $\{PU^{-}(\alpha^{*})\}$, then the contribution rate increases and PR's profit increases.*
- *(ii) If the loss in PR's profit on $\{PR^{-}(\alpha^{*})\}$ outweighs the gain in profit on $\{PR^{+}(\alpha^{*})\}$, then the contribution rate decreases and PR's profit decreases.*
- *(iii) Else the contribution rate decreases and PR may profit or lose.*

Proposition 5 sorts the wide range of possible systematic improvements of health and income into three categories according to their effect on the contribution rate and PR's profit. Given that the class of improvements we consider unambiguously increase health and income for the entire population, these categories are surprisingly manifold. In particular, there exist cases in which customers have to pay a higher percentage of their income for health insurance. Intuitively, this is because an improvement might allow PR to absorb profitable PU customers which urges PU to increase the contribution rate in order to run a balanced budget.

This observation has important implications. The current organization of the health insurance market might mitigate policy programs and campaigns targeted to improve the population's health in order to decrease the contribution rate.

4.3.5.2 Increase in Correlation Between Health and Income

Motivated by empirical studies which document an increase in correlation between health and income, see e.g. Deaton and Paxson (1998), we investigate how changes in the correlation affect the health insurance market. For a meaningful comparison of the correlation of two health-income distributions, we consider health-income distributions which are ranked according to the supermodular order and thus have the same marginal distributions of health and income.[15] Specifically, start with a distribution f and consider a distribution g that is larger than f in the supermodular order. For the case when the opt-out threshold exceeds the contribution cap, we obtain the following clear-cut result.

Proposition 6. *If the opt-out threshold exceeds the contribution cap, an increase in the correlation between health and income increases the contribution rate.*

To gain intuition for our result, it is instructive to decompose the transition from f to g into several substeps. Consider the two-dimensional space of health and income types. Start with the health-income distribution f. Now fit a rectangle into the health income space and consider a transformation that shifts probability mass from the bottom right corner of the rectangle to the bottom left corner and the same probability mass from the upper left corner to the upper right corner. This transformation increases correlation between health and income while keeping the marginal distributions of health and income fixed. Intuitively, we can construct g from f by applying several of these transformations to f.[16]

If the correlation-increasing transformation is such that all four corners of the rectangle lie within within the set of PR customers, PU is unaffected and the contribution rate remains the same. Similarly, in case the four corners of the rectangle lie within the set of PU customers, PU does not need to adjust the contribution rate because the marginal distributions of health and income conditional on being a PU customer are unchanged. Lastly, consider the case when the left corners of the rectangle are in the set of PU customers whereas the right corners of the rectangle lie within the set of PR customers. As a consequence of this transformation, the income distribution of PU customers is not altered, however, the health distribution of PU customers worsens. Therefore, PU has to increase the contribution rate to run a balanced budget. Taking all three cases together, we see that PU increases the contribution rate if correlation between health and income increases. From a broader perspective, PU customer have comparatively low income, whereas PR customer have comparatively high income. If

[15] See the Appendix for a formal definition.
[16] Cf. Meyer and Strulovici (2015) for a way to make this intuition rigorous.

the correlation between health and income increases, PU's low-income customers have also a worse health type which forces PU to increase the contribution rate.

If the opt-out threshold lies below the contribution cap, there exists an income range where PR cream skims. Graphically, PU and PR customers are not separated any longer by a single cutoff in the income dimension. Thus, we have to consider additional correlation-increasing transformations. Consider the transformation where only the upper right corner of the rectangle is in the set of PR customers; all other corners lie in the set of PU customers. For this transformation there are two conflicting effects. On the one hand, the income distribution of PU customers worsens. On the other hand, unprofitable customers leave PU. It depends on the distribution f which of the two effect dominates, and, consequently, whether PU adjusts the contribution rate upward or downward. All other transformations entail a decrease in the contribution rate. Overall, in this case, it depends on the specific increase in correlation and on how much weight is put on which transformation whether PU adjusts the contribution rate upward or downward.

4.4 Applications

We apply our model to address two policy questions. First, we study how customer welfare changes if the health insurance market changes from the current contribution-based system to a premium-based system. We identify the population subgroups that benefit from such a change and the population subgroups that suffer. Second, we characterize the theoretically welfare-optimal fees. Understanding the properties of welfare-optimal fees yields additional insights into how to adjust the organization of the health insurance market to improve customer welfare.

4.4.1 Health Premia

In recent years, discussions to change the health insurance market have centered around two ideas. First, an abolishment of the difference between private and public insurers. Second, a change from an income-dependent, contribution-based fee schedule to a premium-based fee schedule, i.e., a schedule in which fees do not depend on the customer's income and health.

We adjust the model outlined in Section 4.2 to accommodate these two features of a premium-based health insurance market. Subsequently, we apply our two models to compare the premium-based health insurance market with the contribution-based health insurance market. To make the two models comparable, we alter only the insurance provision sector and leave all other characteristics, such as the population's health and income distribution and the customers' objective, unchanged.

In the premium-based health insurance market any customer must insure with either of two identical premium insurers, henceforth PM_i, $i \in \{1,2\}$.[17] Customers can choose freely between PM_1 and PM_2, independent of their income and health. PM_1 and PM_2 offer the same benefit level and face the same health costs. PM_i aims at balancing its budget by charging all its customers premium A_i, i.e.,

$$\mathbb{E}[\min(c(h), L)\mathbf{1}_{\{\mathrm{PM}_i(A_i)\}}] \;=\; \mathbb{E}[A_i \mathbf{1}_{\{\mathrm{PM}_i(A_i)\}}], \tag{5}$$

where $\{\mathrm{PM}_i(A_i)\}$ denotes the set of PM_i's customers given premium A_i.[18] The timing is unchanged. First, PM_1 and PM_2 simultaneously set their premium, then customers choose their preferred insurer. We define equilibrium analogously to Section 4.3.

Proposition 7. *There exists an equilibrium in the premium-based health insurance market. In every equilibrium, all customers pay the premium*

$$A^* \;=\; \mathbb{E}[\min(c(h), L)]. \tag{6}$$

As before, customers choose the insurance that gives them a higher net benefit. Because the benefit level is equal, customers choose the insurer with lower fee, i.e., the insurer with lower premium. Thus, if PM_i's premium is strictly lower than PM_{-i}'s premium, all customers choose PM_i. As PM_i has to cover the health benefits of the entire population and charges every customer the same premium, budget-balancing implies that the equilibrium premium is equal to the average health benefit, i.e., (6). If PM_1 and PM_2 charge the same premium, customers might split between insurers. Together, PM_1 and PM_2 bear the health costs of the entire population. Because PM_1 and PM_2 charge the same premium and each insurer charges all its customers the same premium, there is a single premium in equilibrium. Budget-balancing implies that also in this case the equilibrium premium is equal to the average health benefit.

How does the change to a premium-based system affect redistribution in the population? Proposition 7 reveals that in a premium-based system every customer pays the average health benefit of the population, independent of his income. This implies that redistribution occurs only along the health dimension, i.e., customers with a good health type subsidize customers with a bad health type. In contrast to the contribution-based system, there is no redistribution across the income dimension. Thus, the premium-based system disentangles the mixture of redistribution across health and redistribution across income which is inherent to the contribution-based system. As a consequence, we obtain the following corollary

---

[17] Analogous results hold if there are more than two premium insurers.

[18] We model premium insurers in the spirit of the public insurer in the contribution-based system. Results are virtually unchanged if we assume that premium insurers maximize profits.

Corollary 1. *There exists an income threshold such that all customers with income below the threshold have higher utility in the contribution-based system, and all customers with income above the threshold have higher utility in the premium-based system.*

Ceteris paribus, customers with higher income pay more in the contribution-based system, and customers with lower income pay less in the contribution-based system. As health benefits remain equal, customers with higher income have a lower utility and customers with lower income have a higher utility in the contribution-based system.

To maintain the current level of income redistribution, the introduction of a premium-based system would thus have to be accompanied by an adjustment of income taxation. This provides a theoretical underpinning for the observations of Breyer (2004) and Wissenschaftlicher Beirat beim Bundesministerium der Finanzen (2004).

We further assess and quantify the implications of redistribution on the population's utility in the premium-based system and in the contribution-based system. To this end, we compare the population's welfare in the two systems. We adopt the utilitarian welfare criterion, i.e., our welfare function is the sum (integral) of utilities of all customers in the population. Recall from Proposition 2 that an increase of the opt-out threshold decreases the contribution rate and consequently the fee of all customers in the population. Thus, as the opt-out threshold increases, all customers enjoy a higher utility, i.e., welfare increases. Welfare in the contribution-based system reaches its maximum once the opt-out threshold is so high that the entire population must insure with PU. We refer to this specific contribution-based system as *contribution-based system without PR*.

Proposition 8.

(i) For high levels of the opt-out threshold, the contribution-based system has higher welfare than the premium-based system.

(ii) In the premium-based system, there exists a budget-balanced income redistribution scheme such that welfare is the same as in the contribution-based system without PR.

To understand the result, observe that there are two opposing effects. On the one hand, in the contribution-based system PR makes profits and thereby extracts surplus that is not used to cover the population's health benefits. In the premium-based system neither insurer makes profits because premium insurers compete on equal grounds over the entire population.[19] This competition effect reduces welfare in the contribution-based system compared to the premium-based system. On the other hand, due to concavity of the population's utility function, it increases welfare if low-income customers pay relatively less and high-income customers pay relatively more. This redistribution effect speaks in favor of the income-dependent fee of the contribution-based system compared to the flat fee of the premium-based system.

[19] Recall, that this is true even if we assume that both premium insurers are profit-maximizing.

With a higher opt-out threshold, PR extracts less surplus which attenuates the first effect. Thus, for a sufficiently high opt-out threshold, the second effect dominates, and the contribution-based system yields higher welfare. Conversely, combining the introduction of a premium-based system with a redistribution of income from high incomes to low incomes compensates for the second effect. Consequently, the premium-based system accompanied by an appropriate income redistribution scheme yields higher welfare than the contribution-based system.

We conclude that an easy-to-implement policy recommendation to increase welfare is to increase the opt-out threshold. If the health insurance market is changed more fundamentally to a premium-based system, it would increase welfare to make the redistribution of income that is implicit in the contribution-based system explicit by, e.g., adjusting income taxation in favor of low incomes. Overall, when taking optimal policy parameter choice in the contribution-based system and income taxation in the premium-based system into account, there is no inherent advantage of one system over the other. This finding is in line with the observations of Haufler (2004) and Buchholz (2005). However, the advantage of one system over the other may lie beyond our model. Breyer and Haufler (2000) and Wissenschaftlicher Beirat beim Bundesministerium der Finanzen (2004), for example, advocate a switch to a premium-based system on grounds of smaller transaction costs, less administrative costs, more transparency, and the possibility to adopt more incentive-compatible insurance contracts.

4.4.2 Welfare-Optimal Fees

In view of Proposition 8, we are now interested in welfare-optimal fee schedules that finance a given level of health benefits. Understanding why these fee schedules maximize welfare, gives us further insights into how to adjust the health insurance market in order to increase welfare. Specifically, we consider the following problem: Health insurance is exclusively provided by a benevolent authority that chooses a fee schedule to maximize welfare subject to the constraint of providing a given benefit level. First, we explicitly derive the welfare-optimal fee schedule that may condition on both health and income.

Proposition 9. *For given health benefits, the welfare-maximizing fee schedule that may condition on health and income, $p_{opt}(h,e)$, is given by*

$$p_{opt}(h,e) \; = \; \min(c(h),L) + e - c(h) - \mathbb{E}[e - c(h)]. \tag{7}$$

Observe that $p_{opt}(h,e)$ conditions on customer's health. Intuitively, $p_{opt}(h,e)$ consists of two components. The first component charges each customer his health benefits. The second component associates to every customer his net income, i.e., his income minus health costs. If a customer's net income is high relative to the average net income, his fee is increased by the difference. If a customer's net income is low relative to the average net income, his fee is reduced by the difference. The first component guarantees that the population's health benefits are covered; the second component is a budget-balanced redistribution scheme from customers with high net income to customers with low net income. The redistribution scheme accounts for the positive effect of equating utility across customers on welfare, that stems from the concavity of customers' utility function.

In the contribution-based system, PU's fee only depends on income. Therefore, we are now interested in the properties of welfare-optimal fee schedules that only depend on income, and whether PU's fee satisfies these properties. Also, recall that welfare in the premium-based system can be increased if the introduction of a premium-based system is combined with an appropriate redistribution in the income dimension. Studying characteristics of welfare-optimal fee schedules that only depend on income, translates one-to-one into studying the characteristics of welfare-optimal income redistribution schemes in the premium-based system. Technically, we assume for the next result that the density of the distribution of health conditional on income is continuously differentiable.

Proposition 10. *For the welfare-maximizing fee schedule that only depends on income, $\hat{p}_{opt}(e)$, it holds that*

$$\frac{\mathrm{d}\hat{p}_{opt}(e)}{\mathrm{d}e} \geq 1.$$

Clearly we have that welfare under $p_{opt}(h,e)$ is higher than welfare under $\hat{p}_{opt}(e)$. However, Proposition 10 shows that $\hat{p}_{opt}(e)$ takes into account the correlation between higher income and better health. Recall that concavity of customers' utility function implies that it increases welfare to give customers similar utility levels. That is why $p_{opt}(h,e)$ redistributes income from customers with high net income to customers with low net income. An increase in customer's income by one increases his net income by more than one because a higher income is associated with better health and thus lower health costs through the correlation of health and income. Consequently, $\hat{p}_{opt}(e)$, which tries to balance net incomes across customers, has not only to neutralize the increase of income but also the positive effect on health. Hence, $\hat{p}_{opt}(e)$ increases faster than income itself.

Observe that this stands in marked contrast to PU's fee. PU's fee increases at a rate equal to the contribution rate, which is less than one, and remains constant above the contribution cap. This indicates that a reform to adjust PU's fee schedule to take the positive correlation of health and income into account has the potential of increasing welfare. On a similar note, if the introduction of a premium-based system is combined with an adjustment of income taxation to compensate for the redistribution that is lost through the abolishment of the contribution-based system, it increases welfare if the adjustment of income taxation accounts for the correlation between health and income.

4.5 Extensions

4.5.1 Health Signals

As noted in the remarks following the description of the model, we assume that insurers observe customers' characteristics, in particular their health type, to single out the effect of the organizational structure of the health insurance market. As mentioned, we think that this assumption is plausible given that (private) insurers ask their customers to complete health questionnaires and can make statistical predictions about the probability that a customer becomes sick.

Nevertheless, to demonstrate robustness of our findings to private information of customers, we consider the following variation of our model. In addition to his income and health type, each customer is characterized by a health signal. We can interpret the health signal, e.g., as the customer's answer to a health questionnaire. The customer's health signal is positively correlated with his health type.[20] Insurers observe customers' health signal but not their health type. As PU discriminates only based on income, PU is not directly affected. However, PR's ability to discriminate across health types is hampered; PR has to devise a fee schedule that only depends on income and health signal to maximize profits.

We can reproduce our findings from Section 4.3 following the same steps: As before, customers choose the insurer which charges the lower fee. For insurers, observe that we can replicate our analysis by replacing the health benefit by the expected health benefit conditional on income and health signal. Intuitively, as PR cannot observe customers' health, it estimates health using income and health signal. Due to positive correlation, high income and a favorable health signal are indicative of a good health type. Consequently, PR partially retains its ability to distinguish profitable and unprofitable customers.

4.5.2 Endogenous Health Benefits

So far we have assumed that insurers provide the same benefit level. A careful inspection of the arguments in the proof of Lemma 1 reveals that customers choose the insurance which offers the higher net benefit, i.e., health benefit minus fee. It is convenient to set benefit levels equal to focus on fees. However, note that the analysis is virtually unchanged if, say, PR provides an exogenously higher benefit level. In equilibrium PR charges higher fees such that the net benefit is unchanged. Does this conclusion remain true if PR chooses the benefit level endogenously?

To answer this question, consider the following variant of our model. PR offers customers two contracts: a simple contract reminiscent of PU's contract and a more

[20] Specifically, we assume that health signal and health type are affiliated.

elaborate contract tailored to its customers. Specifically, the first contract provides the same benefit level as PU, and the contract's fee corresponds to the highest fee which PU charges, i.e., contribution rate times contribution cap.[21] For the second contract, PR chooses benefit level and devises an income- and health-dependent fee schedule.

The equilibrium in this health insurance market parallels the equilibrium derived in Section 4.3. The sets of PU and PR customers are unchanged. PR finetunes the elaborate contract to cream skim profitable customers with income above the opt-out threshold. Unprofitable customers with income above the opt-out threshold and the contribution cap choose PR's simple contract. In equilibrium only the net benefit of the elaborate contract is determined, echoing the remarks made at the beginning of this section. Thus, without additional assumptions no prediction about the relative benefit level of PU and PR can be made.

4.5.3 Private Competition

Hitherto, we assumed that there is a single, profit-maximizing private insurer. Motivated by reforms to increase the competition between private insurers by allowing customers to transfer their accumulated savings when switching between private insurers, we are now interested in how redistribution streams change if there is competition among private insurers.

To this end, assume that in addition to PU and the population there are two private insurers, PR_1 and PR_2.[22] All insurers provide the same benefit level. PR_1 and PR_2 maximize profits by devising a fee schedule that conditions on customers' health and income. In the first period, insurers simultaneously design their contracts. In the second period, customers choose the contract that maximizes their utility. For simplicity assume that customers randomize with equal probability if they are indifferent between PR_1 and PR_2.

In equilibrium the set of PU customers is unchanged. Former PR customers split equally between PR_1 and PR_2. The fee of PR_1 customers and PR_2 customers is equal to the minimum of their health benefit and the upper bound on PRs' fee, in particular, they pay less than before. Intuitively, competition pushes the fee which PR_1 customers and PR_2 customers have to make down to the cost they impose on the insurer, i.e., their health benefit. Redistribution streams are as follows. Profitable customers with income below the opt-out threshold pay more than their health benefit and subsidize unprofitable parts of the population. Customers with income above the opt-out threshold opt out of this redistribution scheme by insuring with one of the PRs and pay an amount equal to or less than their health benefit.

[21] In the German health insurance system, private insurers have to offer this baseline contract to every customer.

[22] Analogous results hold if there are more than two private insurers.

4.6 Conclusion

This chapter develops a model of the German health insurance system for analyzing how it redistributes income between different parts of the population and evaluating policy proposals aimed at changing the system. We first show that if health insurance is voluntary, the health insurance market collapses because profitable customers decide to bear their health costs themselves. When health insurance is obligatory, we prove existence of equilibrium. In equilibrium, the private insurer tries to cream skim, i.e., it tries to attract profitable customers and deter unprofitable customers. Redistribution occurs across health and income; profitable public insurance customers with intermediate income subsidize all other public insurance customers. The surplus of profitable, high-income customers is transformed into profits for the private insurer and lost for redistribution. We argue, therefore, that, if a larger part of the population is forced to insure with the public insurer, then the entire population enjoys higher utility. Having analyzed the current German health insurance system in detail, we compare it to a system in which there is only one type of insurer, and insurers compete in premia, i.e., fees that are independent of income and health. In this premium-based system redistribution occurs only across health types. We argue that combining the introduction of a premium-based system with an appropriate adjustment of income taxation increases the population's welfare. Turning to a more careful welfare comparison of the current German system and the premium-based system, we identify two effects that determine the welfare ranking of the two systems: the redistribution effect and the competition effect. In the last part of the chapter, we study the properties of welfare-maximizing fee schedules.

We identify four avenues for future research. First, we think it would be interesting to apply our model to investigate further policy proposals to change the organization of the German health insurance market and their redistributional implications. Second, as repeatedly switching between the public and the private insurance sector is difficult in the current system, we focused on a static model in which customers choose their insurance once. Policy proposals to increase the competition and mobility between the public and the private insurance sector make it interesting to analyze a dynamic model in which customers can switch back and forth between the two sectors. Third, whereas private and public insurance are substitutes in Germany, many countries such as France and England have health insurance systems in which private insurance complements public insurance. Analyzing the organizational structure of these health insurance systems would shed light on how redistribution streams in these systems differ from the redistribution streams in the German system. Last, we believe that it would be fruitful to introduce private information more thoroughly into our model to study how it interacts with the organizational structure of the market, in particular, introducing private information could yield additional insights into the comparison of the premium-based system and the contribution-based system.

References

Atkinson, Anthony Barnes and Joseph Stiglitz (1976): "The design of tax structure: direct versus indirect taxation." *Journal of Public Economics*, 6 (1), 55–75. [119]

Bauhoff, Sebastian (2012): "Do health plans risk-select? An audit study on Germany's Social Health Insurance." *Journal of Public Economics*, 96 (9), 750–759. [122, 127]

Besley, Timothy (1989): "Publicly provided disaster insurance for health and the control of moral hazard." *Journal of Public Economics*, 39 (2), 141–156. [119]

Blomqvist, Åke and Henrik Horn (1984): "Public health insurance and optimal income taxation." *Journal of Public Economics*, 24 (3), 353–371. [118]

Blomqvist, Åke and Per-Olov Johansson (1997): "Economic efficiency and mixed public/private insurance." *Journal of Public Economics*, 66 (3), 505–516. [119]

Boadway, Robin, Manuel Leite-Monteiro, Maurice Marchand, and Pierre Pestieau (2006): "Social insurance and redistribution with moral hazard and adverse selection." *The Scandinavian Journal of Economics*, 108 (2), 279–298. [118]

Breyer, Friedrich (2001): "Income redistribution and the political economy of social health insurance: Comparing Germany and Switzerland." [120]

Breyer, Friedrich (2004): "How to finance social health insurance: issues in the German reform debate." *The Geneva Papers on Risk and Insurance Issues and Practice*, 29 (4), 679–688. [114, 117, 119, 122, 128, 135]

Breyer, Friedrich and Andreas Haufler (2000): "Health care reform: Separating insurance from income redistribution." *International Tax and Public Finance*, 7 (4-5), 445–461. [119, 136]

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston (2014): "How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program." *The American Economic Review*, 104 (10), 3335–3364. [116]

Buchholz, Wolfgang (2005): "A note on financing health-care reform: Some simple arguments concerning marginal tax burden." *FinanzArchiv: Public Finance Analysis*, 61 (3), 438–446. [119, 136]

Bünnings, Christian and Harald Tauchmann (2015): "Who opts out of the statutory health insurance? A discrete time hazard model for Germany." *Health Economics*, 24 (10), 1331–1347. [116]

Cohen, Alma and Peter Siegelman (2010): "Testing for adverse selection in insurance markets." *Journal of Risk and Insurance*, 77 (1), 39–84. [122]

Cremer, Helmuth and Pierre Pestieau (1996): "Redistributive taxation and social insurance." *International Tax and Public Finance*, 3 (3), 281–295. [118]

Deaton, Angus S and Christina H Paxson (1998): "Aging and inequality in income and health." *The American Economic Review*, 88 (2), 248–253. [116, 132]

Einav, Liran, Amy Finkelstein, and Jonathan Levin (2010): "Beyond testing: Empirical models of insurance markets." *Annual Review of Economics*, 2 (1), 311–336. [122]

Fehr, Hans and Heinrich Jess (2006): "Health premiums or health contributions? An evaluation of health care reform options in Germany." *Schmollers Jahrbuch: Journal of Applied Social Science Studies/Zeitschrift für Wirtschafts-und Sozialwissenschaften*, 126 (1), 20–57. [119]

Finkelstein, Amy and James Poterba (2004): "Adverse selection in insurance markets: Policyholder evidence from the UK annuity market." *Journal of Political Economy*, 112 (1), 183–208. [122]

Franc, Carine and Laurence Abadie (2004): "Opting out of public insurance: Is it socially acceptable?" *The Geneva Papers on Risk and Insurance Theory*, 29 (2), 115–136. [118]

Frijters, Paul, John P Haisken-DeNew, and Michael A Shields (2005): "The causal effect of income on health: Evidence from German reunification." *Journal of health economics*, 24 (5), 997–1017. [117]

Gouveia, Miguel (1997): "Majority rule and the public provision of a private good." *Public choice*, 93 (3-4), 221–244. [119]

Grunow, Martina and Robert Nuscheler (2014): "Public and private health insurance in Germany: The ignored risk selection problem." *Health Economics*, 23 (6), 670–687. [114, 116]

Hansen, Bodil O and Hans Keiding (2002): "Alternative health insurance schemes: a welfare comparison." *Journal of Health Economics*, 21 (5), 739–756. [120]

Haufler, Andreas (2004): "Welche Vorteile bringt eine Pauschalprämie für die Finanzierung des Gesundheitswesens? Einige einfache Äquivalenzresultate." *Schmollers Jahrbuch-Journal of Applied Social Science Studies-Zeitschrift für Wirtschafts-und Sozialwissenschaften*, 22 (4), 539–556. [119, 136]

Jacobs, Klaus and Sabine Schulze (2004): "Systemwettbewerb zwischen gesetzlicher und privater Krankenversicherung: Idealbild oder Schimäre." *GGW, H*, 1, 7–18. [119, 122]

Kemnitz, Alexander (2013): "A simple model of health insurance competition." *German Economic Review*, 14 (4), 432–448. [119]

Kifmann, Mathias (2005): "Health insurance in a democracy: Why is it public and why are premiums income related?" *Public Choice*, 124 (3-4), 283–308. [120]

Meyer, Margaret and Bruno Strulovici (2015): "Beyond correlation: Measuring interdependence through complementarities." *mimeo*. [132]

Milgrom, Paul R and Robert J Weber (1982): "A theory of auctions and competitive bidding." *Econometrica: Journal of the Econometric Society*, 1089–1122. [120]

Mirrlees, James A (1971): "An exploration in the theory of optimum income taxation." *Review of Economic Studies*, 38 (2), 175–208. [118]

Neudeck, Werner and Konrad Podczeck (1996): "Adverse selection and regulation in health insurance markets." *Journal of Health Economics*, 15 (4), 387–408. [118]

Newhouse, Joseph P, Mary Price, John Hsu, J Michael McWilliams, and Thomas G McGuire (2015): "How much favorable selection is left in Medicare Advantage?" *American Journal of Health Economics*, 1, 1–26. [116]

Newhouse, Joseph P, Mary Price, Jie Huang, J Michael McWilliams, and John Hsu (2012): "Steps to reduce favorable risk selection in Medicare Advantage largely succeeded, boding well for health insurance exchanges." *Health Affairs*, 31 (12), 2618–2628. [116]

Nishimura, Yukihiro (2009): "Redistributive taxation and social insurance under adverse selection in the insurance market." *International Tax and Public Finance*, 16 (2), 176–197. [118]

Panthöfer, Sebastian (2016): "Risk selection under public health insurance with opt-out." *Health Economics*, 25 (9), 1163–1181. [114, 119]

Pauly, Mark V (1984): "Is cream-skimming a problem for the competitive medical market?" *Journal of Health Economics*, 3 (1), 87–95. [116, 122]

Petretto, Alessandro (1999): "Optimal social health insurance with supplementary private insurance." *Journal of Health Economics*, 18 (6), 727–745. [119]

Rochet, Jean-Charles (1991): "Incentives, redistribution and social insurance." *The Geneva Papers on Risk and Insurance Theory*, 16 (2), 143–165. [118]

Rothschild, Michael and Joseph Stiglitz (1976): "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information." *The Quarterly Journal of Economics*, 90 (4), 629–649. [115, 118, 124]

Schickner, Benjamin and Linda Schilling (2016): "Redistributional effects of health insurance in Germany: Private and public insurance, premia and contribution rates." *mimeo*. [113]

Schubert, Stefanie and Reinhold Schnabel (2009): "Curing Germany's health care system by mandatory health premia?" *Journal of Health Economics*, 28 (5), 911–923. [119]

Shaked, Moshe and J George Shanthikumar (2007): *Stochastic orders*. Springer Science & Business Media. [152, 154]

Siegelman, Peter (2004): "Adverse selection in insurance markets: an exaggerated threat." *The Yale Law Journal*, 113 (6), 1223–1281. [122]

Thomson, Sarah, Reinhard Busse, and Elias Mossialos (2002): "Low demand for substitutive voluntary health insurance in Germany." *Croatian Medical Journal*, 43 (4), 425–432. [121]

Thomson, Sarah and Elias Mossialos (2006): "Choice of public or private health insurance: learning from the experience of Germany and the Netherlands." *Journal of European Social Policy*, 16 (4), 315–327. [114, 119]

Wasem, Jürgen, Stefan Greß, and Kieke GH Okma (2004): "The role of private health insurance in social health insurance countries." In *Social health insurance systems in western Europe*. Ed. by RB Saltmann, R Busse, and J Figuearas. New York: Open University Press. Chap. 10, 227–247. [119]

Wissenschaftlicher Beirat beim Bundesministerium der Finanzen (2004): "Nachhaltige Finanzierung der Renten- und Krankenversicherung." *Schriftenreihe des BMF 77*. [117, 119, 135, 136]

Wörz, Markus and Reinhard Busse (2005): "Analysing the impact of health-care system change in the EU member states-Germany." *Health Economics*, 14 (S1), S133–S149. [114, 117]

Appendix 4.A Appendix: Proofs

4.A.1 Proofs for Voluntary Health Insurance

Proof of Proposition 1. If health insurance is voluntary, every customer type's contract set contains, in addition to PU's and PR's contract, contract $(0,0)$, i.e., $(0,0) \in \mathscr{C}(h,e)$, $\forall(h,e)$. Assume there exists an equilibrium $(\alpha^*, p^*_{pr}, (L^*, p^*(h,e)))$. Optimal choice of customers requires:

$$u(e + \min(L^* - c(h), 0) - p^*(h,e)) \;\geq\; u(e + \min(L' - c(h), 0) - p(h,e)),$$

for all contracts $(L', p(h,e)) \in \mathscr{C}(h,e)$. As the utility function is strictly increasing, this is equivalent to

$$\min(L^*, c(h)) - p^*(h,e) \;\geq\; \min(L', c(h)) - p(h,e),$$

for all contracts $(L', p(h,e)) \in \mathscr{C}(h,e)$. In particular, with voluntary health insurance we have that

$$\min(L^*, c(h)) - p^*(h,e) \;\geq\; 0.$$

Rearranging terms gives

$$p^*(h,e) - \min(L^*, c(h)) \;\geq\; 0. \tag{8}$$

(8) implies that PU and PR incur a weak loss for every insured customer. We will now argue that (8) holds with strict inequality on a set of PU customers with positive measure. Thus, PU's equilibrium condition

$$\alpha^* \ \mathbb{E}[\min(K_2, e)\mathbf{1}_{\{PU(\alpha^*)\}}] \ = \ \mathbb{E}[\min(L, c(h))\mathbf{1}_{\{PU(\alpha^*)\}}]$$

does not hold, a contradiction. Consider the part of the population with $e < K_2$ and

$$\alpha^* \min(K_2, e) - \min(L, c(h)) \ < \ 0.$$

Because $K_2, L, c(h) > 0$ and the support of income is $[0, \bar{e}]$ this part of the population has positive measure. These customers prefer to be insured with PU over remaining uninsured. Also, PR does not want to attract this part of the population because PR would need to set $p_{PR}(h, e) \leq \alpha^* \min(K_2, e)$, incurring a loss on these customers. As $e < K_2$, PR can set $p_{PR}(h, e) \geq \alpha^* \min(K_2, e)$ to deter these unprofitable customers. Therefore, we have $p^*_{PR}(h, e) \geq \alpha^* \min(K_2, e)$. To sum up, we have argued that this part of the population will insure with PU which concludes the proof. □

4.A.2 Proofs for Equilibrium Existence

Proof of Lemma 1. Given any contribution rate $\alpha \in [0, 1]$ and any feasible choice of $p_{PR}(\cdot)$, the contract set of a customer with type (h, e) is

$$\mathscr{C}(h, e) \ = \ \begin{cases} \{(L, \alpha \min(K_2, e))\} & \text{if } e \ < \ K_1, \\ \{(L, \alpha \min(K_2, e)), \ \ (L, p_{PR}(h, e))\} & \text{else.} \end{cases}$$

It is optimal for a type-(h, e) customer to choose $(L, p^*(h, e)) \in \mathscr{C}(h, e)$ if and only if

$$u(e + \min(L - c(h), 0) - p^*(h, e)) \ \geq \ u(e + \min(L - c(h), 0) - p(h, e)),$$

for all $(L, p(h, e)) \in \mathscr{C}(h, e)$. As $u(\cdot)$ is strictly increasing, this is equivalent to

$$\min(L, c(h)) - p^*(h, e) \ \geq \ \min(L, c(h)) - p(h, e),$$

for all $(L, p(h, e)) \in \mathscr{C}(h, e)$. Because health benefits L are equal, the latter expression reduces to

$$p^*(h, e) \ \leq \ p(h, e),$$

which concludes the proof. □

Proof of Lemma 2. Fix any contribution rate α. Consider a feasible fee schedule $p(\cdot)$. Optimal customer choice, Lemma 1, implies that the set of PR customers is given by

$$\{PR(\alpha)\} \ = \ \{(h, e) \ : \ e \ \geq \ K_1, \ \ p(h, e) \ \leq \ \alpha \min(K_2, e)\}.$$

Thus, spelling out the expectation, we can rewrite PR's objective as

$$p_{\text{PR}}(\cdot) \in \underset{p(\cdot)\text{ feasible}}{\arg\max} \int_{\mathcal{E}}\int_{\mathcal{H}} (p(h,e) - \min(L, c(h))\mathbf{1}_{\{e\geq K_1,\ p(h,e)\leq\alpha\min(K_2,e)\}}(h,e)\, f(h,e)\, dh\, de.$$

Because PR's objective involves no derivatives of $p(h,e)$, we can solve it pointwise. Carefully inspecting

$$(p(h,e) - \min(L, c(h))\mathbf{1}_{\{e\geq K_1,\ p(h,e)\leq\alpha\min(K_2,e)\}}(h,e)\, f(h,e)$$

reveals that

$$p_{\text{PR}}(h,e) = \begin{cases} \alpha\min(K_2,e) & \text{if } \alpha\min(K_2,e) \geq \min(L,c(h)), \\ \alpha K_2 & \text{else,} \end{cases}$$

is an optimal policy. □

Proof of Theorem 1. Let customers' and PR's behavior be as described in Lemma 1 and Lemma 2, respectively. Fix a contribution rate α. Formally, the set of PU customers is

$$\{PU(\alpha)\} = \{(h,e) : e < K_1\}\dot{\cup}\{(h,e) : K_1 \leq e < K_2,\ \alpha e < \min(L,c(h))\},$$

and the set of PR customers is

$$\begin{aligned}\{PR(\alpha)\} =& \{(h,e) : e \geq K_1, \alpha\min(K_2,e) \geq \min(L,c(h))\} \\ &\dot{\cup}\{(h,e) : e \geq \max(K_1,K_2),\ \alpha K_2 < \min(L,c(h))\}.\end{aligned}$$

PU seeks a contribution rate α^* such that

$$\alpha^*\ \mathbb{E}[\min(K_2,e)\mathbf{1}_{\{PU(\alpha^*)\}}] = \mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU(\alpha^*)\}}].$$

Reformulating gives

$$\alpha^* = \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU(\alpha^*)\}}]}{\mathbb{E}[\min(K_2,e)\mathbf{1}_{\{PU(\alpha^*)\}}]}.$$

Define the function $T(\alpha)$:

$$T(\alpha) := \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU(\alpha)\}}]}{\mathbb{E}[\min(K_2,e)\mathbf{1}_{\{PU(\alpha)\}}]}.$$

An equilibrium contribution rate α^* corresponds to a fixed point of $T(\cdot)$. First, we show that $T(\cdot)$ is well-defined, i.e., that the denominator cannot become zero:

$$\mathbb{E}\big[\min(K_2,e)(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha e<\min(L,c(h))\}})\big] \geq \mathbb{E}\big[e\mathbf{1}_{\{e<\min(K_1,K_2)\}}\big] > 0,$$

where the last inequality follows from the assumption that $f(h,e)$ has full support and the fact that K_1 and K_2 are strictly positive.

Existence. We prove existence of a fixed point using the intermediate value theorem. First, note

$$T(0) = \frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2\}})]}{\mathbb{E}[\min(K_2,e)(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2\}})]} > 0.$$

The inequality follows from both, numerator and denominator, being strictly positive because of full support of $f(h,e)$ and $K_1, K_2, L, c(\cdot) > 0$. Second, we have

$$T(1) = \frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ e<\min(L,c(h))\}})]}{\mathbb{E}[\min(K_2,e)(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ e<\min(L,c(h))\}})]} \leq \frac{\mathbb{E}[\min(L,c(h))]}{\mathbb{E}[e\mathbf{1}_{\{e<\min(K_1,K_2)\}}]} < 1,$$

where the last inequality follows from Assumption 1. It remains to be shown that $T(\cdot)$ is continuous.[23] First, consider the numerator of $T(\cdot)$. The first addend does not depend on α, thus, we only need to check continuity of

$$g(\alpha) := \int_{\mathcal{H}}\int_{\mathcal{E}} \min(L,c(h))(\mathbf{1}_{\{K_2>e\geq K_1,\ \alpha e<\min(L,c(h))\}})f(h,e)\,de\,dh\,.$$

Fix α and $\tilde{\alpha}$, and assume without loss of generality that $\alpha > \tilde{\alpha}$.

$$\begin{aligned}
&|g(\alpha) - g(\tilde{\alpha})| \\
&\leq \int_{\mathcal{H}}\int_{\mathcal{E}} \min(L,c(h))\,\mathbf{1}_{\{K_2>e\geq K_1\}}|\mathbf{1}_{\{\alpha e<\min(L,c(h))\}} - \mathbf{1}_{\{\tilde{\alpha} e<\min(L,c(h))\}}|f(h,e)\,de\,dh \\
&= \int_{\mathcal{H}}\int_{\mathcal{E}} \min(L,c(h))\,\mathbf{1}_{\{K_2>e\geq K_1\}}\ \mathbf{1}_{\{\tilde{\alpha}<\frac{\min(L,c(h))}{e}\leq\alpha\}}\ f(h,e)\,de\,dh. \qquad (9)
\end{aligned}$$

Because $f(h,e)$ has no atoms, the integrand converges pointwise to zero as $\tilde{\alpha} \to \alpha$. Thus, by the dominated convergence theorem, (9) converges to zero as $\tilde{\alpha} \to \alpha$. Continuity of the denominator follows from an analogous argument. Hence, $T(\cdot)$ is continuous, and the existence of a fixed point follows from the intermediate value theorem.

Uniqueness. If $K_1 \geq K_2$, $T(\cdot)$ is constant in α and thus the fixed point is unique. For $K_2 > K_1$ we argue that

(i) $T(\cdot)$ is increasing left of the first fixed point,

(ii) $T(\cdot)$ is decreasing right of the first fixed point,

[23] Continuity does not follow from standard results for parameter integrals because these require that the integrand is a continuous function of α for almost all h, e.

together with the existence result above, this yields uniqueness of α^*. Note the following elementary equivalence for $a, b, c, d > 0$

$$\frac{a+c}{b+d} < \frac{a}{b} \Leftrightarrow \frac{a}{b} > \frac{c}{d}. \tag{10}$$

For (i) recall that $T(0) > 0$. Let $\tilde{\alpha}, \alpha$ be left of the first fixed point and $\tilde{\alpha} > \alpha$, then $T(\tilde{\alpha}) > \tilde{\alpha} > \alpha$. We argue that $T(\tilde{\alpha}) > T(\alpha)$:

$$\begin{aligned} &T(\tilde{\alpha}) - T(\alpha) \\ &= \frac{\int_{\mathcal{H}}\int_{\mathcal{E}} \min(L, c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_2>e\geq K_1,\ \tilde{\alpha}<\frac{\min(L,c(h))}{e}\}}) f(h,e)\, de\, dh}{\int_{\mathcal{H}}\int_{\mathcal{E}} e(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_2>e\geq K_1,\ \tilde{\alpha}<\frac{\min(L,c(h))}{e}\}}) f(h,e)\, de\, de} \\ &- \frac{\int_{\mathcal{H}}\int_{\mathcal{E}} \min(L, c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_2>e\geq K_1\}}(\mathbf{1}_{\{\alpha<\frac{\min(L,c(h))}{e}\leq\tilde{\alpha}\}} + \mathbf{1}_{\{\tilde{\alpha}<\frac{\min(L,c(h))}{e}\}})) f(h,e)\, de\, dh}{\int_{\mathcal{H}}\int_{\mathcal{E}} e(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_2>e\geq K_1\}}(\mathbf{1}_{\{\alpha<\frac{\min(L,c(h))}{e}\leq\tilde{\alpha}\}} + \mathbf{1}_{\{\tilde{\alpha}<\frac{\min(L,c(h))}{e}\}})) f(h,e)\, de\, dh}. \end{aligned} \tag{11}$$

Analyzing the indicator functions, we see that

$$\alpha \leq \frac{\int_{\mathcal{H}}\int_{\mathcal{E}} \min(L, c(h))\ \ \mathbf{1}_{\{K_2>e\geq K_1,\ \alpha<\frac{\min(L,c(h))}{e}\leq\tilde{\alpha}\}} f(h,e)\, de\, dh}{\int_{\mathcal{H}}\int_{\mathcal{E}} e\mathbf{1}_{\{K_2>e\geq K_1,\ \alpha<\frac{\min(L,c(h))}{e}\leq\tilde{\alpha}\}} f(h,e)\, de\, dh} \leq \tilde{\alpha}. \tag{12}$$

Using $T(\tilde{\alpha}) > \tilde{\alpha}$ and (12), we apply (10) to obtain $T(\tilde{\alpha}) > T(\alpha)$.

For (ii) assume that $T(\alpha)$ is not decreasing right of the first fixed point. Because $T(1) < 1$ and $T(\cdot)$ is continuous, there exist $\tilde{\alpha}, \alpha$, $\tilde{\alpha} > \alpha$, such that $\alpha > \max(\tilde{T}(\tilde{\alpha}), \tilde{T}(\alpha))$ and $\tilde{T}(\tilde{\alpha}) > \tilde{T}(\alpha)$. However, replicating the computations in (11) and (12), we observe

$$\tilde{T}(\tilde{\alpha}) - \tilde{T}(\alpha) \leq 0,$$

a contradiction. We conclude that there exists a unique contribution rate α^* that balances PU's budget.

Profits of PR. Start by observing that

$$\mathbf{1}_{\{PR(\alpha)\}} = \mathbf{1}_{\{e\geq\max\left(K_1,\ \min\left(\frac{\min(L,c(h))}{\alpha},\ K_2\right)\right)\}}$$

is increasing in h and e. Analogously,

$$\mathbf{1}_{\{PU(\alpha)\}} = \mathbf{1}_{\{e<\max\left(K_1,\ \min\left(\frac{\min(L,c(h))}{\alpha},\ K_2\right)\right)\}}$$

is decreasing in h and e. Furthermore, $\min(K_2, e)$ is increasing in e, and $\min(L, c(h))$ is decreasing in h. These observations together with the fact that $f(h, e)$ is affiliated, i.e., log-supermodular, allow us to apply the Fortuin-Kasteleyn-Ginibre (FKG) inequality to obtain:

$$\begin{aligned}
\mathbb{E}\left[\min(L,c(h))\mathbf{1}_{\{PR(\alpha)\}}\right] &\le \mathbb{E}\left[\min(L,c(h))\right]\ \mathbb{E}\left[\mathbf{1}_{\{PR(\alpha)\}}\right],\\
\mathbb{E}\left[\min(K_2,e)\mathbf{1}_{\{PR(\alpha)\}}\right] &\ge \mathbb{E}\left[\min(K_2,e)\right]\ \mathbb{E}\left[\mathbf{1}_{\{PR(\alpha)\}}\right],\\
\mathbb{E}\left[\min(L,c(h))\mathbf{1}_{\{PU(\alpha)\}}\right] &\ge \mathbb{E}\left[\min(L,c(h))\right]\ \mathbb{E}\left[\mathbf{1}_{\{PU(\alpha)\}}\right],\\
\mathbb{E}\left[\min(K_2,e))\mathbf{1}_{\{PU(\alpha)\}}\right] &\le \mathbb{E}\left[\min(K_2,e)\right]\ \mathbb{E}\left[\mathbf{1}_{\{PU(\alpha)\}}\right].
\end{aligned}$$

The four inequalities above yield

$$\frac{\mathbb{E}\left[\min(L,c(h))\mathbf{1}_{\{PR(\alpha)\}}\right]}{\mathbb{E}\left[\min(K_2,e)\mathbf{1}_{\{PR(\alpha)\}}\right]} \le \frac{\mathbb{E}\left[\min(L,c(h))\right]}{\mathbb{E}\left[\min(K_2,e)\right]} \le \frac{\mathbb{E}\left[\min(L,c(h))\mathbf{1}_{\{PU(\alpha)\}}\right]}{\mathbb{E}\left[\min(K_2,e)\mathbf{1}_{\{PU(\alpha)\}}\right]}. \tag{13}$$

In equilibrium we have

$$\frac{\mathbb{E}\left[\min(L,c(h))\mathbf{1}_{\{PR(\alpha^*)\}}\right]}{\mathbb{E}\left[\min(K_2,e)\mathbf{1}_{\{PR(\alpha^*)\}}\right]} \le \frac{\mathbb{E}\left[\min(L,c(h))\mathbf{1}_{\{PU(\alpha^*)\}}\right]}{\mathbb{E}\left[\min(K_2,e)\mathbf{1}_{\{PU(\alpha^*)\}}\right]} = \alpha^*.$$

Rearranging terms gives

$$\mathbb{E}\left[(\alpha^* \min(K_2,e) - \min(L,c(h)))\,\mathbf{1}_{\{PR(\alpha^*)\}}\right] \ge 0,$$

which concludes the proof. □

4.A.3 Proofs for Comparative Statics in Policy Parameters

Proof of Proposition 2. Consider an increase of K_1 to $\tilde{K}_1$, $K_1 \le \tilde{K}_1$. For this proof, we make the dependence of $T(\cdot)$ on K_1 explicit and write $T_{K_1}(\cdot)$. Similarly, we denote the set of PU customers by $\{PU_{K_1}(\alpha)\}$. Let the contribution rates α^* and $\tilde{\alpha}^*$ be the unique fixed points of $T_{K_1}(\cdot)$ and $T_{\tilde{K}_1}(\cdot)$ respectively. We argue that $T_{K_1}(\tilde{\alpha}^*) \ge T_{\tilde{K}_1}(\tilde{\alpha}^*)$ which implies $\alpha^* \ge \tilde{\alpha}^*$ because the fixed point is unique. Spelling out $T_{K_1}(\tilde{\alpha}^*) \ge T_{\tilde{K}_1}(\tilde{\alpha}^*)$, we obtain

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{K_1}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{PU_{K_1}(\tilde{\alpha}^*)\}}]} \ge \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}. \tag{14}$$

We distinguish two cases.

Case 1. First, let $K_1 \le \tilde{K}_1 \le K_2$. Observe that

$$\begin{aligned}
\mathbf{1}_{\{PU_{K_1}(\tilde{\alpha}^*)\}} &= \mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\le e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}}\\
&= \mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\le e<\tilde{K}_1\}} - \mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^* e\ge\min(L,c(h))\}} + \mathbf{1}_{\{\tilde{K}_1\le e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}}\\
&= \mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}} - \mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^* e\ge\min(L,c(h))\}}.
\end{aligned}$$

Hence, we can rewrite (14) as

$$\frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}-\mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^*e\ge\min(L,c(h))\}})]}{\mathbb{E}[\min(K_2,e)(\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}-\mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^*e\ge\min(L,c(h))\}})]}\ \ge\ \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}. \tag{15}$$

Similarly as in (10), we have

$$\frac{a-b}{c-d}\ \ge\ \frac{a}{c}\ \Leftrightarrow\ \frac{b}{d}\ \le\ \frac{a}{c}$$

for $c-d>0$, $a,b,c,d\ge 0$. Therefore, (15) is equivalent to

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^*e\ge\min(L,c(h))\}}]}{\mathbb{E}[\min(K_2,e)\mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^*e\ge\min(L,c(h))\}}]}\ \le\ \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}.$$

Exploiting the indicator function of term on the left side of the above inequality and the fact that $\tilde{\alpha}^*$ is a fixed point of $T_{\tilde{K}_1}(\cdot)$, we obtain

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^*e\ge\min(L,c(h))\}}]}{\mathbb{E}[\min(K_2,e)\mathbf{1}_{\{K_1\le e<\tilde{K}_1,\ \tilde{\alpha}^*e\ge\min(L,c(h))\}}]}\ \le\ \tilde{\alpha}^*\ =\ \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{PU_{\tilde{K}_1}(\tilde{\alpha}^*)\}}]}.$$

Thus, (14) holds in this case.

Case 2. Second, consider the case $K_2\le K_1\le\tilde{K}_1$. (14) becomes

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{e<K_1\}}]}{\mathbb{E}[\min(K_2,e)\mathbf{1}_{\{e<K_1\}}]}\ \ge\ \frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{e<K_1\}}+\mathbf{1}_{\{K_1\le e<\tilde{K}_1\}})]}{\mathbb{E}[\min(K_2,e)(\mathbf{1}_{\{e<K_1\}}+\mathbf{1}_{\{K_1\le e<\tilde{K}_1\}})]}.$$

Using (10), the latter inequality is equivalent to

$$\frac{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(L,c(h))\mathbf{1}_{\{e<K_1\}}f(h,e)\ \ dh\,de}{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(K_2,e)\mathbf{1}_{\{e<K_1\}}f(h,e)\ \ dh\,de}\ \ge\ \frac{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(L,c(h))\mathbf{1}_{\{K_1\le e\}}f(h,e)\ \ dh\,de}{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(K_2,e)\mathbf{1}_{\{K_1\le e\}}f(h,e)\ \ dh\,de}. \tag{16}$$

Now, we proceed as in the proof Theorem 1 where we showed that PR's profit is positive. Note that $\mathbf{1}_{\{e<K_1\}}$ is a decreasing function of h,e, and that $\mathbf{1}_{\{K_1\le e\}}$ is an increasing function of h,e. Together with the affiliation of $f(h,e)$ and the monotonicity of $\min(L,c(h))$ and $\min(K_2,e)$, these observations imply, using the FKG inequality,

$$\frac{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(L,c(h))\mathbf{1}_{\{K_1\le e\}}f(h,e)\ \ dh\,de}{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(K_2,e)\mathbf{1}_{\{K_1\le e\}}f(h,e)\ \ dh\,de}\ \le\ \frac{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(L,c(h))f(h,e)\ \ dh\,de}{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}}\min(K_2,e)f(h,e)\ \ dh\,de}$$

and

$$\frac{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}} \min(L,c(h))f(h,e)\ \ dh\,de}{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}} \min(K_2,e)f(h,e)\ \ dh\,de} \leq \frac{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}} \min(L,c(h))\mathbf{1}_{\{e<K_1\}}f(h,e)\ \ dh\,de}{\int_0^{\tilde{K}_1}\int_{\underline{h}}^{\overline{h}} \min(K_2,e)\mathbf{1}_{\{e<K_1\}}f(h,e)\ \ dh\,de}.$$

Thus, (16) holds, implying that (14) holds also in this case. This concludes the proof. □

Proof of Proposition 3. We proceed similar as in the proof of Proposition 2. Consider a decrease of K_2 to $\tilde{K}_2$, $\tilde{K}_2 \leq K_2$. For this proof, we make the dependence of $T(\cdot)$ on K_2 explicit and write $T_{K_2}(\cdot)$. Similarly, we denote the set of PU customers by $\{PU_{K_2}(\alpha)\}$. Let the contribution rates α^* and $\tilde{\alpha}^*$ be the unique fixed points of $T_{K_2}(\cdot)$ and $T_{\tilde{K}_2}(\cdot)$ respectively.

Proof of (i). First, consider the case $K_1 \leq \tilde{K}_2 \leq K_2$. We argue that $T_{K_2}(\tilde{\alpha}^*) \geq T_{\tilde{K}_2}(\tilde{\alpha}^*)$ which implies $\alpha^* \geq \tilde{\alpha}^*$ because the fixed point is unique. Spelling out $T_{K_2}(\tilde{\alpha}^*) \geq T_{\tilde{K}_2}(\tilde{\alpha}^*)$, we obtain

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{K_2}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{PU_{K_2}(\tilde{\alpha}^*)\}}]} \geq \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[\min(\tilde{K}_2,e))\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}. \tag{17}$$

Observe that

$$\begin{aligned}
\mathbf{1}_{\{PU_{K_2}(\tilde{\alpha}^*)\}} &= \mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}} \\
&= \mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<\tilde{K}_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}} - \mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}} \\
&= \mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}} - \mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}}.
\end{aligned}$$

Hence, we can rewrite (17) as

$$\frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}} + \mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}})]}{\mathbb{E}[e(\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}} + \mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}})]} \geq \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[e\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}, \tag{18}$$

where we used the indicator functions to simplify the denominators. The latter inequality is equivalent to

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}})]}{\mathbb{E}[e\mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}}]} \geq \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[e\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}$$

by (10). Exploiting the indicator function of term on the left side of the above inequality and the fact that $\tilde{\alpha}^*$ is a fixed point of $T_{\tilde{K}_2}(\cdot)$, we obtain

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}})]}{\mathbb{E}[e\mathbf{1}_{\{\tilde{K}_2\leq e<K_2,\ \tilde{\alpha}^* e<\min(L,c(h))\}}]} \geq \tilde{\alpha}^* = \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}{\mathbb{E}[e\mathbf{1}_{\{PU_{\tilde{K}_2}(\tilde{\alpha}^*)\}}]}.$$

Thus, (17) holds.

Proof of (ii). Second, consider the case $\tilde{K}_2 \leq K_2 \leq K_1$. Observe that $T(\cdot)$ is constant in α in this case. We argue that $T_{\tilde{K}_2}(\cdot) \geq T_{K_2}(\cdot)$ which implies $\tilde{\alpha}^* \geq \alpha^*$. Spelling out $T_{\tilde{K}_2}(\cdot) \geq T_{K_2}(\cdot)$, we obtain

$$\frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{e<K_1\}}]}{\mathbb{E}[\min(\tilde{K}_2,e))\mathbf{1}_{\{e<K_1\}}]} \geq \frac{\mathbb{E}[\min(L,c(h))\mathbf{1}_{\{e<K_1\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{e<K_1\}}]},$$

which holds as $\tilde{K}_2 < K_2$. □

Proof of Proposition 4. Consider an increase of L to $\tilde{L}$, $L \leq \tilde{L}$. For this proof, we make the dependence of $T(\cdot)$ on L explicit and write $T_L(\cdot)$. Similarly, we denote the set of PU customers by $\{PU_L(\alpha)\}$. Let the contribution rates α^* and $\tilde{\alpha}^*$ be the unique fixed points of $T_L(\cdot)$ and $T_{\tilde{L}}(\cdot)$ respectively. We argue that $T_{\tilde{L}}(\alpha^*) \geq T_L(\alpha^*) = \alpha^*$ which implies $\tilde{\alpha}^* \geq \alpha^*$ because the fixed point is unique. $T_{\tilde{L}}(\alpha^*)$ is given by

$$\frac{\mathbb{E}[\min(\tilde{L},c(h))\mathbf{1}_{\{PU_{\tilde{L}}(\alpha^*)\}}]}{\mathbb{E}[\min(K_2,e))\mathbf{1}_{\{PU_{\tilde{L}}(\alpha^*)\}}]}.$$

Consider the numerator of this fraction:

$$\begin{aligned}
&\mathbb{E}[\min(\tilde{L},c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(\tilde{L},c(h))\}})] \\
&= \mathbb{E}[\min(\tilde{L},c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h))\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}})] \\
&\geq \mathbb{E}[\min(\tilde{L},c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h))\}})] + \mathbb{E}[\alpha^* e\,\mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}})] \\
&\geq \mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h))\}})] + \mathbb{E}[\alpha^* e\,\mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}})].
\end{aligned}$$

To get from the second to the third line, we exploit the third indicator function. From the third to the fourth line we use $\min(L,c(h)) \leq \min(\tilde{L},c(h))$. Thus, we obtain

$$\begin{aligned}
&T_{\tilde{L}}(\alpha^*) \\
&\geq \frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h))\}})] + \mathbb{E}[\alpha^* e\,\mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}})]}{\mathbb{E}[e\,(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(\tilde{L},c(h)\}})]} \\
&= \frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h))\}})] + \mathbb{E}[\alpha^* e\,\mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}})]}{\mathbb{E}[e\,(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h)\}})] + \mathbb{E}[e\mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}}]}.
\end{aligned}$$

Observe that

$$\frac{\mathbb{E}[\min(L,c(h))(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h))\}})]}{\mathbb{E}[e\,(\mathbf{1}_{\{e<K_1\}} + \mathbf{1}_{\{K_1\leq e<K_2,\ \alpha^* e<\min(L,c(h)\}})]} = T_L(\alpha^*) = \alpha^*$$

and

$$\frac{\mathbb{E}[\alpha^* e\,\mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}})]}{\mathbb{E}[e\mathbf{1}_{\{K_1\leq e<K_2,\ \min(L,c(h))\leq\alpha^* e<\min(\tilde{L},c(h))\}}]} = \alpha^*.$$

Because

$$\frac{a+b}{c+d} = \frac{a}{c} \Leftrightarrow \frac{a}{c} = \frac{b}{d}$$

for $a, b, c, d > 0$, we conclude that

$$T_{\tilde{L}}(\alpha^*) \geq \alpha^*.$$

□

4.A.4 Proofs for Structural Population Changes

4.A.4.1 Proofs for Systematic Improvement of Health and Income

Preliminaries. We make the dependence of the expectation operator on the distribution f explicit and write $\mathbb{E}_f[\cdot]$. Throughout the proof we use the following characterization of (multivariate) first-order stochastic dominance, cf. Shaked and Shanthikumar (2007),

Theorem 2. *Consider two probability distributions over $\mathbb{R}^n$ with densities $\tilde{f}$ and f respectively. $\tilde{f}$ first-order stochastically dominates f if and only if $\mathbb{E}_{\tilde{f}}[\phi] \geq \mathbb{E}_f[\phi]$ for all increasing functions $\phi : \mathbb{R}^n \to \mathbb{R}$ for which the expectations exist.*

Let $\tilde{f}(h,e)$ first-order stochastically dominate $f(h,e)$. Other than that, we assume that $\tilde{f}(h,e)$ satisfies the same assumptions as $f(h,e)$. As before, we are interested in fixed points of the function

$$T_f(\alpha) = \frac{\mathbb{E}_f[\min(c(h),L)\mathbf{1}_{\{PU(\alpha)\}}]}{\mathbb{E}_f[\min(K_2,e)\mathbf{1}_{\{PU(\alpha)\}}]}, \tag{19}$$

where we made the dependence of $T(\cdot)$ on the distribution explicit. By the proof of Theorem 1, $T_f(\alpha)$ has a unique fixed point α^* and is increasing for $\alpha \leq \alpha^*$ and decreasing for $\alpha \geq \alpha^*$. Denote by α^* the equilibrium contribution rate associated with $f(h,e)$ and by $\tilde{\alpha}^*$ the equilibrium contribution rate associated with $\tilde{f}(h,e)$. If we argue that

$$T_{\tilde{f}}(\alpha^*) \geq (\leq) T_f(\alpha^*) = \alpha^*,$$

then we know that $\tilde{\alpha}^* \geq (\leq)\alpha^*$.

Start by observing that

$$\mathbb{E}_{\tilde{f}}[\alpha^* \min(K_2,e) - \min(c(h),L)] \geq \mathbb{E}_f[\alpha^* \min(K_2,e) - \min(c(h),L)] = 0. \tag{20}$$

because $\alpha^* \min(K_2,e) - \min(c(h),L)$ is an increasing function of (h,e). Hence, if the entire population would insure with PU, the contribution rate could be adjusted downward. Also, note that PR's profit from an (h,e)-type customer,

$(\alpha \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR(\alpha)\}}$, is an increasing function of α, for all h, e.

Formally, the decomposition outlined in the text is given by

$$\{PU^+(\alpha^*)\} = \{(h,e)|\ \alpha^* \min(K_2, e) - \min(c(h), L) \geq 0,\ e < K_1\}, \tag{21}$$
$$\{PU^-(\alpha^*)\} = \{(h,e)|\ \alpha^* \min(K_2, e) - \min(c(h), L) < 0,\ e < \max(K_1, K_2)\}, \tag{22}$$
$$\{PR^+(\alpha^*)\} = \{(h,e)|\ \alpha^* \min(K_2, e) - \min(c(h), L) \geq 0,\ e \geq K_1\}, \tag{23}$$
$$\{PR^-(\alpha^*)\} = \{(h,e)|\ \alpha^* \min(K_2, e) - \min(c(h), L) < 0,\ e \geq \max(K_1, K_2)\}. \tag{24}$$

Define $\mathbb{E}_{\tilde{f}-f}[\cdot] := \mathbb{E}_{\tilde{f}}[\cdot] - \mathbb{E}_f[\cdot]$. Consider the difference in insurers' profit under $\tilde{f}(h,e)$ and $f(h,e)$ on each set of customers (21)-(24),

$$\mathbb{E}_{\tilde{f}-f}[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PU^+(\alpha^*)\}}], \tag{25}$$
$$\mathbb{E}_{\tilde{f}-f}[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PU^-(\alpha^*)\}}], \tag{26}$$
$$\mathbb{E}_{\tilde{f}-f}[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR^+(\alpha^*)\}}], \tag{27}$$
$$\mathbb{E}_{\tilde{f}-f}[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR^-(\alpha^*)\}}]. \tag{28}$$

By (20) we have

$$(25) + (26) + (27) + (28) \geq 0. \tag{29}$$

We verify the statements about the impact of customers' movements on insurers' profit from each subgroup made in the main body of the text. Checking monotonicity of the appropriate functions and applying Theorem 2 yields (26)≥ 0, (26)+(28)≥ 0, (27)≥ 0, and (27)+(25)≥ 0.

Proof of Proposition 5.
Proof of (i). By assumption (25)+(26)$\leq$0, which is equivalent to

$$\alpha^* = T_f(\alpha^*) \leq T_{\tilde{f}}(\alpha^*),$$

hence, $\tilde{\alpha}^* \geq \alpha^*$. Furthermore, by (25)+(26)$\leq$0 and (29), we have 0$\leq$(27)+(28), i.e.,

$$\mathbb{E}_{\tilde{f}}[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR(\alpha^*)\}}] - \mathbb{E}_f[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR(\alpha^*)\}}] \geq 0. \tag{30}$$

(30) and $\tilde{\alpha}^* \geq \alpha^*$, together with monotonicity of PR's profit in α show that PR's profit increase under $\tilde{f}$.

Proof of (ii). By assumption (27)+(28)$\leq$ 0. This implies (25)+(26)$\geq$ 0, i.e.,

$$\alpha^* = T_f(\alpha^*) \geq T_{\tilde{f}}(\alpha^*),$$

therefore, $\tilde{\alpha}^* \leq \alpha^*$. Also, by (27)+(28)$\leq$ 0,

$$\mathbb{E}_{\tilde{f}}[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR(\alpha^*)\}}] - \mathbb{E}_f[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR(\alpha^*)\}}] \leq 0. \tag{31}$$

(31) and $\tilde{\alpha}^* \leq \alpha^*$ show that PR's profit decreases.

Proof of (iii). By assumption (25)+(26)≥ 0, which is equivalent to

$$\alpha^* = T_f(\alpha^*) \geq T_{\tilde{f}}(\alpha^*),$$

and hence, $\tilde{\alpha}^* \leq \alpha^*$. Furthermore, (27)+(28)≥ 0, i.e.,

$$\mathbb{E}_{\tilde{f}}[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR(\alpha^*)\}}] - \mathbb{E}_f[(\alpha^* \min(K_2, e) - \min(c(h), L))\mathbf{1}_{\{PR(\alpha^*)\}}] \geq 0. \quad (32)$$

The positive effect of the shift from f to $\tilde{f}$ on PR's profit, (32), may be mitigated by the decrease of the equilibrium contribution rate. The exact effect on PR's profit depends on the specific shift $\tilde{f}$. □

4.A.4.2 Proofs for Increase in Correlation Between Health and Income

Preliminaries. We make the dependence of the expectation operator on the distribution f explicit and write $\mathbb{E}_f[\cdot]$. Throughout the proof we use the following characterization of the supermodular order, cf. Shaked and Shanthikumar (2007).

Theorem 3. *Consider two probability distributions over $\mathbb{R}^n$ with densities g and f respectively. g is larger than f in the supermodular order if and only if $\mathbb{E}_g[\phi] \geq \mathbb{E}_f[\phi]$ for all supermodular functions $\phi : \mathbb{R}^n \to \mathbb{R}$ for which the expectations exist.*

Proof of Proposition 6. Assume that $g(h, e)$ is larger than $f(h, e)$ in the supermodular order. Other than that, we assume that $g(h, e)$ satisfies the same assumptions as $f(h, e)$. As before, we are interested in fixed points of the function

$$T_f(\alpha) = \frac{\mathbb{E}_f[\min(c(h), L)\mathbf{1}_{\{PU(\alpha)\}}]}{\mathbb{E}_f[\min(K_2, e)\mathbf{1}_{\{PU(\alpha)\}}]}, \quad (33)$$

where we made the dependence of $T(\cdot)$ on the distribution explicit. By the proof of Theorem 1, $T_f(\alpha)$ has a unique fixed point α^* and is increasing for $\alpha \leq \alpha^*$ and decreasing for $\alpha \geq \alpha^*$. Denote by α^* the equilibrium contribution rate associated with $f(h, e)$ and by $\tilde{\alpha}^*$ the equilibrium contribution rate associated with $g(h, e)$. If we argue that

$$T_g(\alpha^*) \geq T_f(\alpha^*) = \alpha^*,$$

then we know that $\tilde{\alpha}^* \geq \alpha^*$. Observe that if $K_1 \geq K_2$ we have

$$T_g(\alpha) = \frac{\mathbb{E}_g[\min(c(h), L)\mathbf{1}_{\{e<K_1\}}]}{\mathbb{E}_g[\min(K_2, e)\mathbf{1}_{\{e<K_1\}}]}. \quad (34)$$

First, consider the denominator of the latter expression

$$\mathbb{E}_g[\min(K_2,e)\mathbf{1}_{\{e<K_1\}}] \;=\; \int_0^{\bar{e}} \min(K_2,e)\mathbf{1}_{\{e<K_1\}} g(e)\;\; de \;=\; \mathbb{E}_f[\min(K_2,e)\mathbf{1}_{\{e<K_1\}}], \tag{35}$$

where the last equality follows from the fact that g and f have the same marginals. Second, we analyze the numerator of (34). Let $e' \geq e$ and $h' \geq h$, then

$$\min(c(h'),L)\mathbf{1}_{\{e'<K_1\}} + \min(c(h),L)\mathbf{1}_{\{e<K_1\}} \;\geq\; \min(c(h'),L)\mathbf{1}_{\{e<K_1\}} + \min(c(h),L)\mathbf{1}_{\{e'<K_1\}}$$

because $\min(c(h),L)$ is decreasing in h. Consequently, $\min(c(h),L)\mathbf{1}_{\{e<K_1\}}$ is supermodular, and by definition of the supermodular order we obtain

$$\mathbb{E}_g[\min(c(h),L)\mathbf{1}_{\{e<K_1\}}] \;\geq\; \mathbb{E}_f[\min(c(h),L)\mathbf{1}_{\{e<K_1\}}]. \tag{36}$$

Putting (35) and (36) together, we get

$$T_g(\alpha^*) \;=\; \frac{\mathbb{E}_g[\min(c(h),L)\mathbf{1}_{\{e<K_1\}}]}{\mathbb{E}_g[\min(K_2,e)\mathbf{1}_{\{e<K_1\}}]} \;\geq\; \frac{\mathbb{E}_f[\min(c(h),L)\mathbf{1}_{\{e<K_1\}}]}{\mathbb{E}_f[\min(K_2,e)\mathbf{1}_{\{e<K_1\}}]} \;=\; T_f(\alpha^*) \;=\; \alpha^*$$

which concludes the proof. □

4.A.5 Proofs for Applications

4.A.5.1 Proofs for Health Premia

Proof of Proposition 7. Fix any two premia A_1 and A_2 set by PM_1 and PM_2, respectively. The contract set $\mathscr{C}(h,e)$ of a customer with type (h,e) is

$$\mathscr{C}(h,e) \;=\; \{(L,A_1),(L,A_2)\}.$$

Because health benefits are equal, it is optimal for every customer to choose contract (L,A_i) with $A_i = \min(A_1,A_2)$.

Start by observing that the following is an equilibrium: $A_1 \geq A_2$, $A_2 = \mathbb{E}[\min(c(h),L)]$ and all customers choose PM_2. We now deduce more generally that the premium paid by all customers is $\mathbb{E}[\min(c(h),L)]$ in any equilibrium of the premium-based health insurance market. It is convenient to denote by $\beta(h,e) \in \{0,1\}$ customer-(h,e)'s choice of insurance, where $\beta(h,e) = 1$ means that the customer chooses PM_1, and $\beta(h,e) = 0$ means that the customer chooses PM_2. Let $(A_1^*, A_2^*, \beta^*(h,e))$ be an equilibrium of the premium-based health insurance market.

Case 1. If $A_1^* > A_2^*$, then $\beta^*(h,e) = 0$, for all (h,e). PM_2's equilibrium condition requires $A_2^* = \mathbb{E}[\min(c(h),L)]$. The case $A_1^* < A_2^*$ is symmetric.

Case 2. If $A_1^* = A_2^*$, PM_1's and PM_2's equilibrium condition requires

$$\mathbb{E}[\beta^*(h,e)\min(c(h),L)] = \mathbb{E}\left[A_1^*\beta^*(h,e)\right] \tag{37}$$

and

$$\mathbb{E}[(1-\beta^*(h,e))\min(c(h),L)] = \mathbb{E}\left[A_2^*(1-\beta^*(h,e))\right]. \tag{38}$$

Adding up (37) and (38) and using $A_1^* = A_2^*$ yields

$$\mathbb{E}[\min(c(h),L)] = A_1^* = A_2^*,$$

which concludes the proof.[24] □

Proof of Corollary 1. Comparing the income-dependent fee of the contribution-based system

$$\alpha^* \min(K_2, e)$$

to the constant fee of the premium-based system

$$\mathbb{E}[\min(L, c(h))]$$

yields the existence of a threshold $e^* \in [0, \bar{e}]$ such that for all $e < e^*$ we have $\alpha^* \min(K_2, e) < \mathbb{E}[\min(L, c(h))]$, and for all $e > e^*$ we have $\alpha^* \min(K_2, e) > \mathbb{E}[\min(L, c(h))]$. As health benefits are equal in both system customers with income $e > e^*$ enjoy a higher utility and customers with income $e < e^*$ enjoy a lower utility in the premium-based system.

We now argue that $e^* \in (0, \bar{e})$. First, observe that

$$\alpha^* = \frac{\mathbb{E}[\min(L, c(h))\mathbf{1}_{\{PU(\alpha^*)\}}]}{\mathbb{E}[\min(K_2, e)\mathbf{1}_{\{PU(\alpha^*)\}}]} \geq \frac{\mathbb{E}[\min(L, c(h))]}{\mathbb{E}[\min(K_2, e)]},$$

where the equality follows from α^* being a fixed point of $T(\cdot)$, and the inequality follows from (13). Therefore, we can conclude that

$$\alpha^* \min(K_2, \bar{e}) \geq \frac{\mathbb{E}[\min(L, c(h))]}{\mathbb{E}[\min(K_2, e)]} \min(K_2, \bar{e}) > \mathbb{E}[\min(L, c(h))].$$

Second, note that

$$\alpha^* \min(K_2, 0) = 0 < \mathbb{E}[\min(L, c(h)],$$

which concludes the proof. □

---

[24] As before, we restrict attention to equilibria where $\beta(\cdot,\cdot)$ is measurable with respect to (h,e). Note that the result still holds if customers are allowed to randomize, i.e., if $\beta(h,e) \in [0,1]$ denotes the probability that customer-(h,e) chooses PM_1.

Proof of Proposition 8. Fix a fee $p(h,e)$ for each customer type. Given this set of fees, welfare is

$$\mathscr{W}(p(h,e)) \;=\; \mathbb{E}[u(\min(c(h),L) - c(h) + e - p(h,e))].$$

Proof of (i). Set $K_1 = \bar{e}$. Recall that the fee in the contribution-based system is $\alpha^* \min(K_2, e)$, whereas it is $A^* = \mathbb{E}[\min(L, c(h))]$ in the premium-based system. As $K_1 = \bar{e}$, budget-balancing of PU implies

$$\alpha^* \;\mathbb{E}[\min(K_2, e)] \;=\; \mathbb{E}[\min(L, c(h))] \;=\; A^*.$$

To save on notation, define $\psi(h,e) = \min(c(h),L) - c(h) + e$ and note that $\psi(\cdot,\cdot)$ is increasing in both arguments. Consider the welfare difference between the premium-based system and the contribution-based system

$$\begin{aligned} &\mathbb{E}[u(\psi(h,e) - \alpha^* \;\mathbb{E}[\min(K_2,e)])] - \mathbb{E}[u(\psi(h,e) - \alpha^* \min(K_2,e))] \\ &\quad < \;\mathbb{E}\big[u'(\psi(h,e) - \alpha^* \min(K_2,e))(\alpha^* \min(K_2,e) - \alpha^* \;\mathbb{E}[\min(K_2,e)])\big], \end{aligned} \tag{39}$$

where the inequality follows from strict concavity of $u(\cdot)$. Observe that

1. $u'(\psi(h,e) - \alpha^* \min(K_2,e))$ is decreasing in (h,e) because $u'(\cdot)$ is decreasing and $\psi(h,e) - \alpha^* \min(K_2,e)$ is increasing in (h,e) as $\alpha^* \leq 1$.
2. $\alpha^* \min(K_2,e) - \alpha^* \;\mathbb{E}[\min(K_2,e)]$ is weakly increasing in (h,e).

Hence, the FKG inequality implies that (39) is bounded above by

$$\mathbb{E}\big[u'(\psi(h,e) - \alpha^* \min(K_2,e))\big]\mathbb{E}[\alpha^*(\min(K_2,e) - \mathbb{E}[\min(K_2,e)])] \;=\; 0,$$

where the last equality follows from

$$\mathbb{E}[\alpha^*(\min(K_2,e) - \mathbb{E}[\min(K_2,e)])] \;=\; 0. \tag{40}$$

Therefore, the contribution-based system with $K_1 = \bar{e}$ gives the population a strictly higher welfare than the premium-based system. Recall that welfare is increasing in K_1. Thus, for sufficiently high K_1 the contribution-based system is welfare-dominant.

Proof of (ii). Consider the income redistribution scheme that is defined by the transfer $\tau(e)$ to agent with income e, where

$$\tau(e) \;=\; \alpha^* \;\mathbb{E}[\min(K_2,e)] - \alpha^* \min(K_2,e).$$

By definition, the premium-based system together with this income redistribution scheme gives the population the same welfare as the welfare-optimal contribution-based system, i.e., the system with $K_1 = \bar{e}$. Furthermore, (40) implies that the income redistribution scheme is budget-balanced. □

4.A.5.2 Proofs for Welfare-Optimal Fees

Proof of Proposition 9. Let $A := \mathbb{E}[\min(c(h), L)]$ be the aggregate health benefits of the population. Formally, we consider the problem

$$\max_{p(h,e)} \quad \mathbb{E}[u(\min(c(h), L) - c(h) + e - p(h,e))], \tag{41}$$

$$s.t. \quad A \leq \mathbb{E}[p(h,e)]. \tag{42}$$

The Lagrangian

$$\mathbb{E}[u(\min(c(h), L) - c(h) + e - p(h,e)) + \lambda(p(h,e) - A)]$$

yields the first-order condition

$$u'(\min(c(h), L) - c(h) + e - p(h,e)) = \lambda. \tag{43}$$

Note that $u'(\cdot)$ is strictly decreasing. Solving for $p(h,e)$ and inserting into the constraint, (42), gives

$$A = \mathbb{E}\left[-u'^{-1}(\lambda) + \min(c(h), L) - c(h) + e\right].$$

Using the definition of A, we obtain

$$\lambda = u'(\mathbb{E}[e - c(h)]). \tag{44}$$

Equating (43) and (44) yields

$$u'(\min(c(h), L) - c(h) + e - p(h,e)) = u'(\mathbb{E}[e - c(h)]). \tag{45}$$

Again exploiting that $u'(\cdot)$ is strictly decreasing and after rearranging terms we obtain

$$p_{opt}(h,e) = \min(c(h), L) + e - c(h) - \mathbb{E}[e - c(h)].$$

□

Proof of Proposition 10. We start by rewriting (41) to account for the fact that the fee may not depend on h. For clarity we spell out all expectations explicitly.

$$\max_{p(e)} \int_{\mathcal{E}} \int_{\mathcal{H}} u(\min(c(h), L) - c(h) + e - p(e)) \; f(h|e) \; \mathrm{d}h \; f(e) \; \mathrm{d}e, \tag{46}$$

$$s.t. \quad A \leq \int_{\mathcal{E}} p(e) f(e) \; \mathrm{d}e. \tag{47}$$

The Lagrangian for the problem is

$$\int_{\mathcal{E}}\int_{\mathcal{H}} u(\min(c(h),L)-c(h)+e-p(e))f(h|e)\ \mathrm{d}h+\lambda(p(e)-A)f(e)\ \mathrm{d}e.$$

Using Leibniz's integral rule, we obtain the first-order condition

$$\int_{\mathcal{H}} u'(\min(c(h),L)-c(h)+e-p(e))f(h|e)\ \mathrm{d}h-\lambda\ =\ 0. \tag{48}$$

(48) defines p as an implicit function of e. Denote the left side of (48) by $G(e,p)$. Then

$$\frac{\partial G(e,p)}{\partial p}\ =\ \int_{\mathcal{H}} -u''(\min(c(h),L)-c(h)+e-p)f(h|e)\ \mathrm{d}h\ >\ 0, \tag{49}$$

where the last inequality follows from strict concavity of $u(\cdot)$. Furthermore,

$$-\frac{\partial G(e,p)}{\partial e}\ =\ \frac{\partial G(e,p)}{\partial p}+\int_{\mathcal{H}} -u'(\min(c(h),L)-c(h)+e-p)\ \frac{\partial f(h|e)}{\partial e}\ \mathrm{d}h. \tag{50}$$

Rewrite the second term on the right side of inequality (50) as

$$\int_{\mathcal{H}} -u'(\min(c(h),L)-c(h)+e-p)\ \frac{\partial \log f(h|e)}{\partial e}\ f(h|e)\ \mathrm{d}h.$$

Observe that:

1. By affiliation $\frac{\partial \log f(h|e)}{\partial e}$ is increasing in h. Indeed, we have

$$0\ \leq\ \frac{\partial^2 \log f(h,e)}{\partial e\ \partial h}\ =\ \frac{\partial^2 \log(f(h|e)f(e))}{\partial e\ \partial h}\ =\ \frac{\partial}{\partial h}\left(\frac{\partial \log f(h|e)}{\partial e}\right).$$

2. $-u'(\min(c(h),L)-c(h)+e-p)$ is increasing in h because $\min(c(h),L)-c(h)$ is increasing and $-u'(\cdot)$ is increasing by concavity.

Neglecting the argument of $-u'(\cdot)$ for convenience and applying the FKG inequality, we get

$$\int_{\mathcal{H}} -u'(\cdot)\ \frac{\partial \log f(h|e)}{\partial e}\ f(h|e)\ \mathrm{d}h \geq \int_{\mathcal{H}} -u'(\cdot)\ f(h|e)\ \mathrm{d}h\ \int_{\mathcal{H}} \frac{\partial \log f(h|e)}{\partial e}\ f(h|e)\ \mathrm{d}h. \tag{51}$$

Rewriting the second term on the right hand side of inequality (51) and using Lebesgue's dominated convergence theorem, we see that

$$\int_{\mathcal{H}} \frac{\partial \log f(h|e)}{\partial e}\ f(h|e)\ \mathrm{d}h\ =\ \int_{\mathcal{H}} \frac{\partial f(h|e)}{\partial e}\ \mathrm{d}h\ =\ \frac{\partial}{\partial e}\left(\int_{\mathcal{H}} f(h|e)\ \mathrm{d}h\right)\ =\ 0.$$

Consequently, note that

$$\int_{\mathscr{H}} -u'(\cdot) \ \frac{\partial \log f(h|e)}{\partial e} \ f(h|e) \ \mathrm{d}h \ \geq \ 0. \tag{52}$$

Applying the implicit function theorem, we conclude that

$$\frac{\mathrm{d}\hat{p}_{opt}}{\mathrm{d}e} = \frac{-\frac{\partial G(e,p)}{\partial e}}{\frac{\partial G(e,p)}{\partial p}} = \frac{\frac{\partial G(e,p)}{\partial p}}{\frac{\partial G(e,p)}{\partial p}} + \frac{\int_{\mathscr{H}} -u'(\cdot) \ \frac{\partial \log f(h|e)}{\partial e} \ f(h|e) \ \mathrm{d}h}{\frac{\partial G(e,p)}{\partial p}} \ \geq \ 1,$$

where the inequality follows from (49) and (52). □