Essays on the Strategic Response to Demand with (In-)Correct Beliefs

Inaugural-Dissertation

zur Erlangung des akademischen Grades
doctor rerum politicarum (Dr. rer. pol.)

eingereicht an der Wirtschaftswissenschaftlichen Fakultät
der Heinrich-Heine Universität Düsseldorf
von

Jana Gieselmann, M.Sc.

Mai 2025

Erstgutachter: Prof. Paul Heidhues, PhD
Zweitgutachter: Prof. Dr. Alexander Rasch

Acknowledgments

Writing this dissertation would not have been possible without the support and feedback of many people. So, while this section is highly incomplete, I am deeply grateful to anyone who has supported me along the way.

First, I thank my supervisor Paul Heidhues for his invaluable feedback, patience and continuous support. I certainly learned the most during my PhD in his courses, while co-instructing his courses, and through the many conversations we had. Lastly, I am thankful for his guidance and care during the academic job market. I have gratefully acknowledged financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation 462020252).

Second, I thank my second supervisor Alexander Rasch for his supervision and collaboration over the years. He previously supervised my Master's thesis and is also my coauthor for the third chapter of this dissertation.

Besides my supervisors, I thank Philipp Strack for his invitation to visit Yale, his feedback, and his support during the academic job market. Spending time at Yale was one of the most special experiences of my PhD, and my thesis has greatly benefited from the feedback and interactions with the faculty and fellow PhD students. Among others, I am especially grateful to Teemu Pekkarinen and Jasmina Simon for their patience in listening to and discussing my projects with me.

I thank my coauthors Johannes Muthers and Matthias Hunold for the good teamwork and fruitful discussions. Moreover, I want to thank my colleagues at DICE for their feedback, enjoyable discussions, and patience when listening to my talks—especially David Zeimentz and Simon Schulten, who always took the time to listen to my ideas, even when they were busy. Lastly, I thank my year-long office mate Sabrina Herzog, who always had an open ear and time for a coffee. As mentioned before, many more people have improved the research in this thesis through their thoughtful comments, and I am deeply grateful to all of you.

To my family and friends, I am grateful for their support along the way. Special thanks go to my mother, Ilka, who has supported me and my career over all these years. Last but not least, I am indebted to Leo, who has always believed in me—even when I did not—and who took on all other responsibilities when I needed to focus on work. I am truly thankful for his company over the years.

Contents

List of Figures

List of Tables

Introduction

In many markets, beliefs of buyers are important and the recent literature has documented that such beliefs need not be correct. In this dissertation, I shed light on how firms respond to the possibly incorrect beliefs of market participants and study the broader implications for the functioning and regulation of such markets. The different chapters of this dissertation are motivated by recent cases, for example, competition or consumer protection cases against large platforms such as dating or cryptocurrency platforms in digital markets, or antitrust cases against colluding manufacturers in traditional retail markets.

The first two chapters study (exploitative) platform markets. Platform markets are ubiquitous in our modern economy and shape how people meet and interact. For example, matching platforms have become the most common means of finding a (romantic) partner or a job. These platforms, however, have repeatedly come under scrutiny for engaging in deceptive and exploitative business models.

In Chapter 1, I examine how a matching platform, which has increasing access to user data, tailors its algorithm when — as I argue is commonly the case — incentives between the platform and users are misaligned. Users pay to be matched by the platform, while the platform makes money as long as users continue to search for partners. Contrary to the intuition that more data about users might improve matching efficiency and speed, I show that more data allows the platform to raise profits by designing a matching rule that increases search time. Restating the platform's problem as a linear programming problem allows me to characterize the optimal matching mechanism. I show that random matching is, but for knife-edge cases, suboptimal for the platform. Instead, the platform strategically lowers match quality to increase search time and thus profits, leading to unnecessary delays. In addition, the optimal matching mechanism often induces inefficient matches to leave the platform together.

Finally, I provide two explanations for why platforms adopt business models with misaligned incentives: targeted advertising and the presence of overconfident users. In principle, the platform can set up a business model that extracts the entire surplus from users by collecting high personalized fees and providing users with the socially optimal match. Under the realistic assumption, however, that users are reluctant to pay upfront but are willing to consume ads I show that an ad-based model can outperform the former business model if targeted advertising is sufficiently efficient. Alternatively, if users are overconfident about their desirability, this belief leads users to underestimate their search time. Therefore, under the pay-as-you-search business model they spend a higher amount ex post than anticipated ex ante. This, in turn, favors the prevailing business model.

In Chapter 2, I investigate the use of so-called fake profiles by platforms. One core implication of network effects is that a platform wants to convince potential users that it has a large network. In a number of prominent cases involving dating or cryptocurrency platforms, firms did so by actively engaging in activities that artificially inflate the size of the network. To study the impact on users who are either suspecting or unsuspecting, I develop a model in which a profit-maximizing platform wants to convince users of its large network size when there is uncertainty about the number of potential users. To do so, the platform uses prices and fake profiles to signal its size both when users are sophisticated and when they are unaware of the platform's ability to do so.

If — as in some real-world examples — users are naïve about the platform's ability to use fake profiles, the platform exploits users' misperceptions by using fake profiles to deceive users about the size of its network, and thus its value. By raising prices in equilibrium, the platform profits from the artificial increase in demand. When users are sophisticated on the other hand, larger platforms use fake profiles to differentiate themselves from smaller platforms. Sophisticated users correctly anticipate the platform's incentives and thus discount the perceived network size by the expected number of fake profiles. In contrast to the case of naïve consumers, platforms would benefit from a ban on fake profiles when users are sophisticated. Rather than seeing platforms commit to not using fake profiles, we observe that platforms actually hide the use of fake profiles in their terms of service. This, however, likely indicates that users are mainly naïve in these markets, which makes fake profiles profitable.

In Chapter 3, I investigate jointly with Matthias Hunold, Johannes Muthers, and Alexander Rasch how incorrect beliefs of retailers in vertically related markets affect a manufacturer's ability to collude. Evidence from cartel cases in vertically related markets indicates that manufacturers often fail to achieve collusive price increases, because unsuspecting retailers refuse to accept higher prices for the fear of being outcompeted by fellow downstream firms that (continue to) receive better offers.

To address this issue, we show how the (market) beliefs of retailers affect the market outcome when manufacturers try to collude. With two exclusive manufacturer-retailer pairs and private contracting, we show that potential strategic misinterpretations and misunderstandings by retailers are important for the feasibility of manufacturer collusion in vertically related markets. We model the retailers' (potentially incorrect) expectations about their competitors' wholesale price offers. If retailers believe collusion to be infeasible or do not foresee manufacturers' punishment strategies, it is impossible for manufacturers to collude. By contrast, if retailers anticipate the collusive strategy and condition their action on past offers, collusion becomes feasible. We introduce the property of opportunism-proofness that excludes profitable joint deviations by the collusive entity and discuss adaptive beliefs for which retailers heuristically learn about the market conduct.

To conclude, this dissertation examines how firms strategically respond to demand with incorrect beliefs. This analysis contributes to the growing field of behavioral industrial organization, see Heidhues and Kőszegi (2018) for a survey. While recent policy papers suggest that behavioral effects are particularly important in digital settings (Crémer et al., 2019; Scott Morton et al., 2019; Fletcher et al., 2021), academic research on this topic is scarce. Chapter 1 and 2 contribute to filling this gap. In Chapter 3, we propose that certain challenges in traditional retail markets can be naturally understood through the lens of incorrect beliefs, which classic competition models abstract from.

Chapter 1

(Mis-)Matchmaker

1.1 Introduction

The emergence of digital matchmakers has revolutionized the way people meet and interact. By reducing search frictions, these platforms have the potential to more efficiently match users. With the help of algorithms based on detailed user data, they promise to facilitate the search for suitable partners in many areas of life. In fact, online dating has become the most common way to meet potential partners in recent years, and for more than a decade, job searches have been conducted predominantly through such online platforms (Rosenfeld et al., 2019; Kircher, 2022). This paper investigates the impact of a platform with detailed user data on the resulting speed and assortativity of matching in the society. It highlights a novel source of mismatching: profit-driven, purposeful mismatching of platforms.

To do so, I study the matching rule of a profit-maximizing platform on which users search for a suitable match. To capture the two most prominent business models, I assume that the platform commits to either an amount of advertising or a payment per period in which the user is active.[1] In either case, spending their time searching is costly for users. To attract and keep users' attention, the platform offers users a recommended match in each period. First, I show that the most prominent search protocols used to study centralized or decentralized matching markets — the positive assortative matching rule (PAM) and a random matching rule — are strictly suboptimal. Instead, the platform uses its knowledge about users to strategically lower the quality of recommended matches. This induces agents to search longer and thereby increases the payments the platform can collect. Besides prolonging search, the resulting matching outcomes can be drastically different from the socially optimal outcome — positive assortative matching — and induce a substantial welfare loss.

Why do platforms then rely on business models that induce misaligned incentives? I provide two plausible explanations. First, when, as in many online markets, users are reluctant to make monetary payments but are willing to consume ads,[2] offering an ad-based model can be more profitable. Second, when users have arguably well-documented misperceptions such as being overconfident regarding their desirability,[3] they underesti-

[1]See Appendix C for evidence on the business model of dating and job search apps.

[2]Advertising-based models play a key role in online markets, including both fully ad-supported and "freemium" business models. Freemium refers to business models, where users can use a basic service for free in exchange for consuming ads, but need to pay a fee to use the premium service (without ads). Freemium has become the most popular pricing strategy for many apps (see ACM (2019) or `https://www.statista.com/chart/1733/app-monetization-strategies/`).

[3]Overconfidence has been widely documented in the experimental literature, see for example Burks et al. (2013) and Dubra (2015). Especially overconfidence with respect to one's own attractiveness is common (Greitemeyer, 2020). Psychologists argue that such overconfidence determines how individuals look and compete for potential partners (Murphy et al., 2015). In labor markets, Spinnewijn (2015) and Mueller et al. (2021) find that the unemployed overestimate how quickly they will find a job. Moreover, beliefs are not revised (sufficiently) downward after remaining unemployed. Both findings suggest that job seekers are persistently overconfident about their desirability to firms.

mate their expected search duration and hence payments to the platforms for existing pay-per-month schemes.

After discussing the related literature in Section 1.2, Section 1.3 presents the model. A monopoly platform organizes a two-sided matching market in which users search for a partner on the opposite side. The platform commits to a matching rule that determines the probability that two users — each characterized by a vertical type — will meet. Additionally, the platform commits to a per-period cost that it collects from active users, which are either an amount of advertising or a search fee per period. After active users have paid the per-period cost, they receive a recommendation from the platform. Upon meeting, users simultaneously decide whether to accept or reject the proposed match. After rejecting, a user can continue to search. The analysis focuses on steady states; in these the inflow of new agents must equal the outflow under the platform's matching rule.

Section 1.4 starts by characterizing the users' search behavior. Then, fixing search costs, the platform's problem is to choose matching probabilities conditional on each users' type subject to participation constraints regarding the users' decision to join the platform, incentive constraints on the users acceptance decisions, feasibility constraints on the matching mechanism as well as steady-state constraints. This original problem is highly non-linear. Instead of analyzing the original problem, I make use of an auxiliary problem. This auxiliary problem is a linear programming problem in which the platform chooses masses of recommended matches and matched pairs accepting each other using the facts that: (i) the objective function is linear in steady-state masses, and (ii) the constraints are linear in the mass of recommended and matched pairs by using appropriate transformations. The profit-maximizing solution to this auxiliary problem is then transformed back to the solution of the original problem. Given the profit-maximizing matching rule, the platform chooses its advertising level or search fee. In the most general setting for any given finite set of users' types, I prove that an optimal solution to the platform's profit-maximization problem exists using the auxiliary problem. Based on the reformulation, I show that the widely analyzed matching rules are suboptimal. Random matching is suboptimal, when at least two types on each side of the market participate. Moreover, whenever both market sides are fully symmetric I show that the positive assortative matching rule — where each user meets a user of their own type — can be suboptimal.

Considering the special case with two types on each side of the market and symmetric inflows, Section 1.4.2 illustrates the main insight of the model — the platform's incentive to recommend and foster mismatches. To induce users to search, the platform frequently recommends mismatches to users, i.e., a high type meets a low relatively more often than a high type. The socially efficient matching outcome in which users sort positively is only implemented by the platform if significantly more low than high types enter the market. Otherwise, the platform induces a weakly, or even non-assortative, matching outcome.

The platform's matching thus creates two intertwined inefficiencies: it distorts matching outcomes by inducing mismatches that deviate from the socially optimal outcome, and it increases users' search time, leading to higher search costs than necessary. Both inefficiencies have implications for real-world markets such as dating and labor markets. In particular, in labor markets, the extent of mismatch has a significant impact on productivity and long-term unemployment (Şahin et al., 2014; McGowan and Andrews, 2015). Moreover, prolonged search duration, i.e., time spent unemployed or in a mismatched job, has high economic and social costs (e.g., unemployment insurance). In marriage markets, sorting has been found to have important implications for income inequality and household decisions (Lee, 2016). In addition, the quality of the relationship or marriage is a determinant of overall well-being and health (Robles et al., 2014; Sharabi and Dorrance-Hall, 2024). In the special case with two types, I find that the socially efficient matching outcome can induce the longest search time of agents, while the search time of agents decreases when the platform implements a weakly assortative or non-assortative outcome.

Finally, Section 1.5 turns to the question of why platforms rely on business models in which the incentives between the platform and the users are misaligned. For example, a simple potential business model for platforms would be to collect high personalized search fees from each type and provide them with the socially optimal match in the first period. In principle, this business model extracts the entire surplus from users. Under the realistic assumption that users are reluctant to pay upfront but are willing to consume ads, however, I show that an ad-based model can outperform the former business model if targeted advertising is sufficiently efficient. Alternatively, if users are overconfident about their desirability, this belief leads users to underestimate their search time when incentivized to search. Therefore, under the pay-as-you-search business model they spend a higher amount ex post than anticipated ex ante. This, in turn, favors the prevailing business model.

Section 1.6 concludes and highlights that the tension arising from the misalignment of incentives becomes more important as the platform collects more data and develops more predictive algorithms.

1.2 Related Literature

This article contributes to two central strands of literature, which I detail below. In contrast to the literature, I consider the profit-maximizing incentives of a matchmaker when agents are vertically differentiated and characterize the matching rule and resulting matching outcome.

Matching and Search Theory The vast literature on search-and-matching models, see for instance Burdett and Coles (1999), Eeckhout (1999), Bloch and Ryder (2000), and Smith (2006), provides insights into the functioning of decentralized markets in which agents meet at "random".[4] These matching models with heterogeneous agents build the foundation to investigate sorting and mismatch in markets such as labor and marriage markets when search frictions are present. In line with these models, agents in my model have vertical preferences that result in a unique stable matching. I follow Lauermann and Nöldeke (2014) and suppose that types are finite. The model at hand crucially departs from the literature on decentralized matching, which assumes that agents meet according to a random matching technology, by explicitly accounting for the design of the matching rule. With increasing access to user data about preferences and machine-learning tools, matching platforms can design their own recommendation and matching algorithms to maximize profits. While many platforms do not disclose the specifics of their matching algorithms, it is evident that their algorithms are far more sophisticated than random matching.[5]

The question of how to design the matching rule is related to the literature on centralized matching as pioneered by Gale and Shapley (1962), Roth (1982), and Roth and Sotomayor (1992), which studies match quality and implementation of efficient matching rules in two-sided markets.[6] The principal considers properties such as stability, strategy-proofness and Pareto efficiency of the matching rule. In contrast, I characterize the profit-maximizing solution for different given business models.

Search problems are widely studied not only on an individual level but researchers also rely on these to better understand job search and its implications on the functioning of the economy. Early articles include Pissarides (1985), Mortensen and Pissarides (1994), and Mortensen and Pissarides (1999), which focus on wage bargaining and unemployment dynamics and on-the-job search when agents are ex-ante homogeneous. Dolado et al. (2009) introduces heterogeneous types of workers and firms into job search models, which are also crucial in my model. A recent treatment on how job search has changed in the digital era is provided by Kircher (2022).

Finally, my paper is related to papers investigating biased beliefs of agents in matching and search markets. Closely related in a dating market, Antler and Bachi (2022) show that agents' coarse reasoning leads to overoptimism about their prospects in the market

[4]The aforementioned literature assumes that agents have non-transferable utility. Search-and-matching models with transferable utility have been analyzed, for example, by Becker (1973, 1974) and Shimer and Smith (2000). For an overview of the literature on search-and-matching models see Chade et al. (2017).

[5]Dating platforms such as Tinder or bumble provide a general description of their algorithm, see for example `https://www.help.tinder.com/hc/en-us/articles/7606685697037-Powering-Tinder-The-Method-Behind-Our-Matching`, whereas the dating platform "Hinge" claims to use the Gale-Shapley algorithm designed to find stable matchings.

[6]The literature on matching in two-sided markets can be divided into centralized and decentralized matching (see Echenique et al. (2023) for a recent overview).

and induces them to search inefficiently long. In labor markets, Spinnewijn (2015) and Mueller et al. (2021) document that job seekers often hold overoptimistic beliefs and thereby underestimate their time to find a job. I contribute to this literature by showing how current platform business models exploit overconfident types.

Platform Markets Central to the literature that studies platform and (online) two-sided markets is the presence of network effects and how these shape the incentives and price setting of a platform that enables the interaction between two groups (Caillaud and Jullien, 2003; Rochet and Tirole, 2003, 2006). As a result, in most models agents are assumed to care only about the number of matches instead of match quality.

With the emergence of digital matchmakers, the literature extended to analyzing (customized) matching on platforms with a focus on the interaction between pricing and matching efficiency (Damiano and Li, 2007; Damiano and Hao, 2008), price discrimination (Gomes and Pavan, 2016, 2024), and auctions (Johnson, 2013; Fershtman and Pavan, 2022), all abstracting from search frictions and dynamics. In my model, the platform designs the matching rule in its online market place, but in contrast to the aforementioned articles, the platform has an incentive to not implement the efficient and full surplus extracting matching rule.

Within the analysis of digital matchmakers, Halaburda et al. (2018a) and Antler et al. (2023, 2024) also focus on applications to dating platforms. Most closely related is Antler et al. (2024) who study a matchmaker's incentives in a model with horizontally differentiated types, which determine the fit of agents. The platform charges a single "upfront" fee in the second period after agents have joined and received their first match for free. The authors draw a similar conclusion: the platform has an incentive to invest into a technology that increases the speed of search but not into improving match quality. The main difference lies in modeling the matching technology. The authors restrict attention to a truncated random matching technology under which agents meet at random above a threshold and do not meet if their fit is below the threshold; in contrast, I solve for the optimal matching rule.

Within the platform literature models on platforms intermediating consumer search Hagiu and Jullien (2011, 2014), Eliaz and Spiegler (2011b, 2016), and Nocke and Rey (2024) are closely related. Hagiu and Jullien (2011) provide a rationale for intermediaries to divert search of their consumers away from preferred stores. Although the insight is closely related to the mismatching incentive in my model, the (one-sided) market in Hagiu and Jullien (2011) does not include the strategic component on the other side as stores would never reject a consumer willing to buy. Hence, there is no analogue to my finding that the platform prolongs search of lower types by recommending them to higher types knowing that they will reject those lower types. Additionally, there is no equivalent to overconfident users in their model. Finally, my model of a two-sided matching market

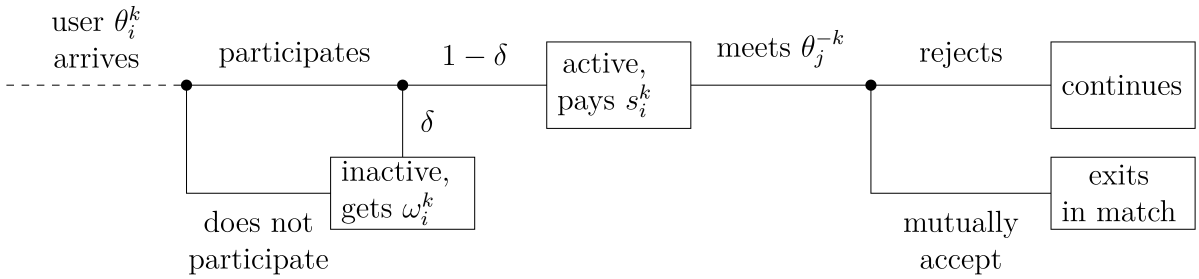

Figure 1.1: Within-Period Timing

offers insights into the allocative inefficiency and the length of search for labor and dating markets intermediated by matching platforms.

1.3 Model

A monopolist platform organizes a matching market in which a continuum of agents from two sides, $k = A, B$, search for a partner from the opposite side. The market operates in discrete time with an infinite horizon. I focus on steady state analysis. In slight abuse of notation, I therefore suppress time indices whenever it does not lead to confusion.

Agents An agent of each side is characterized by a type $\theta_i^k \in \Theta^k$, with $\Theta^k = \{\theta_1^k, \theta_2^k, ..., \theta_{N^k}^k\}$ finite. At the beginning of each period, an agent θ_i^k decides whether to enter the market or to exit and take outside option ω_i^k. An agent that participates in the market becomes inactive with an exogenous probability $\delta > 0$ and also leaves the search process. The platform charges an active agent of type θ_i^k a search cost s_i^k. Then, each active agent receives a single recommendation from the platform. After receiving a recommendation, two agents who meet observe each other's type and simultaneously decide whether to accept or reject the other agent. The following payoffs are realized based on their actions in the current period: (i) mutual acceptance yields a match utility of $u(\theta_i^k, \theta_j^{-k}) = \theta_i^k \theta_j^{-k}$, and (ii) (one-sided) rejection yields a utility of zero in the current period. After a rejection, an agent can continue searching in the next period. The timing within each period is summarized in Figure 1.1.

Agents are assumed to use time- and history-independent strategies. A pair of functions $\sigma_k : \Theta^k \times \Theta^{-k} \to [0,1]$ and $\sigma_{-k} : \Theta^k \times \Theta^{-k} \to [0,1]$ describe the acceptance strategies, where $0 \leq \sigma_k(\theta_i^k, \theta_j^{-k}) \leq 1$ is the probability that an agent of type θ_i^k on side k accepts a match with type θ_j^{-k} on the other side. The function $\eta_i^k : (\theta_i^k, \omega_i^k) \to [0,1]$ describes the participation strategy of an agent of type θ_i^k with outside option ω_i^k. In other words, without loss of generality, I focus on strategies where identical agents, active on the same side of the market and of the same type, use the same acceptance and participation strategy.

Then,

$$\alpha(\theta_i^k, \theta_j^{-k}) = \sigma_k(\theta_i^k, \theta_j^{-k}) \cdot \sigma_{-k}(\theta_i^k, \theta_j^{-k})$$

denotes the probability of a mutual acceptance by type θ_i^k and θ_j^{-k}.

Matching A matching mechanism $\mathcal{M} := \{\phi^k(\cdot)\}_{k=A,B}$ consists of (potentially stochastic) matching rules $\phi^k(\cdot)$. Let $\hat{\Theta}^k$ be the set of participating types from side $k = A, B$. For $\theta_i^k \in \hat{\Theta}^k$, $\phi^k(\cdot|\theta_i^k) \in \Delta(\hat{\Theta}^{-k} \cup \omega_i^k)$, which is a probability measure over $\hat{\Theta}^{-k} \cup \omega_i^k$. Intuitively, this describes the probability of meeting the various types of the opposing side as well as the outside option. Any $\theta_i^k \in \Theta^k \setminus \hat{\Theta}^k$ who does not participate is assumed to be meet their outside option with probability one, $\phi(\omega_i^k|\theta_i^k) = 1$. Denote the mass of agents of type θ_i^k on side k by $f(\theta_i^k)$. Matching mechanism $\mathcal{M}$ induces a distribution of matched pairs M

$$\left(\begin{pmatrix} f(\theta_1^k) \\ \vdots \\ f(\theta_{N^k}^k) \end{pmatrix}, \begin{pmatrix} f(\theta_1^{-k}) \\ \vdots \\ f(\theta_{N^{-k}}^{-k}) \end{pmatrix}\right) \mapsto \begin{pmatrix} \Phi(\theta_1^k, \theta_1^{-k}) & \cdots & \Phi(\theta_1^k, \theta_{N^{-k}}^{-k}) \\ \vdots & & \vdots \\ \Phi(\theta_{N^k}^k, \theta_1^{-k}) & \cdots & \Phi(\theta_{N^k}^k, \theta_{N^{-k}}^{-k}) \end{pmatrix} \equiv M.$$

An entry of matrix M is the mass of agents that are recommended to each other under matching mechanism $\mathcal{M}$ and is given by

$$\Phi(\theta_i^k, \theta_j^{-k}) = f(\theta_i^k)\phi(\theta_j^{-k}|\theta_i^k) = f(\theta_j^{-k})\phi(\theta_i^k|\theta_j^{-k}),$$

where the masses are symmetric, i.e. the mass of agents of type θ_i^k on side k being matched to agents of type θ_j^{-k} on side $-k$ is equal to the mass of agents of type θ_j^{-k} on side $-k$ being matched to type θ_i^k on side k: $\Phi(\theta_i^k, \theta_j^{-k}) = \Phi(\theta_j^{-k}, \theta_i^k)$. Under matching mechanism $\mathcal{M}$, the mass of agents of type θ_i^k that are unmatched, i.e. do not receive a recommendation in a given period, is

$$\Phi(\theta_i^k, \omega_i^k) = f(\theta_i^k) - \sum_{\theta_j^{-k} \in \Theta^{-k}} \Phi(\theta_i^k, \theta_j^{-k}).$$

To capture the idea that the platform can only generate revenue by keeping users' attention and, hence, wants to match as many agents as possible, I impose the following assumption.

Assumption 1. *Let $\hat{k}$ be the short side of the market. For each agent on side $\hat{k}$, $\phi(\omega_i^k|\theta_i^{\hat{k}}) = 0$.*

Under Assumption 1, feasibility of the matching rule can be expressed in terms of the masses of matched pairs.

Definition 1. A matching mechanism $\mathcal{M}$ is feasible if

$$\sum_{\theta_j^{-k} \in \Theta^{-k}} \Phi(\theta_i^k, \theta_j^{-k}) + \mathbf{1}_{k=\hat{k}} \Phi(\theta_i^k, \omega_i^k) = \eta_i^k f(\theta_i^k), \forall \theta_i^k \in \Theta^k, k = A, B. \tag{1.1}$$

Timing and Population Dynamics At the beginning of a period t, agents who did not find a match in the last period arrive and a (time-invariant) inflow of new agents of type θ_i^k given by the mass $\{\beta_i^k\}_i^{k=A,B}$ enters the platform. Agents decide whether to participate on the platforms. Those who decide to participate become inactive with probability δ, while active agents are matched according to matching mechanism $\mathcal{M}$ resulting in matrix M_t. Based on their recommended match, agents make their acceptance decision resulting in mutual acceptance probabilities $\{\alpha_t(\theta_i^k, \theta_j^{-k})\}_{ij}$. At the end of the period, agents that mutually accepted each other exit in pairs. The total outflow of agents is then given by pairs that exit together in a match, agents that become inactive with probability δ and agents that decided not to participate.

Platform The platform commits to a matching mechanism $\mathcal{M} := \{\phi^k(\cdot)\}_k$. To capture the two most prominent business models, I suppose that the platform either commits to an extent of advertising or a given payment per period. Formally, this choice induces the type-dependent search cost s_i^k while generating revenue per search of type θ_i^k of $\nu(s_i^k)$. In case of payments, $\nu(s_i^k)$ is the identity function. In case of advertisements, $\nu(s_i^k)$ is an increasing and strictly concave function of the search costs, which for example captures the intuition that the agents' disutility of advertising is convex in the number of ads shown while the platform's profit is constant per ad. Let $s_i^k \in [0, \overline{u}]$, where $\overline{u}$ is the maximum match utility that the highest type can achieve on the platform. The platform discounts future profits according to ρ and thus maximizes

$$\Pi = \sum_{k=A,B} \sum_{\theta_i^k \in \Theta^k} \frac{(1-\delta)\eta_i^k}{1-\rho} \nu(s_i^k) f(\theta_i^k).$$

Equilibrium Concept The model focuses on a steady state analysis in which the market is balanced: that is, the inflow of agents is equal to the outflow of agents under matching mechanism $\mathcal{M}$. Formally:

Definition 2. (Steady State) For given matching mechanism $\mathcal{M}$, a *steady state* is a tuple $(f(\theta_i^k), \alpha(\theta_i^k, \theta_j^{-k}), \eta_i^k)_{ij}^k$ that satisfies

$$\beta_i^k = f(\theta_i^k) \left[(1 - \eta_i^k) + \eta_i^k \left(\delta + (1-\delta) \sum_{\theta_j^{-k} \in \Theta^{-k}} \alpha(\theta_i^k, \theta_j^{-k}) \phi(\theta_j^{-k} | \theta_i^k) \right) \right], \tag{1.2}$$

for all $\theta_i^k \in \Theta^k, k = A, B$. The left-hand side describes the inflow of agents of type θ_i^k, where the right-hand side is the outflow. The outflow is the mass of type θ_i^k agents times the probability that agents do not participate plus the probability of becoming inactive or exiting in a match.

A steady state is an equilibrium if the following is satisfied.

Definition 3. (Equilibrium) A steady state is an equilibrium if — given that agents anticipate other agents' strategies correctly — the profile of stationary strategies (σ, η) satisfies:

1. Agents accept a match if and only if the match yields a higher payoff than the expected utility from continuing to search.

2. Agents participate if and only if the expected utility from participating yields a higher payoff than their outside option.

Under the usual Nash assumption of correctly anticipating other players' strategies, the definition captures that agents maximize expected utility with respect to their acceptance strategy implicitly ruling out the case that a valuable pair is rejected because everyone is certain that their partner rejects.[7] The third part captures that agents maximize expected utility when deciding to participate on the platform.

1.3.1 Discussion of Assumptions

Search Costs Agents incur additive search costs s_i^k in each period, which are designed by the platform. They either represent the nuisance costs from advertising as, for example, in Anderson and Coate (2005), which are positively related to the advertising intensity, or the search fee that the platform charges periodically. Search frictions are modeled by introducing the exogenous exit probability δ. Following a literal interpretation, δ is the probability with which agents become inactive, i.e. the probability that an agent finds a job or a partner offline through other means. More generally, δ can be thought of as modeling the force that leads agents to discount the future, which makes delayed matching more costly.

Business Model The platform is assumed to be a monopolist in the matching market. Following evidence from the dating market, the most popular dating platforms have a common owner. For simplicity, I assume that the dominant owner only offers one platform

[7]This allows the current match partner to tremble with small probability. Alternatively, acceptance decisions could be made sequential in which case agents would have to accept a valuable match.

in my model.[8] More generally, we often observe platforms with large market power in two-sided markets, where joining a new platform is worthwhile only if others join. My monopoly setup is a simple setting capturing such market power.

The model examines two prevalent business models: an advertisement-based approach and periodic search fees. Many platforms adopt the former— (targeted) advertising — by monetizing user attention through selling advertising slots to firms. In return for users' attention, the platform provides its matching service for free. In this setup, keeping user attention is crucial for the platform's revenue.[9] This is why I assume that the platforms earns no revenue when not capturing the user's attention through offering a potential match. Alternatively, platforms implement search fees, which they collect from active users. Examples include "pay-per-click" or "pay-per-contact" fees, though monthly subscription plans are also common. These fees are typically low, distinguishing them significantly from participation fees, which are far less common but used by some selective matching platforms.[10]

An advertising-based stream of revenues continues to be a prominent part of platform business models, especially with transaction costs. Platforms have transaction costs when setting up a payment system, while many users are reluctant to give their credit card data to platforms. Overall, privacy concerns, risk aversion and uncertainty about new products (platforms) can play a role why users (initially) prefer to use the matching service for "free" while watching advertisement over signing up to a subscription plan or paying a participation fee. As a consequence, many platforms rely on these so-called "freemium" business models, which have become even more popular since the emergence of mobile applications (apps). Here, "freemium" describes business models where a basic service is available to users for free (with advertisement), whereas an upgraded service can be accessed through purchases.[11] Other platforms, however, rely only on advertising or fees. I return to the question of why platforms refrain from collecting a fixed fee for a certain promised match in Section 1.5.

[8]The dating market is highly concentrated with the Match Group Inc. owning many of the most popular dating platforms: Tinder, Hinge, PlentyofFish, Match, OkCupid etc. (see `https://www.bamsec.com/filing/89110323000114?cik=891103`), while other dating platforms are highly differentiated and for example, cater to specific religious groups. Recent experimental evidence from Dertwinkel-Kalt et al. (2024) suggest that even the closest competitors, Tinder and bumble, are viewed to be almost independent instead of substitutes by consumers.

[9]Recent papers that study different aspects of attention on platforms are for example Prat and Valletti (2022), Chen (2022), and Srinivasan (2023).

[10]For an overview of the most common platforms and their fee structure see Appendix C.

[11]For empirical evidence see for example, Kummer and Schulte (2019) for studying privacy concerns in the mobile app market and Deng et al. (2023) for studying freemium pricing of mobile applications.

No Agent is Unmatched The key assumption of the matching rule, Assumption 1, states that if possible each agent receives a recommended match in any period.[12] As many online platforms take on a dual role as attention intermediaries and need to attract consumers' attention to sell to advertisers, providing a constant stream of potential matches aims at grabbing and keeping consumers' attention.[13]

To grab users' attention, the platform makes a recommendation any time the agent enters and is active on the platform. The recommendation of a potential match can be viewed as being part of a menu that the platform offers. Following the idea of Eliaz and Spiegler (2011a), the platform offers a menu that consists of an attention-grabbing component and its true value of the service. In reality, the attention grabbing component is supported by push notifications or emails, while the value from the platform's service is determined by the expected utility from getting a match. The modeling choice is further supported by a recent lawsuit against the MatchGroup Inc., owner of a majority of the most popular dating platforms.[14] In the complaint, the plaintiff accuses Match to monopolize users' attention and claim that "Push Notifications prey on users' fear of missing out on any potential matches with a strategic notification system designed to capture and retain attention throughout the day".

1.4 Analysis

To analyze the equilibrium, I need to characterize the agents' behavior and the platform's optimization problem. The agents' search process is characterized by a set of participation and incentive constraints that determine whether an agent is willing to incur the search costs as well as accepts or rejects a recommended match.

Agents' Search Process. Consider the strategy of agent θ_i^k being active in the matching market. Upon meeting an agent θ_j^{-k}, the agent decides whether to accept or reject the *recommended* match. Mutual acceptance results in a *match* and both agents leave the market as a pair. If at least one of the agents rejects the match, agent θ_i^k continues to search.

[12] In the literature on search-and-matching models time is often continuous, such that matching opportunities arrive at a constant rate. Similarly, Antler et al. (2023, 2024) make the assumption that matches arrive at a constant rate even in the presence of a matchmaking platform.

[13] In a recent experiment, Aridor (Forthcoming) provides evidence that users allocate their attention across product categories and offline when facing restriction in their time spent on a specific platform. The results suggest that competition for attention spans across multiple markets.

[14] Oksayan v. MatchGroup Inc., N.D. Cal., No. 3:24-cv-00888, 2/14/24.

Due to the stationarity of the environment, the continuation value of agent θ_i^k, $V^C(\theta_i^k)$, is defined by the following recursive equation

$$\begin{aligned} V^C(\theta_i^k) =& \delta\omega_i^k + (1-\delta)\Bigg[-s_i^k + \sum_j \alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k}|\theta_i^k)\theta_i^k\theta_j^{-k} \\ &+ (1 - \sum_j \alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k}|\theta_i^k))V^C(\theta_i^k)\Bigg]. \end{aligned}$$

The first term represents the case in which agent θ_i^k will become inactive with probability δ and gets its outside option ω_i^k. If the agent remains active with probability $1-\delta$, it incurs the search cost s_i^k. The expected utility from leaving in a match is given by the utility from a match with type θ_j^{-k}, which is equal to the product of both types, and the probability of meeting and mutually accepting type θ_j^{-k}. With the counterprobability, the match was not mutually accepted and agent θ_i^k continues to search.

Solving for the continuation value yields

$$V^C(\theta_i^k) = \frac{\delta\omega_i^k + (1-\delta)\left(-s_i^k + \sum_j \alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k}|\theta_i^k)\theta_i^k\theta_j^{-k}\right)}{\delta + (1-\delta)\left(\sum_j \alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k}|\theta_i^k)\right)}. \tag{1.3}$$

The continuation value then characterizes the payoff of an agent who rejects a match and returns to the search process, whereas the match payoff $\theta_i^k\theta_j^{-k}$ characterizes the payoff of an agent who accepts a match with type θ_j^k (and is accepted by them). By Definition 3, if the match value $\theta_i^k\theta_j^{-k}$ is smaller (larger) than the continuation value $V^C(\theta_i^k)$, agent-θ_i^k will reject (accept) a recommended match with agent-θ_j^{-k}.

The optimal strategy of an agent who uses a time-and history-independent strategy satisfies:

$$\sigma_k(\theta_i^k, \theta_j^{-k}) = \begin{cases} 0 & \text{if } \theta_i^k\theta_j^{-k} < V^C(\theta_i^k) \\ r \in [0,1] & \text{if } \theta_i^k\theta_j^{-k} = V^C(\theta_i^k) \\ 1 & \text{if } \theta_i^k\theta_j^{-k} > V^C(\theta_i^k) \end{cases}, \text{ for } k = A, B. \tag{1.4}$$

If the match value with a type $\hat{\theta}_j^{-k}$ is larger than the continuation value, agent θ_i^k will accept a recommended match with agent $\hat{\theta}_j^{-k}$ and all agents of types higher than $\hat{\theta}_j^{-k}$. The optimality of this strategy follows directly from the supermodularity of the match payoff.

An agent participates if the continuation value is larger than the agent's outside option. Due to the stationarity and history-independence of strategies, if an agent decides to participate in the matching market, they will not exit during the search process and search until they exit in a match or become inactive with probability δ.

Remark. The strategy of an agent of type θ_i^k is increasing in its second argument $\sigma_k(\theta_i^k, \theta_{N^{-k}}^{-k}) \geq \sigma_k(\theta_i^k, \theta_{N-1}^{-k}) \geq \cdots \geq \sigma_k(\theta_i^k, \theta_1^{-k})$, but may be neither in- nor decreasing in its first argument.

The fact that the agent's strategy is increasing in its second argument follows directly from Equation 1.4. If the agent's outside options are weakly increasing in type, for matching rules such as random or positive assortative matching rules $\sigma_k(\theta_i^k, \theta_j^{-k})$ is additionally decreasing in its first argument: $\sigma_k(\theta_{N^k}^k, \theta_j^{-k}) \leq \cdots \leq \sigma_k(\theta_1^k, \theta_j^{-k})$. A random matching rule yields the same meeting probabilities for all types. Due to the supermodularity of the payoff function, higher types will reject (weakly) higher types than lower types do. With positive assortative matching, the matching probabilities conditional on being a higher type first-order stochastically dominates the matching probabilities conditional on being a lower type. Hence, higher types will reject strictly higher types than lower types do. In contrast, a negative assortative matching rule, which recommends (almost exclusively) higher types to lower types, and vice versa, can cause lower types to reject lower types while higher types are willing to accept them. Indeed, I will explicitly provide an example of such an equilibrium in Section 1.4.2.

Given the agent's strategy in Equation 1.4, the acceptance probabilities satisfy

$$\alpha(\theta_i^k, \theta_j^{-k}) = \begin{cases} 0 & \text{if } \theta_i\theta_j < V^C(\theta_i^k) \text{ or } \theta_i\theta_j < V^C(\theta_j^{-k}) \\ 1 & \text{if } \theta_i\theta_j > V^C(\theta_i^k) \text{ and } \theta_i\theta_j > V^C(\theta_j^{-k}) \end{cases}. \tag{1.5}$$

Equation 1.5 establishes the relationship between acceptance probabilities and matching outcomes. Mutual acceptance requires that whenever two types of agents meet, both must find it optimal to stop searching.

1.4.1 Multiple Types

Consider the case with N^k types of agents such that $\Theta^k = \{\theta_1^k, ..., \theta_{N^k}^k\}$ on side $k = A, B$, where $\theta_{N^k}^k > ... > \theta_1^k$. The following section provides general results on the existence of an equilibrium, optimal solution and their properties. Let s_i^k be exogenous.

Lemma 1. *For a given feasible matching mechanism, a steady-state equilibrium exists if and only if Equation 1.2 and 1.5 are satisfied.*

Suppose for a feasible matching mechanism, an equilibrium exists. Then, it must give rise to (i) a steady state and (ii) optimal strategies of agents, i.e. satisfy Definition 2 and Definition 3. Hence, by (i) Equation 1.2 (balance condition) must hold, and (ii) implies Equation 1.5 (optimal mutual acceptance) must hold. Conversely, if Equation 1.2 is violated the steady state (balance) condition fails and if Equation 1.5 is violated at least some agent behaves suboptimal. Thus, a feasible matching rule gives rise to an equilibrium if and only if Equation 1.2 and 1.5 hold.

Lemma 2. *There exists a feasible matching rule that gives rise to an equilibrium.*

In the most simple case consider the matching rule $\phi(\omega_i^k|\theta_i^k) = 1$ for all types $\theta_i^k \in \Theta^k$ on side $k = A, B$. Given that agents are matched with their outside option, no agent is willing to incur search costs. With no agent participating in the steady state, the matching rule is feasible and gives rise to a steady state equilibrium.

Next, to determine the profit-maximizing matching rule $\mathcal{M}$, it is useful to define the matching outcome. Intuitively, the matching outcome is defined as the matrix that describes the distribution of pairs under matching rule $\mathcal{M}$ that exit in a match. Recall that matrix M describes the masses of recommended pairs under matching rule $\mathcal{M}$ and let A denote the matrix of agents' mutual acceptance probabilities

$$A \equiv \begin{pmatrix} \alpha(\theta_1^k, \theta_1^{-k}) & \cdots & \alpha(\theta_1^k, \theta_{N^{-k}}^{-k}) \\ \vdots & & \vdots \\ \alpha(\theta_{N^k}^k, \theta_1^{-k}) & \cdots & \alpha(\theta_{N^k}^k, \theta_{N^{-k}}^{-k}) \end{pmatrix}.$$

Formally, the matching outcome is defined as the componentwise multiplication (Hadamard product) of matrix A and M:

Definition 4. The matching outcome is defined by the matrix

$$A \odot M = \begin{bmatrix} \alpha(\theta_1^k, \theta_1^{-k})\Phi(\theta_1^k, \theta_1^{-k}) & \cdots & \alpha(\theta_1^k, \theta_{N^{-k}}^{-k})\Phi(\theta_1^k, \theta_{N^{-k}}^{-k}) \\ & \vdots & \\ \alpha(\theta_{N^k}^k, \theta_1^k)\Phi(\theta_{N^k}^k, \theta_1^{-k}) & \cdots & \alpha(\theta_{N^k}^k, \theta_{N^{-k}}^{-k})\Phi(\theta_{N^k}^k, \theta_{N^{-k}}^{-k}) \end{bmatrix} \equiv O(\cdot).$$

Matching outcomes are (i) assortative if $O(\cdot)$ has positive entries only along the main diagonal, (ii) weakly assortative if $O(\cdot)$ has positive entries along the main diagonal and to the right if and only if all entries below are also positive, and (iii) non-assortative otherwise.

If a matching outcome is assortative, this implies that lower types are matched with strictly lower types than higher types while the definition of weakly assortative implies that lower types can be matched with the same types as higher types. The definition is weak in the sense that it does not require that lower types accept with a higher probability than higher types. Other matching outcomes are called non-assortative and entail negative assortative outcomes where higher types are matched with strictly lower types than lower types.

Denote by $m(\theta_i^k, \theta_j^{-k}) = \alpha(\theta_i^k, \theta_j^{-k})\Phi(\theta_i^k, \theta_j^{-k})$ an entry of matrix $O(\mathcal{M})$. Each entry is therefore the mass of *matched* pairs that exit the market together in every period. For a given matching rule, an equilibrium induces at most one matching outcome since the mutual acceptance probabilities and steady state masses are pinned down in equilibrium.

To find the profit-maximizing matching rule and the associated matching outcome, I proceed in two steps. First, I fix a matrix of acceptance probabilities and determine the optimal feasible matching rule that implements the mutual acceptance probabilities. Second, supposing the optimal matching rule from step one is used to implement any chosen matrix of acceptance probabilities, I choose the matrix that yields the highest platform profits.

First, note that the platform finds it optimal to induce either full participation of a type or no participation.

Lemma 3. *It is without loss of generality to consider $\eta_i^k \in \{0, 1\}$.*

Suppose the platform charges type-dependent search fees, and type θ_i^k, who is indifferent between participating and not participating, participates with probability less than one. Then, the platform makes the same profit if type θ_i^k participates with probability one, the platform sometimes matches them to their outside option, and reduces their search fee such that they make the same payments in expectation. If the platform uses an advertising-based business model, the platform will strictly increase its profit by this procedure due to the concavity of advertising. Therefore, from now on I will focus on $\eta_i^k \in \{0, 1\}$, which allows to focus on the set of participating types. Then, suppose the platform induces a set $\hat{\Theta}^k$ for $k = A, B$ to participate.

In the following, I will transform the platform's profit-maximization problem into a linear program. For given search cost s_i^k, recall that the platform's objective is to maximize

$$\max_{\mathcal{M}} \sum_{k=A,B} \sum_{\theta_i^k \in \hat{\Theta}^k} \frac{(1-\delta)s_i^k}{1-\rho} f(\theta_i^k),$$

i.e., the platform maximizes the steady state mass of active agents with weight s_i^k. Note that the platform does not earn revenue from agents that are inactive or do not participate in the market in the first place. The maximization problem underlies a set of constraints. First, the matching rule must implement a steady state. The steady state condition (Equation 1.2) implies

$$\beta_i^k = f(\theta_i^k)\delta + (1-\delta) \sum_{\theta_j^{-k} \in \Theta^{-k}} \underbrace{\alpha(\theta_i^k, \theta_j^{-k})\Phi(\theta_j^{-k}|\theta_i^k)}_{=m(\theta_i^k, \theta_j^{-k})}. \qquad \text{(Steady State)}$$

In the steady state, the inflow of agents of θ_i^k is equal to the mass of agents that become inactive in a period with probability δ and the mass of active agents that exit in matched pairs. In the steady state, the mass of agents of type θ_i^k can be restated as

$$f(\theta_i^k) = \frac{\beta_i^k - (1-\delta)\sum_j m(\theta_i^k, \theta_j^{-k})}{\delta}, \qquad \text{(Steady-State Mass)}$$

and therefore, depends positively on the inflow, β_i^k, and negatively on the mass of matched pairs that include type θ_i^k. Second, the matching rule determines whether agents participate in the market and whether agents search according to the platform's recommendations. For participating agents, it must hold that the agent prefers participating in the market to accepting the outside option, i.e.

$$\omega_i^k \leq \frac{\delta\omega_i^k + (1-\delta)\left(-s_i^k + \sum_j \alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k}|\theta_i^k)\theta_i^k\theta_j^{-k}\right)}{\delta + (1-\delta)\left(\sum_j \alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k}|\theta_i^k)\right)} = V^C(\theta_i^k).$$

Since the match payoffs are supermodular, there exists a critical lowest type that an agent θ_i^k is willing to accept (Equation 1.4). Agent θ_i^k rejects (accepts) all types below (above) the critical lowest type. The incentive constraint for agent θ_i^k to follow the recommendation of the platform to (weakly) reject an agent θ_j^{-k} reads[15]

$$\theta_i^k\theta_j^{-k} \leq V^C(\theta_i^k).$$

By using the steady state condition, the participation and incentive constraints can be reformulated. Note that the denominator of the continuation value is equal to the probability that an agent exists, which is equal to $\beta_i^k/f(\theta_i^k)$ by Equation 1.2. Inserting into the continuation value and rearranging yields

$$\beta_i^k\omega_i^k \leq \delta f(\theta_i^k)\omega_i^k - (1-\delta)f(\theta_i^k)s_i^k + (1-\delta)\sum_j \underbrace{\alpha(\theta_i^k, \theta_j^{-k})\Phi(\theta_j^{-k}|\theta_i^k)}_{=m(\theta_i^k, \theta_j^{-k})}\theta_i^k\theta_j^{-k}, \tag{PC}$$

$$\beta_i^k\theta_i^k\theta_j^{-k} \leq \delta f(\theta_i^k)\omega_i^k - (1-\delta)f(\theta_i^k)s_i^k + (1-\delta)\sum_j \underbrace{\alpha(\theta_i^k, \theta_j^{-k})\Phi(\theta_j^{-k}|\theta_i^k)}_{=m(\theta_i^k, \theta_j^{-k})}\theta_i^k\theta_j^{-k}. \tag{IC}$$

Lastly, the platform's matching rule must satisfy the feasibility constraints. Without loss of generality, let side B be of smaller or same size as side A. Then on side A, the sum over the mass of each recommended pair that includes type θ_i^A must be equal to the steady state mass of θ_i^A. On side B, the sum over the mass of each recommended pair that includes type θ_i^B and the mass of agents of type θ_i^B that are unmatched must be equal to the steady state mass of type θ_i^B

$$\sum_{\theta_j^{-k}\in\Theta^{-k}} \Phi(\theta_i^k, \theta_j^{-k}) + \mathbf{1}_{k=A}\Phi(\theta_i^k, \omega_i^k) = f(\theta_i^k), k = A, B. \tag{Feasibility}$$

As stated above, for given matrix A the above constraints and the objective function are all linear functions of the steady state masses, matched pairs, and recommended pairs.

[15] In mechanism design, this is often referred to as an obedience constraint because there is no private information throughout the model.

The steady-state mass in turn is also a linear functions of the mass of matched pairs. To complete the reformulation as linear program, it remains to include the indifference constraints for agents who mix when accepting type from the other market side, which implies that the respective incentive constraint must hold with equality. Appendix A.1 formally does so, leading to:

Lemma 4. *The platform's problem can be restated as a linear programming problem in the mass of matched and recommended pairs:* $\{m(\theta_i^k, \theta_j^{-k})\}, \{\Phi(\theta_i^k, \theta_j^{-k})\}_{ij}$.

Note that by Lemma 1, the solution to the linear program is an equilibrium as it fulfills Equation 1.2 and 1.5. Given a solution of the linear program — the auxiliary problem — the optimal matching rule to the original problem results from

$$\phi(\theta_j^{-k}|\theta_i^k) = \frac{\Phi(\theta_i^k, \theta_j^{-k})}{f(\theta_i^k)}.$$

Next, I show that the auxiliary problem has an optimal solution. I say that a matrix A of mutual acceptance probabilities can be implemented if there exists a matching mechanism $\mathcal{M}$ such that $\left((f(\theta_i^k))_{\theta_i^k \in \Theta^k}, A, \eta\right)$ is an equilibrium. Let $\mathcal{A}$ be the set of matrices A that can be implemented. By Proposition 2, $\mathcal{A}$ is non-empty. For every $A' \in \mathcal{A}$, construct a matrix A'' such that

$$\begin{aligned}
&\alpha''(\theta_i^k, \theta_j^{-k}) = \alpha'(\theta_i^k, \theta_j^{-k}) \text{ if } \alpha'(\theta_i^k, \theta_j^{-k}) \in \{0, 1\}, \\
&\alpha''(\theta_i^k, \theta_j^{-k}) = \alpha_{ij} \text{ otherwise,}
\end{aligned}$$

where α_{ij} can take on any value in $[0, 1]$. I use $\alpha_{ij} \in [0, 1]$ whenever an agent is indifferent, which implies that the same constraints in the auxiliary program must hold. Denote the resulting set of matrices as $\mathcal{A}^*$ and note that $\mathcal{A}^*$ is finite. Now, I can solve the linear program over the mass of matched and recommended pairs (ignoring acceptance probabilities). Solving this for all (finite) possible combinations of constraints yields a set of candidate solutions among which I choose the one that maximizes the platform's profit. To find the corresponding acceptance probabilities $\alpha_{ij} \in [0, 1]$ when the agent is indifferent, divide the matched pairs through the recommended ones

$$\alpha_{ij} = \frac{m(\theta_i^k, \theta_j^{-k})}{\Phi(\theta_i^k, \theta_j^{-k})}.$$

Formally, as $\mathcal{A}^*$ is finite, only a finite number of linear problems must be solved. Each linear program returns a set of candidate solutions and a value of the objective function. Fixing $A \in \mathcal{A}^*$, the linear program returns a value $\Pi(A)$, i.e., the profit level, and let $\mathcal{G} = \bigcup_{A \in \mathcal{A}^*} \Pi(A)$ be the set of profit levels for all linear programs with $A \in \mathcal{A}^*$ that implement an equilibrium.

Lemma 5. *The set $\mathcal{G}$ is non-empty and finite with $\Pi(A) < \infty$ for all $A \in \mathcal{A}^*$ and $-\infty < \Pi(A)$ for at least one $A \in \mathcal{A}^*$.*

Key to the proof is to show that the linear program for any given matrix $A \in \mathcal{A}$ is (i) not unbounded and (ii) not infeasible, i.e. the feasible region is non-empty. Given that both (i) and (ii) are satisfied, an optimal solution to the linear program exists and the linear program attains a finite optimal value (Dantzig, 1963).[16]

Theorem 1. *There exists an optimal solution.*

I proceed by showing that an optimal solution exists for any exogenous search costs s_i^k for all $\theta_i^k \in \Theta^k,\ k = A, B$. By Lemma 5, the maximum over set $\mathcal{G}$ is well-defined as $\mathcal{G}$ is finite and bounded such that an optimal solution exists. Next, I show that there exists an optimal solution if the platform chooses search costs s_i^k for all $\theta_i^k \in \Theta^k,\ k = A, B$. Through a series of Lemmas, I prove that the set $\mathcal{G}$ is compact-valued and upper hemicontinuous in the vector of search costs. This implies that the set $\max \mathcal{G}$ is upper semicontinous in the vector of search costs. Therefore, by an extension of the Weierstrass theorem a maximum exists.

To identify properties of the optimal solution, first consider two prominently studied matching rules. As discussed in Section 1.2, in decentralized matching-and-search markets agents are often assumed to meet according to a random matching technology. A natural question to consider is whether a platform that has access to extensive user data would commit to a random meeting technology as well.

Proposition 1. *Suppose $N^k \cdot N^{-k} > 1$. Then, random matching is generically suboptimal for exogenous search costs and endogenous search fees. Consider the class of functions: $\nu(s_i^k) = \kappa(s_i^k)^\alpha$ with $\kappa \in \mathbb{R}^+$ and $\alpha \in (0, 1)$. Random matching is generically suboptimal within this class of functions.*

The proposition shows that random matching is generically suboptimal for the platform if search costs are exogenous or type-dependent search fees are endogenously chosen.[17] For analytical convenience, I consider the class of concave revenue functions in the proof to determine a knife-edge solution.

Consider the nontrivial case in which there are different types to be matched. Under random matching, the conditional probability of meeting a type θ_i^k on side k is the same for all types $\theta_j^{-k} \in \Theta^{-k}$ on side $-k$ and corresponds to the proportion of type θ_i^k in the

[16]Existence follows from the fact that the constraint set is a convex polyhedron. Because the objective is linear and the constraint set is convex, any local extremum will be the global extremum. As the objective is linear, the extremum will be obtained at one of the extreme points of the constraint set, i.e., at the vertices of the polyhedron.

[17]Consider the following definition for *generically suboptimal*. The probability of the case in which random matching is optimal occurs with probability zero when the model parameters are randomly drawn from continuous intervals as defined in the proof.

population. As shown in Appendix A.2, the probability of meeting a type θ_i^k is a function of the inflow, β_i^k, and the probability of exit, δ. In contrast, for given search costs, the optimal solution of the linear program is a function of these and internalizes changes in the search cost. Therefore, random matching is generically suboptimal for given search costs, although it may coincide with the optimal solution for knife-edge $s_i^k, \theta_i^k, \delta$, and β_i^k. This result extends to the case in which the platform chooses a (linear) search fee. The platform does not choose random matching, but chooses a positive assortative matching rule that maximizes the agents match surplus and extracts all surplus via the search fee.

Proposition 1 highlights that a platform, which has increasing access to user data, does not commit to a random matching technology. Proposition 1 immediately implies that the platform values user data as access to data increases the platform's profit.

Corollary 1. *Suppose a platform has access to data about user types. The platform makes higher profits by using the data to discriminate users by conditioning the matching rule on user types instead of refraining from using user data.*

Second, consider the positive assortative matching rule (PAM) under the assumption that both sides are symmetric with respect to the inflow of new agents: $\beta_i^A = \beta_i^B$, their type space $\Theta^k = \Theta$, and outside options. Under symmetry, PAM matches agents if and only if they are of the same type on both sides of the market. In this particular case, PAM is of special interest in the literature as it maximizes total match surplus when the match utility is supermodular, where an agent's individual match surplus is defined as the difference between the expected match utility on the platform and the agent's outside option. Furthermore, the resulting matching outcome, i.e., the positive assortative matching outcome, is equivalent to the set of stable matchings (Roth and Sotomayor, 1992). That is, matches are individually rational, i.e., yield a utility greater than their outside option, and are pairwise stable, i.e., there exists no blocking pair of agent that would prefer to be matched to each other instead of the equilibrium matching. The next proposition shows under which circumstances the positive assortative matching rule (PAM) is not profit-maximizing under type-dependent search fees and advertising.

Proposition 2. *Suppose both market sides are symmetric.*

(i) PAM is profit-maximizing if the platform can charge arbitrary high type-dependent search fees. Conversely, for every type there exists a threshold $\overline{s}_i$ such that if $s_i < \overline{s}_i$, PAM is suboptimal.

(ii) There exists a threshold $\overline{\delta}$ such that if $\delta \leq \overline{\delta}$ and $\nu(\cdot)$ is concave, PAM is suboptimal.

When the platform commits to a (time-constant) deterministic matching rule such as PAM, agents will accept the recommended match in the first period. Therefore, all agents search for exactly one period, which results in a steady state population equal to the inflow for each type.

First, PAM is indeed profit-maximizing if the platform has pricing power. By charging (high) type-dependent fees, the platform can extract the full surplus from agents, i.e., the expected match value of an assortative match over the agent's outside option. In this case, the "search fee" is paid once, since agents search for only one period. The proposition, however, shows that if the platform cannot commit to high search fees, for example due to a (binding) price ceiling $\overline{s}$, then PAM is no longer optimal. Let $\overline{s}$ be such that s_i violates the condition in Proposition 2 for at least one type $\theta_i \in \Theta$. For the sake of exposition, assume that this is not the lowest type. Then the platform can no longer extract the full surplus from an agent of type θ_i. Then, PAM is not profit-maximizing, as the platform has an incentive to deviate to a matching rule under which type θ_i and the lowest type θ_1 meet with mass ε. The price ceiling $\overline{s}$ is such that whenever type θ_i and type θ_1, θ_i (weakly) rejects θ_1 under the new matching rule. This implies that type θ_i searches longer than one period such that the platform earns more from type θ_i. For example, fees for in-app purchases in Apple's App store are capped at 999.99\$, i.e., $\overline{s} = 999.99\$$. The estimated lifetime utility from a match and hence, potential willingness to pay for a partner could be well above 999.99\$. Traditional matchmakers charge over ten times the amount.[18] Alternatively, users may be reluctant to spend large sums online in one payment, such that the platform's pricing power might be limited as well.

Second, suppose the platform follows an advertising-based business model. If the return to advertising is concave and $\delta \leq \overline{\delta}$, then PAM is suboptimal. Under PAM agents search for only one period. Thus, a profit-maximizing platform would need to impose the highest feasible search cost per agent. With concave advertising returns, however, it becomes more profitable to reduce search costs and increase the mass of participating agents. Since $\delta > 0$ implies a loss in profits due to exogenous attrition that increases with longer search times, a high δ reduces the platform's willingness to trade off longer search durations for lower costs.

Proposition 2 raises the question of why we, as users, do not observe high search fees online, and why matching appears to be (anecdotally) worsening rather than improving. If the platform has pricing power and can perfectly identify users' types, Proposition 2 implies that it induces only one period of search and employs PAM to extract the full surplus from users. This raises the question: under what conditions does the platform have an incentive to induce more search?

In Section 1.4.2, I examine pricing under complexity constraints. I present an example with two user types, where the platform is limited to setting a single price, and show that

[18] See `https://www.nytimes.com/2024/02/13/business/dating-bounty-roy-zaslavskiy.html?unlocked_article_code=1.VU0.XqAb.q2iJT-p0bHz1&smid=nytcore-ios-share&referringSource=articleShare`

under these conditions, the platform prefers not to use PAM. In Section 1.5, I demonstrate that even when the platform has full pricing power and can implement complex pricing schemes, it does not use PAM and instead relies on advertising——provided it is sufficiently efficient. Furthermore, when users are overconfident, I show that the platform has an incentive to induce search by lowering fees for high types.

1.4.2 Binary Types

Suppose now that market sides are symmetric. There are only two types on each side of the market and with slight abuse of notation denote the type set by $\Theta = \{\theta_h, \theta_l\}$ with $\theta_h > \theta_l$. Each type has an outside option of zero.[19] In the previous section, I showed that random matching is suboptimal for the platform, while PAM is optimal if the platform charges $s_h = \theta_h^2$ and $s_l = \theta_l^2$.

This section examines the case in which the platform is constrained in setting agents' search costs. In reality, a platform serves many types of users, which would require complex pricing schemes to extract each agent's surplus. I therefore consider a setting in which both types of agents face the same search cost designed by the platform, $s_h = s_l = s$. One possible interpretation is that both types use the basic service of a (freemium) platform. In this case, the platform is assumed to determine the amount of advertising shown to each agent using the basic service. Alternatively, if payments are involved, agents may choose among (discrete) pricing tiers, with all agents on the same tier paying the same amount—as is common on dating platforms. On job platforms, for example, firms often pay the same price per click when advertising a job in a given submarket. To determine how the matching outcome is affected by the platform-chosen matching rule, the analysis fully characterizes all possible matching outcomes in this example.

As in Section 1.4.1, I proceed in two steps. First, I characterize the optimal matching rule that implements the mutual acceptance probability matrices that are consistent with Equation 1.5. Given the first step, I find the optimal matrix of mutual acceptance probabilities that maximize the platform's profit. To identify the optimal matching rule for the platform, suppose for now that s is exogenous.

The first result, Lemma 11,[20] characterizes the optimal matching rule that implements the mutual acceptance probability matrices. With two types, the mutual acceptance matrix takes the following form

$$A = \begin{bmatrix} \alpha(\theta_h, \theta_h) & \alpha(\theta_h, \theta_l) \\ \alpha(\theta_h, \theta_l) & \alpha(\theta_l, \theta_l) \end{bmatrix},$$

[19] The following analysis qualitatively unaffected as long as the outside options are $\omega_l < \theta_l^2$ and $\omega_h < \theta_h \theta_l$. The platform's profit, however, is quantitatively affected as the platform can extract less rent from each agent.

[20] Lemma 11 and its proof can be found in Appendix B.

where the mutual acceptance probability of the assortative matches are along the diagonal and the mutual acceptance probability of mismatches are off the diagonal. Trivially with one type, the mutual acceptance matrix consists only of one entry. With two types, only three possible matrices can be implemented as part of an equilibrium

$$A_{PAM} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, A_{WPAM} = \begin{bmatrix} 1 & \alpha'_{ij} \\ \alpha'_{ij} & 1 \end{bmatrix}, A_{NAM} = \begin{bmatrix} 1 & \alpha''_{ij} \\ \alpha''_{ij} & 0 \end{bmatrix}, \alpha'_{ij}, \alpha''_{ij} \in [0,1].$$

Given the platform's matching rule, high type agents can either accept only other high types, or accept low types with positive probability. This results in three possible constellations of mutual acceptance probabilities and thus matching outcomes. If high type only accept high types, low types will always accept high and types, resulting in a positive assortative matching outcome — only agents of the same type accept each other (A_{PAM}). Depending on the matching rule if high types accept low types with positive probability, low types may accept low types, resulting in a weakly assortative matching outcome — high and low types mutually accept the same types of agents (A_{WPAM}). Alternatively, low types may reject low types, resulting in a non-assortative matching outcome — high types accept low types, but low types do not (A_{NAM}).

For each of the three possible matching outcomes, there exists an optimal matching rule that implements the outcome for a range of parameters (Lemma 11). The implementation of the matching outcomes depends crucially on feasibility. Given the total mass of agents that join, the ratio of new high to low type agents, $0 < \beta_h/\beta_l < \infty$, determines which outcome can be implemented, as the ratio affects the steady state population of both types. The positive assortative matching outcome can be implemented for all $0 < \beta_h/\beta_l < \infty$, whereas the weakly assortative and non-assortative outcomes cannot.

Given the existence of an optimal matching rule, which matrix A maximizes the platform's profit for fix search costs? The next proposition summarizes the results.

Proposition 3. *(i) Let $0 \leq s \leq \theta_l^2$. The platform implements A_{PAM} and the matching outcome, $\mathcal{O}_{PAM}$, is positive assortative if*

$$0 \leq \frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(a)}, \ \ or \ \left(\frac{\beta_h}{\beta_l}\right)^{(b)} \equiv \frac{(1-\delta)(\theta_h^2 - s)}{\theta_h(\theta_h - \theta_l) - s + \delta(\theta_h^2 - s)} \leq \frac{\beta_h}{\beta_l}.$$

The platform implements A_{WPAM} and the matching outcome, $\mathcal{O}_{WPAM}$, is weakly positive assortative if

$$\left(\frac{\beta_h}{\beta_l}\right)^{(a)} \leq \frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(b)}.$$

(ii) Let $\theta_l^2 \leq s \leq \theta_h\theta_l$. If $\beta_h \geq \beta_l$, the platform implements A_{WPAM} and the matching outcome is either weakly assortative, $\mathcal{O}_{WPAM}$, or non-assortative for large enough s,

$\mathcal{O}_{NAM}$. If $\beta_h < \beta_l$, the platform implements A_{WPAM} and the matching outcome is weakly assortative, $\mathcal{O}_{WPAM}$, or only high types participate if s is large enough.
(iii) Lastly, if $\theta_h\theta_l \leq s \leq \theta_h^2$, low types do not participate on the platform. The mutual acceptance matrix and matching outcome is positive assortative.

First, consider the maximum rent that the platform can extract when the positive assortative matrix, A_{PAM}, is implemented. A high type agent is willing to search the longest for a match with another high type. In this case, the maximum rent the platform can extract from a high type agent is proportional to $\theta_h(\theta_h - \theta_l)$, which is the value of its own type times the *match premium*. The match premium is the gain from being in a match with a high type instead of leaving with a low type. If the platform were to extract more rent, high types would start accepting low types as well, and thus only search for one period. Conversely, if high types always reject low types, the maximum rent the platform can extract from low types is proportional to θ_l^2.

Due to feasibility constraints, the platform is constrained by the ratio of high to low types when choosing the matching rule. The platform can extract the rent from both types — as described above — at

$$\left(\frac{\beta_h}{\beta_l}\right)^{(a)} = \frac{(1-\delta)(\theta_l^2 - s)(s + \delta(\theta_h^2 - s))}{(\theta_h(\theta_h - \theta_l) - s - \delta(\theta_h^2 - s))(s + \delta(\theta_l^2 - s))}, \tag{1.6}$$

At this "optimal" ratio, high types are just indifferent between accepting and rejecting low types, while low types are just indifferent between participating or not, which results in

$$\phi(\theta_h|\theta_h) = \frac{s + \delta(\theta_h\theta_l - s)}{(1-\delta)\theta_h(\theta_h - \theta_l)}, \tag{1.7}$$

$$\phi(\theta_l|\theta_l) = \frac{s}{\theta_l^2}. \tag{1.8}$$

Due to feasibility constraints, the incentive and participation constraints cannot generally bind at the same time while implementing a positive assortative matching outcome. As the ratio increases, relatively more high type agents enter compared to low type agents. In this case, high types inevitably meet high types more often, so the platforms makes the participation constraint binding for low types. The platform must increase the probability of a high type meeting a high type such that high types are left with a rent greater than $\theta_h\theta_l$. As the ratio decreases, relatively few high type agents enter compared to low type agents. The platform makes the incentive constraint binding for high types, leaving a positive rent for low types by increasing the probability of a low type meeting a low type. In both cases, the platform potentially forgoes a significant amount of rent when moving away from the "optimal" ratio.

Second, consider the maximum rent that the platform can extract when the weakly positive assortative matrix, A_{WPAM}, is implemented. Suppose the ratio of high to low types is greater than in Equation 1.6. Then, the platform can commit to a matching rule in which high types randomize over accepting and rejecting low types, while low types remain indifferent between participating and their outside option. The expected match utility of high types decreases, while the expected match utility of low types increases. For a ratio of high to low types greater than in Equation 1.6, implementing A_{WPAM} yields a higher profit than A_{PAM}. When implementing A_{PAM}, the platform must increase the meeting probability of assortative pairs as the ratio β_h/β_l increases, otherwise low types will no longer be willing to participate. This implies, however, that the platform forgoes rent from high types. Inducing high types to accept mismatches with positive probability, $\alpha(\theta_h, \theta_l) > 0$, leads to a longer search of low types as they receive a higher expected match utility. Extending the search of low types, implies that there are more low types on the platform, so the platform can also extend the search time of high types.

Third, consider the maximum rent that the platform can extract when the negative assortative matrix, A_{NAM}, is implemented. High types accept both types with positive probability, while low types reject low types and only enter in (mis-)matches with high types. The rent extracted from low types is then proportional to $\theta_l(\theta_h - \theta_l)$, the value of their own type times the *match premium.* The platform, however, never finds it profitable to implement A_{NAM} when it can implement A_{WPAM} as the platform can extract all rent from low types in the latter case, whereas it can only extract the rent premium in the former case. Lastly, if search costs are large, the platform can implement A_{WPAM}, but match low types only to high types if feasible. This in turn results in a non-assortative matching outcome albeit mutual acceptance would be weakly assortative.

For a given inflow of high and low types, $\frac{\beta_h}{\beta_l}$, Proposition 3 presents the matching outcomes that the platform prefers to implement. The assortativity of the matching outcomes is non-monotonic in the ratio of high to low types. For example, the platform can implement the positive assortative matching outcome in markets in which one type dominates. In contrast, the platform implements mismatch in relatively balanced markets.

Corollary 2. *The platform strategically lowers the quality of (recommended) matches. The platform's matching creates two economic inefficiencies: delayed matching and mismatched pairs.*

In other words, the platform recommends mismatches to agents when feasible, i.e., the platform fosters *mismeetings* to delay agent's matches. By delaying matches, the platform increases the payments that it collects from agents per period. Extending users' search, such as in one-sided (matching) markets, has also been shown by for example, Hagiu and Jullien (2011). In addition to mismeetings, the platform also fosters actual mismatches by inducing users to leave in mismatched (inefficient) pairs. Without allowing side payments, mismatches are a way of shifting utility from high to low types to incentivize low type participation.

Next, consider the inefficiencies measured as (i) the amount of mismatch compared to the socially optimal matching and (ii) the length of search for agents. Let the (welfare) loss from mismatch be given by

$$\mathcal{W} = \sum_{(\theta_i,\theta_j)\in\Theta\times\Theta} \alpha(\theta_i,\theta_j)\Phi(\theta_i,\theta_j)(\theta_i\theta_j - \theta_i^2),$$

i.e., the sum over the mass of mismatches times the difference in match utilities between the mismatches and the assortative matches. The expected usage time of an agent is given by their stopping time

$$\mathcal{T}(\theta_i) = \frac{1}{\delta + (1-\delta)\sum_{j=h,l}\alpha(\theta_i,\theta_j)\phi(\theta_j|\theta_i)},$$

such that the total length of search is $\mathcal{T} = \mathcal{T}(\theta_h) + \mathcal{T}(\theta_l)$.

Proposition 4. *(i) If the platform implements A_{PAM} together with matching outcome $\mathcal{O}_{PAM}$, mismatch is $\mathcal{W}_{PAM} = 0$ and $\mathcal{T}(\theta_i)$ is decreasing in s and δ.*
(ii) If the platform implements A_{WPAM} together with matching outcome $\mathcal{O}_{WPAM}$, mismatch is $\mathcal{W}_{WPAM}$ is increasing in s if $\beta_l > \beta_h$ and in- or decreasing in s otherwise as well as decreasing in δ for $s \leq \theta_l^2$ and in- or decreasing in δ otherwise. $\mathcal{T}(\theta_i)$ is decreasing in s and δ.
(iii) If the platform implements A_{WPAM} together with matching outcome $\mathcal{O}_{NAM}$, mismatch is $\mathcal{W}_{NAM} = -\beta_l(\theta_h - \theta_l)^2$ and $\mathcal{T}(\theta_i) = 1$.

By definition, welfare loss is zero under positive assortative matching, as it maximizes total surplus. As search cost or friction δ increases—–both of which lower agents' continuation values—–the platform must raise assortativity and decrease agents' search time to keep low types participating and high types rejecting low types. In the weakly assortative case, assortativity rises with δ, reducing mismatches as long as $s \leq \theta_l^2$. Since the mass of assortative matches varies with s, the mass of mismatches may increase or decrease depending on whether β_h or β_l is larger. In contrast, welfare loss in the non-assortative case is unaffected by search cost or δ, and the platform induces only one period of search.

1.4.3 Advertising and Search Fee

Search Fee If $\nu(s) = s$, the platform charges a linear search fee. It maximizes profit by choosing the fee, given a matching matrix A and the optimal rule that implements it. According to Proposition 3, the platform considers three cases where both types participate: (i) implementing A_{PAM} with the positive assortative outcome $\mathcal{O}(A_{PAM})$, (ii) implementing A_{WPAM} with the weakly positive assortative outcome $\mathcal{O}(A_{WPAM})$, and (iii) implementing A_{WPAM} with the non-assortative outcome $\mathcal{O}(A_{PAM})$. Additionally, there is case (iv) where only high types participate. The next proposition states that the platform can implement any of the four cases depending on the exogenous parameters.

Proposition 5. *For a range of parameters, the platform chooses*
(i) s to maximize $\Pi(A_{PAM})$ *s.t.* $s \in [0, \theta_l^2] : {}^{\beta_h}/_{\beta_l} \leq ({}^{\beta_h}/_{\beta_l})^{(a)}$ *and implements* $\mathcal{O}_{PAM}$*.*
(ii) s to maximize $\Pi(A_{WPAM})$ *s.t.* $s \in [0, \overline{s}] : ({}^{\beta_h}/_{\beta_l})^{(a)} \leq {}^{\beta_h}/_{\beta_l} \leq ({}^{\beta_h}/_{\beta_l})^{(b)}$ *and implements* $\mathcal{O}_{WPAM}$*.*
(iii) $s = \theta_h\theta_l$ *and implements* $\mathcal{O}_{NAM}$*.*
(iv) $s = \theta_h^2$ *and excludes low types from participating.*

The proposition characterizes the platform's optimal solution when s is a uniform search fee paid by agents. When ${}^{\beta_h}/_{\beta_l}$ is relatively low, the platform chooses to implement A_{PAM}. In this case, its profit is bounded by

$$\Pi_{PAM} < \frac{2(1-\delta)}{1-\rho}\left(\beta_h\theta_h(\theta_h-\theta_l)+\beta_l\theta_l^2\right),$$

which corresponds to the maximum surplus the platform can extract as $\delta \to 0$, when high types are indifferent between accepting or rejecting low types, and low types are indifferent between participating or opting out. As in Proposition 3 if ${}^{\beta_h}/_{\beta_l} \geq ({}^{\beta_h}/_{\beta_l})^{(a)}$, the platform can implement A_{WPAM} and outcome $\mathcal{O}_{WPAM}$.

For ${}^{\beta_h}/_{\beta_l} \geq 1$, the platform can implement the non-assortative outcome. As agents only search for one period, the profit given $\mathcal{O}_{NAM}$ is maximized if the search fee is set as high as possible. Thus, the platform chooses $s = \theta_h\theta_l$ yielding a profit of

$$\Pi_{NAM} = \frac{2(1-\delta)}{1-\rho}\left(\beta_h+\beta_l\right)\theta_h\theta_l.$$

With increasing ${}^{\beta_h}/_{\beta_l}$, the platform finds it profitable to charge the highest possible search fee $s = \theta_h^2$ to extract the full surplus from high types, while excluding low types from participating on the platform. This holds as with increasing ${}^{\beta_h}/_{\beta_l}$, the share of revenue from high types grows larger and hence, it becomes more profitable exploiting only one type of users. The platform makes a profit of

$$\Pi_{\theta_h} = \frac{2(1-\delta)}{1-\rho}\beta_h\theta_h^2.$$

Advertising Now, suppose the platform sells the attention of its users to advertisers. The platform decides on the advertising intensity, which is related to the search cost that users experience. Let $\nu(s)$ be the revenue per unit of search cost to users. Recall that $\nu(s)$ is an increasing, concave function of search cost s with $\nu(0) = 0$. This assumption excludes functions that are convex, i.e., under which the platform could prefer an advertising intensity that induces users to stay for only one period, thereby significantly reducing the mass of active users. The platform maximizes its profit with respect to s

$$\max_s \frac{2\nu(s)(1-\delta)}{1-\rho}\left(\sum_{\theta_i\in\Theta} f(\theta_i^k)(s)\right),$$

where the mass of agents of type θ_i is given by Lemma 3 subject to the conditions in Proposition 3. The platform chooses $s = s^A$ such that

$$\frac{\nu(s^A)}{\nu'(s^A)} = -\frac{\sum_k \sum_i f(\theta_i^k)}{\partial \sum_k \sum_i f(\theta_i^k) / \partial s}\Big|s = s^A \,. \tag{1.9}$$

Under the above condition, the marginal cost of an increase in search cost, given by the semi-elasticity of demand on the right-hand side, is equal to the marginal benefit of an increase in search cost, given by the semi-elasticity of advertising revenue.

Proposition 6. *If $\frac{\nu(s)}{\nu'(s)} \geq s$ for $s \in [0, \theta_h^2]$, the platform chooses search costs that are lower or equal than a uniform search fee. Furthermore, if $\nu(\theta_h\theta_l)/\nu(\theta_h^2) \geq \beta_h/\beta_h+\beta_l$, i.e. $\nu(\cdot)$ is sufficiently concave at high search costs, the platform finds it profitable to never exclude low types from the search process.*

1.5 Explanations

If $\nu(s_i^k) = s_i^k$, the optimal contract is a set of personalized search fees. The platform maximizes the total match surplus as in Appendix A.2 and extracts the surplus from each agent via the fee. Considering the simplified model from Section 1.4.2, the platform commits to the positive assortative matching rule and personalized fees ($s_h = \theta_h^2, s_l = \theta_l^2$) (see Proposition 2). Under the positive assortative matching rule, agents meet their match in the first period. The platform's profit is

$$\Pi^{PAM} = \frac{2(1-\delta)}{(1-\rho)}(\beta_h\theta_h^2 + \beta_l\theta_l^2).$$

1.5.1 Advertisement

Advertisement plays a key role in the digital economy. More specifically, in the light of the application to dating and job search platforms, a substantial share of these platforms rely on advertisement as a source of revenue, see Appendix C for an overview of dating and job search apps that show advertisement. In the following example, I highlight that a (partly) advertising-based business model can outperform profits generated by personalized prices demonstrated by the following example.

Example 1. Consider the concave function $\nu(s) = \kappa s^\alpha$ for $\alpha = \frac{1}{2}$. Figure 1.2 plots the function for different values of κ. Furthermore, assume that $\beta_h < \beta_l$ and let the value of a high type, $\theta_h = 2$, be twice as large as the value of a low type, $\theta_l = 1$. Denote the ${}^{1-\delta}/_{1-\rho} = \gamma$.

For $\beta_h < \beta_l$, the platform either implements A_{PAM} or A_{WPAM}. To maximize advertising profits, the platform chooses $s^A \in [0, \theta_l^2]$ to solve

$$\frac{\beta_h}{\beta_l} = \frac{(1-\delta)(\theta_l^2 - s^A)(s + \delta(\theta_h^2 - s^A))}{(\theta_h(\theta_h - \theta_l) - s^A - \delta(\theta_h^2 - s^A))(s^A + \delta(\theta_l^2 - s^A))},$$

at which the agents' search time under A_{PAM} and A_{WPAM} coincides. Furthermore recall that if the condition is satisfied, the agents' search time is maximized as low types are indifferent between participating or not and high types are indifferent between accepting and rejecting low types (and rejecting with probability one). The platform's advertising profit is

$$\Pi^A = 2\gamma\kappa\sqrt{s^A}\left(\frac{\beta_h\theta_h(\theta_h - \theta_l)}{s^A + \delta(\theta_h^2 - s^A)} + \frac{\beta_l\theta_l^2}{s^A + \delta(\theta_l^2 - s^A)}\right).$$

For the chosen parameters, s^A is equal to $\theta_l^2 = 1$ if $\beta_h = 0$ and strictly larger than zero for β_h approaching β_l. The profits for $\beta_h = 0$ are

$$\begin{aligned}\Pi^A(\beta_h = 0) &= \gamma\kappa\sqrt{\theta_l^2}\beta_l = \gamma\kappa\beta_l,\\ \Pi^{PD}(\beta_h = 0) &= \gamma\beta_l\theta_l^2 = \gamma\beta_l,\end{aligned}$$

which coincide for $\kappa = 1$. Thus, for $\kappa > \underline{\kappa} = 1$, advertising profits are larger than the profits of the optimal contract for some $\beta_l > 0$. Now, let β_h approach β_l, the profits are

$$\begin{aligned}\Pi^A(\beta_h = 0) &= 2\gamma\kappa\sqrt{s^A}\left(\frac{\beta_h\theta_h(\theta_h - \theta_l)}{s^A + \delta(\theta_h^2 - s^A)} + \frac{\beta_l\theta_l^2}{s^A + \delta(\theta_l^2 - s^A)}\right),\\ \Pi^{PD}(\beta_h \to \beta_l) &= 4\gamma\beta_l(\theta_h^2 + \theta_l^2).\end{aligned}$$

Then, there exists a $\kappa > \overline{\kappa}$ such that advertising profits are larger than the profits of the optimal contract for all $\beta_h \in [0, \beta_l)$. For the values in this example, $\overline{\kappa} \approx {}^3/_2$.

For general revenue functions $\nu(s)$, an advertisement-based business model generates higher profits than charging personalized prices if advertisement revenue is sufficiently efficient compared to its nuisance:

$$\frac{\nu(s)}{s} \geq \frac{\beta_h\theta_h^2 + \beta_l\theta_l^2}{s(\mathcal{T}(\theta_h) + \mathcal{T}(\theta_l))},$$

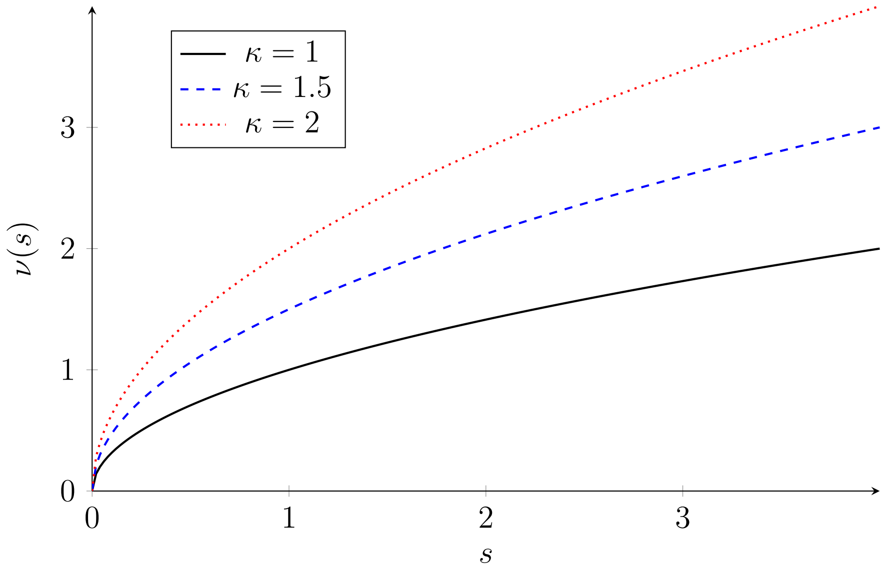

Figure 1.2: $\nu(s) = \kappa\sqrt{s}$ for different κ

where the numerator is the full surplus that can be extracted from agents under PAM with personalized fees and the denominator is the total amount of search cost that agent's pay while searching under advertising. Note that if the market is extremely unbalanced, i.e. if only high types are in the market, advertising is less profitable as long as $\nu(\theta_h^2)/\theta_h^2 \leq 1$.

1.5.2 Overconfidence

Up to this point, the model has assumed that agents behave rationally and have a correct expectation about their own type. In the following, I will introduce a fraction of overconfident agents, i.e., agents who perceive themselves to be of a higher type than they actually are. In the simplest example, an overconfident low type perceives itself as a high type. Overconfidence is a widely documented bias in the psychology and behavioral economics literature.[21]

Especially in dating markets and labor markets overconfidence is thought to be prevalent for example, when it comes to a person's own attractiveness or ability. In dating markets, both women and men prefer attractive over unattractive profiles regardless of their own attractiveness (Egebark et al., 2021). Bruch and Newman (2018, 2019) analyze the structure of online dating markets in US cities and provide suggestive evidence for the fact that the majority of users contacts a partner who is more desirable than they are instead of contacting a partner who is as desirable than they are. One possible explana-

[21] Ample evidence suggests that on average agents overestimate their ability, traits and prospects. Such overconfidence has been documented in laboratory experiments by Burks et al. (2013); Dubra (2015); Charness et al. (2018). Additionally, there is empirical evidence that consumers are overoptimistic regarding future self-control when signing up for a gym membership (DellaVigna and Malmendier, 2006), workers overpredict their own productivity (Hoffman and Burks, 2020), and some CEOs are overoptimistic regarding their firm's performance (Malmendier and Tate, 2005, 2008).

tion is documented by Greitemeyer (2020), that is, more unattractive people are unaware of their (un-)attractiveness from a psychological perspective. Similarly in labor markets, Spinnewijn (2015) and Mueller et al. (2021) find that the unemployed overestimate how quickly they will find a job and are persistently overconfident about their desirability to firms. In line with the empirical evidence, Dargnies et al. (2019) document in a laboratory experiment that agents who are overconfident are less likely to accept earlier job offers in a matching market.

Following this evidence, consider the following simple extension to the model in Section 1.4.2. There exists a symmetric share of λ overconfident users on each side of the market. An overconfident user has type θ_l, but persistently believes to have type θ_h, i.e. is stubborn and does not learn their true type. Denote the overconfident type by $\hat{\theta}_l$. Other agents correctly identify overconfident types as low types. Following Definition 3, an overconfident type chooses their strategy confidently believing in their misperceived type. As a result of overestimating their own type, they, however, are overoptimistic about the likelihood of being accepted by others. As before, users incur search costs and become inactive with probability δ.[22]

As overconfidence has been identified in empirical and experimental setting, I suppose that the platform can perfectly identify overconfident users as well. The platform chooses matching rule $\mathcal{M}$, which consists of $\phi(\cdot|\theta_i)$ for $\theta_i \in \{\theta_l, \theta_h, \hat{\theta}_l\}$, and search costs (s_h, s_l). As a benchmark, suppose the platform induces only one period of search by charging $(s_h = \theta_h^2, s_l = \theta_l^2)$ and choosing the positive assortative matching rule in which high types only meet each other and (true) low types, which includes overconfident types, only meet each other. The platform's profits are

$$\Pi_{PAM}^{OC} = \frac{2(1-\delta)}{1-\rho}(\beta_h \theta_h^2 + \beta_l(1-\lambda)\theta_l^2 + \beta_l \lambda \theta_h^2).$$

To show that the platform can improve on this profit, let the platform induce search by inducing high types to reject low types. The matching rule and search costs must satisfy the participation constraint of low types and the incentive constraint of high types

$$\theta_h \theta_l \leq \frac{(1-\delta)(-s + \phi(\theta_h|\theta_h)\theta_h^2)}{\delta + (1-\delta)\phi(\theta_h|\theta_h)}, \tag{IC-θ_h}$$

$$0 \leq \frac{(1-\delta)(-s + \phi(\theta_l|\theta_l)\theta_l^2)}{\delta + (1-\delta)\phi(\theta_l|\theta_l)}. \tag{1.10}$$

Given both constraints are satisfied, the participation constraint of high types and the incentive constraint of low types (to reject low types) are satisfied as well. Next, consider the acceptance behavior of an overconfident type. Given their perception of the game,

[22]Note that δ can have an additional interpretation in the presence of overconfident users. If overconfident users do not find a match, δ can be interpreted as the probability that an overconfident agent leaves due to growing dissatisfaction with the platform.

rejecting low types is perceived optimal if

$$\theta_h\theta_l \leq \frac{(1-\delta)(-s+\phi(\theta_h|\theta_h)\theta_h^2)}{\delta+(1-\delta)\phi(\theta_h|\theta_h)}, \tag{PIC-$\hat{\theta}_h$}$$

which coincides with the incentive constraint of high types. Similarly, they face the same perceived participation constraint. The payoff from participation is

$$-\frac{s}{\delta} < 0,$$

since overconfident users reject low types, but high types never accept overconfident types. This leads them to search until they exogenously exit with probability δ.

Remark. Overconfident users search too intensively.

Proposition 7. *(Overconfidence) Let* $\lambda^* \equiv \frac{\beta_h}{\beta_l}\frac{\delta\theta_h\theta_l}{(1-\delta)\theta_h^2-\theta_h\theta_l}$. *For* $\lambda < \lambda^*$, *the platform maximizes profits by setting* $(s_h = \theta_h^2, s_l = \theta_l^2)$ *and inducing only one period of search. The platform's profit is* Π_{PAM}^{OC}. *For* $\lambda \geq \lambda^*$, *the platform maximizes profits by setting* $(s_h = \theta_h(\theta_h-\theta_l) - \delta/1-\delta\theta_h\theta_l, s_l = \theta_l^2)$ *and inducing search from overconfident users. The platform's profit is*

$$\Pi_S^{OC} = \frac{2(1-\delta)}{1-\rho}\left(\beta_h(\theta_h(\theta_h-\theta_l) - \frac{\delta}{1-\delta}\theta_h\theta_l) + \beta_l(1-\lambda)\theta_l^2 + \frac{\beta_l\lambda(\theta_h(\theta_h-\theta_l)-\frac{\delta}{1-\delta}\theta_h\theta_l)}{\delta}\right).$$

Anecdotes from Dating Apps, such as Tinder, provide evidence for the fact that less than 10% of users account for a disproportional amount of revenue.[23] On Tinder, an average user spends around 30$ in in-app purchases and subscriptions, whereas "heavy" users would spend 10 times the amount.

Consider the following example to illustrate that in markets with many low types, already a small percentage of overconfident users can be sufficient to achieve higher profits.

Example 2. Let $\beta_h = \frac{1}{4}, \beta_l = \frac{3}{4}, \delta = 1/10, \theta_h = 2, \theta_l = 1$. Then, $\lambda \geq 4.2\%$. For low values of δ, a relatively small percentage of overconfident users is necessary to substantially increase the platforms profit. Note that δ is directly related to the stopping time of overconfident users, i.e. overconfident users search for ten periods before they exit. More generally, consider the following comparative statics.

Corollary 3. λ^* *increases in* δ, *and* $\frac{\beta_h}{\beta_l}$.

Intuitively, the necessary share of overconfident users decreases if δ becomes small as overconfident users search for more periods. If the ratio $\frac{\beta_h}{\beta_l}$ increases, i.e. there are more

[23]See https://uxdesign.cc/how-tinder-drives-over-1-6-billion-in-revenue-8006e718e761 and the referenced podcast therein, https://open.spotify.com/episode/1ZfL2Mq1n0NzyVKKerynvZ?si=UBlpCunARLW8jPfNNYK4dw.

high types than low types in the market, the platform needs to rely more on overconfident users. The reason is that given that the platform lowers the search fee for high types to exploit overconfident users, they become less profitable. Hence, with more high types, there must be more overconfident types to offset the loss from high types.

1.6 Conclusion

On matching platforms, the misalignment of incentives between users and the platform becomes more problematic as platforms collect more data and develop more predictive algorithms. This paper presents a model in which a platform has perfect information about its users' types and matches them to its advantage. In contrast, random matching corresponds to the case where the platform has no information about its users' types. The platform benefits from more information about its users' types: Random matching is strictly suboptimal.

Both sorting and search time have implications for real-world markets. The platform's algorithm can support the socially optimal matching. But even absent exogenous search costs and search frictions, the algorithm can also foster non-assortative matching outcomes in fully symmetric markets resulting in mismatch. Additionally, it increases users' search time by recommending unsuitable matches. While mismatch has a negative impact on productivity and long-term unemployment in labor markets (Şahin et al., 2014; McGowan and Andrews, 2015), assortative mating in marriage markets is a driver of household inequality (Pestel, 2017; Eika et al., 2019; Almar et al., 2023). Therefore, if policies aim to reduce mismatch — as in labor markets — policymakers should be concerned about matching platforms that employ the business models described above. Rather than relying on platforms to reduce search frictions, the platform's algorithm is a potential source of additional mismatch. In contrast, dating apps can make a positive contribution to reducing household inequality.

Empirical evidence on online matching and search platforms is mixed. For example, in dating markets Hitsch et al. (2010) show that matches are approximately efficient and stable. The authors, however, rely on data before the advent of large dating apps. In contrast, more recent evidence, such as Sharabi and Dorrance-Hall (2024), finds that people who meet online are less satisfied in their marriages. In labor markets, Kroft and Pope (2014) show that Craigslist has no effect on the unemployment rate. Similarly, Gürtzgen et al. (2021) provide evidence that online searches do not affect employment stability or wage outcomes, but instead increase the proportion of unsuitable candidates in job applications.

A Appendix

A.1 Linear Programming Formulation

The linear programming formulation of the platform's problem in Lemma 4 is given in the following. For $\alpha(\theta_i^k, \theta_j^{-k}) \in \{0,1\}$, the platform's optimization problem can be represented by the following (mixed integer) linear program:

$$\max_{\{\Phi(\cdot),m(\cdot)\}_{ij}^k} \sum_{k=A,B} \sum_{\theta_i^k \in \Theta^k} \frac{(1-\delta)\nu(s_i^k)}{1-\rho} f(\theta_i^k), \tag{1.11}$$

subject to participation constraints

$$\beta_i^k \omega_i^k \leq f(\theta_i^k)(\delta\omega_i^k - (1-\delta)s_i^k) + (1-\delta)\sum_j m(\theta_i^k, \theta_j^{-k})\theta_i^k\theta_j^{-k}, \forall \theta_i^k \in \Theta^k,\ k = A, B, \tag{1.12}$$

incentive constraints

$$\begin{aligned} \beta_i^k\theta_i^k\theta_j^{-k} + \alpha(\theta_i^k, \theta_j^{-k})(-\beta_i^k\theta_i^k\theta_j^{-k}) \leq & f(\theta_i^k)(\delta\omega_i^k - (1-\delta)s_i^k) + (1-\delta)\sum_j m(\theta_i^k, \theta_j^{-k})\theta_i^k\theta_j^{-k} \\ \leq & \left(\frac{\beta_i^k}{\delta}\theta_i^k\theta_j^{-k} - \beta_i^k\theta_i^k\theta_j^{-k}\right)(1-\alpha(\theta_i^k, \theta_j^{-k})) + \beta_i^k\theta_i^k\theta_j^{-k}, \end{aligned} \tag{1.13}$$

$$\forall \theta_i^k \in \Theta^k,\ k = A, B,$$

feasibility and steady state constraints

$$\sum_{\theta_j^{-k} \in \Theta^{-k}} \Phi(\theta_i^k, \theta_j^{-k}) + \mathbf{1}_{k=B}\Phi(\theta_i^k, \omega_i^k) = f(\theta_i^k), \forall \theta_i^k \in \Theta^k,\ k = A, B, \tag{1.14}$$

$$f(\theta_i^k) = \frac{\beta_i^k - (1-\delta)\sum_j m(\theta_i^k, \theta_j^{-k})}{\delta}, \forall \theta_i^k \in \Theta^k,\ k = A, B, \tag{1.15}$$

and constraints on the matched and recommended pairs $\forall(\theta_i^k, \theta_j^{-k}) \in \Theta^k \times \Theta^{-k}$. First, the mass of recommended and matched pairs must be non-negative and the mass of matched pairs cannot be greater than the mass of recommended pairs

$$\Phi(\theta_i^k, \theta_j^{-k}) \geq 0, m(\theta_i^k, \theta_j^{-k}) \geq 0, \tag{1.16}$$

$$m(\theta_i^k, \theta_j^{-k}) \leq \Phi(\theta_i^k, \theta_j^{-k}). \tag{1.17}$$

Second, the mass of matched pairs must be smaller than the largest possible mass of the agents, i.e. the mass that arises when agents only exit upon becoming inactive β_i^k/δ times the acceptance probability, and larger than the mass of recommended pairs minus the largest possible mass times the probability of a rejection

$$m(\theta_i^k, \theta_j^{-k}) \leq \frac{\min\{\beta_i^k, \beta_j^{-k}\}}{\delta}\alpha(\theta_i^k, \theta_j^{-k}), \tag{1.18}$$

$$m(\theta_i^k, \theta_j^{-k}) \geq \Phi(\theta_i^k, \theta_j^{-k}) - \frac{\min\{\beta_i^k, \beta_j^{-k}\}}{\delta}(1 - \alpha(\theta_i^k, \theta_j^{-k})). \tag{1.19}$$

This ensures that the mass of matched pairs must be smaller than the mass of recommended pairs and that for $\alpha(\theta_i^k, \theta_j^{-k}) = 0$ the mass of matched pairs cannot be greater than zero. To accommodate for mixed acceptance probabilities of agents, consider an agent of type θ_m^k that is indifferent between accepting and rejecting a type θ_s^{-k}. Hence, θ_m^k could randomize over the acceptance probability towards type θ_s^{-k}: $\sigma_k(\theta_m^k, \theta_s^{-k}) \in (0,1)$. Conceptually, this imposes indifference or equality on some constraints rather than inequalities in the original formulation above. For any pair $(\theta_m^k, \theta_s^{-k}) \in \Theta^k \times \Theta^{-k}$ for which $\alpha(\theta_m^k, \theta_s^{-k}) \in (0,1)$, the adjusted incentive constraints are

$$\beta_m^k \theta_m^k \theta_s^{-k} = f(\theta_m^k)(\delta\omega_m^k - (1-\delta)s_m^k) + (1-\delta)\sum_j m(\theta_m^k, \theta_j^{-k})\theta_m^k \theta_j^{-k}, \text{ for } \theta_m^k, \tag{1.20}$$

$$\beta_s^{-k} \theta_m^k \theta_s^{-k} \geq f(\theta_s^{-k})(\delta\omega_s^{-k} - (1-\delta)s_s^{-k}) + (1-\delta)\sum_j m(\theta_m^k, \theta_j^{-k})\theta_m^k \theta_j^{-k}, \text{ for } \theta_s^k, \tag{1.21}$$

where θ_m^k is indifferent between accepting and rejecting θ_s^{-k} and θ_s^{-k} (weakly) accepts θ_m^k. The constraints on the mass of recommended and matched pairs are

$$m(\theta_m^k, \theta_s^{-k}) \leq \frac{\min\{\beta_m^k, \beta_s^{-k}\}}{\delta}, \text{ for } (\theta_m^k, \theta_s^{-k}), \tag{1.22}$$

$$m(\theta_m^k, \theta_j^{-k}) \leq \Phi(\theta_m^k, \theta_s^{-k}), \text{ for } (\theta_m^k, \theta_s^{-k}). \tag{1.23}$$

The linear program can be summarized in the subsequent lemma.

Lemma 6 (Linear Program)**.** *Fix any mutual acceptance matrix A. The platform's maximization problem yields the same profit as linear programming problem with objective function in Equation 1.11 subject to constraints Equation 1.12 through 1.16 for any* $\alpha(\theta_i^k, \theta_j^{-k}) \in \{0,1\}$, *and for any pair* $(\theta_m^k, \theta_s^{-k}) \in \Theta^k \times \Theta^{-k}$ *for which* $\alpha(\theta_s^k, \theta_m^{-k}) \in (0,1)$, *replace Equation 1.13 for* θ_m^k *by Equation 1.20 and replace Equation 1.13 for* θ_s^k *by Equation 1.21 and replace Equations 1.18 to 1.19 for* $(\theta_m^k, \theta_s^{-k})$ *by Equations 1.22 to 1.23.*

Note on Standard Form of a Linear Program To abbreviate future arguments, I relate the linear program by the standard form of a linear program. The matrix notation is

$$\max xc^T,$$

$$s.t. Hx \leq b, x \geq 0,$$

where $c \in \mathbb{R}^n$. The variable vector $x \in X \subset \mathbb{R}^n$ consists of n variables, i.e., the mass of recommended and matched pairs, and is an element of the compact set X as each mass takes a value in $[0, \beta_i/\delta]$. The m inequalities are given by matrix $H \in \mathbb{R}^{m \times n}$. Equalities, such as the feasibility constraints, can be expressed as two opposite inequalities. Vector $b \in \mathbb{R}^m$ captures the right-hand side of the inequalities. $\mathcal{P} \equiv \{x \in \mathbb{R}^n | Hx \leq b\}$ is the feasible region given by the inequality constraints.

A.2 Benchmarks

This section analyzes two polar cases, in which the intermediary has full information about agent's types and is able to extract the full rent from the matching output or the intermediary has no information about agent's types and must match agents at random.

Socially-Optimal Matching The first benchmark constitutes the case in which the intermediary (or a social planner) provides the socially-optimal matching under the premise that agent's types can be identified perfectly. The intermediary or social planner maximizes the sum of total matching outputs given that agents only search for one period. The matching output function is supermodular, i.e. types of both sides are complements. The socially-optimal matching is the solution to the linear program

$$\max_{M} \sum_{k=A,B} \sum_{\theta_j^{-k} \in \Theta^{-k}} \sum_{\theta_i^k \in \Theta^k} (\theta_i^k \theta_j^{-k} - \omega_i^k) m(\theta_i^k, \theta_j^{-k}) \tag{1.24}$$

subject to feasibility

$$\sum_{\theta_j^{-k} \in \Theta^{-k}} m(\theta_i^k, \theta_j^{-k}) \leq \beta_i^k, \forall \theta_i^k \in \Theta^k, \tag{1.25}$$

$$\sum_{\theta_i^k \in \Theta^k} m(\theta_i^k, \theta_j^{-k}) \leq \beta_j^{-k}, \forall \theta_j^{-k} \in \Theta^{-k}, \tag{1.26}$$

$$m(\theta_i^k, \theta_j^{-k}) \geq 0, \forall (\theta_i^k, \theta_j^{-k}) \in \Theta^k \times \Theta^{-k}. \tag{1.27}$$

The linear program follows the optimal assignment problem by Koopmans and Beckmann (1957) and Shapley and Shubik (1971). Both agents that form the match $(\theta_i^k, \theta_j^{-k})$ receive the output $\theta_i^k \cdot \theta_j^{-k}$.

Remark. If markets are fully symmetric, the socially optimal matching is $m(\theta_i^k, \theta_j^{-k}) = \beta_i^k$ if $\theta_i^k = \theta_j^{-k}$. The outcome is said to exhibit positive assortative matching.

If market sides are fully symmetric, $\beta_i^A = \beta_i^B$, the solution to the linear program is attained with $m(\theta_i^k, \theta_j^{-k}) \in \{0, \beta_i^k\}$, that is a pair is either matched with probability one or not matched. Although the linear program permits partial or fractional matching of agents, Dantzig (1963) showed that the maximum value of the objective is attained with probabilities in $\{0, 1\}$.

For symmetric populations of agents, optimality requires that no individual remains unmatched, such that the feasibility constraints must hold with equality. Otherwise, the social planner can increase welfare by assigning an unmatched agent to another unmatched agent as the value of their match is greater than zero. The objective is maximized if $m(\theta_i^k, \theta_j^{-k}) = \beta_i^k$ when $\theta_i^k = \theta_j^{-k}$ by applying the rearrangement inequality (Hardy et al., 1952).

Random Matching The second benchmark is a random matching market. For example, if an intermediary has no information (data) about agents' types, and thus cannot condition on any observables, the intermediary's matching rule incorporates random meetings between agents. A random matching market may also reflect offline meetings between agents that are not intermediated by any platform.

A random matching market is a tuple $(\hat{\Theta}^k, f(\theta_i^k))_{k=A,B}$ with parameters (s_i^k, δ).The analysis builds on the model of Lauermann and Nöldeke (2014).[24]

The total mass of agents on side k is $\overline{f}^k = \sum_{\theta_i^k \in \Theta^k} f(\theta_i^k)$. Since each agent can meet at most one agent per unit of time, the total mass of meetings is given by $\min\{\overline{f}^A, \overline{f}^B\}$. Given that meetings are random, the fraction of meetings that involve type θ_i^k on side k and type θ_j^{-k} on side $-k$ is then

$$\frac{f(\theta_i^k) f(\theta_j^{-k}) \min\{\overline{f}^k, \overline{f}^{-k}\}}{\overline{f}^k \cdot \overline{f}^{-k}}.$$

If $\overline{f}^k > \overline{f}^{-k}$, then the mass of agents on side k that meet their outside option is $\Phi(\theta_i^k, \omega_i^k) = \frac{\overline{f}^k - \overline{f}^{-k}}{\overline{f}^k}$. The probability to meet type θ_j^{-k} on side $-k$ conditional on being an agent of any type on side k is

$$\phi(\theta_j^{-k}) = \frac{f(\theta_j^{-k})}{\overline{f}^{-k}} \frac{\min\{\overline{f}^k, \overline{f}^{-k}\}}{\overline{f}^k},$$

where the probability that type θ_i^k on side k exits the search process in a match with type θ_j^{-k} is

$$\mu(\theta_i^k, \theta_j^{-k}) = \frac{(1-\delta)\alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k})}{\delta + (1-\delta)\sum_{\theta_j^{-k}} \alpha(\theta_i^k, \theta_j^{-k})\phi(\theta_j^{-k})},$$

where $\mu(\theta_i^k, \omega_i^k) = 1 - \sum_{\theta_j^{-k}} \mu(\theta_i^k, \theta_j^{-k})$ is the probability that type θ_i^k remains unmatched.

[24] In contrast to Lauermann and Nöldeke (2014), agents may face explicit search cost s_i^k in addition to δ.

Let $(f(\theta_i^k), \alpha(\theta_i^k, \theta_j^{-k})_{ij})_{k=A,B}$ be a steady state. Then M with entries given by

$$m(\theta_i^k, \theta_j^{-k}) = \frac{\alpha(\theta_i^k, \theta_j^{-k}) f(\theta_i^k) f(\theta_j^{-k}) \min\{\overline{f}^k, \overline{f}^{-k}\}}{\overline{f}^k \cdot \overline{f}^{-k}}. \tag{1.28}$$

is the unique matching outcome induced by the steady state under random matching. Vice versa, if M is a steady state matching outcome then $f(\theta_i^k), \alpha(\theta_i^k, \theta_j^{-k})$ is given by

$$f(\theta_i^k) = \frac{\beta_i^k}{\delta} \mu(\theta_i^k, \omega_i^k), \tag{1.29}$$

$$\alpha(\theta_i, \theta_j) = m(\theta_i^k, \theta_j^{-k}) \frac{\overline{f}^k \cdot \overline{f}^{-k}}{f(\theta_i^k) f(\theta_j^{-k}) \min\{\overline{f}^k, \overline{f}^{-k}\}}, \tag{1.30}$$

where $\alpha(\theta_i^k, \theta_j^{-k}) \leq 1$ for all $(\theta_i^k, \theta_j^{-k}) \in \hat{\Theta}^k \times \hat{\Theta}^{-k}$ and $m(\theta_i^k, \omega_i^k)$ is the probability of ending up with one's outside option. Matching M is an **equilibrium** matching if

$$m(\theta_i^k, \theta_j^{-k}) = \begin{cases} 0 & \text{if } \theta_i^k \theta_j^{-k} < V^C(\theta_i^k) \text{ or } \theta_i^k, \theta_j^{-k} < V^C(\theta_j^{-k}) \\ \frac{f(\theta_i^k) f(\theta_j^{-k}) \min\{\overline{f}^k, \overline{f}^{-k}\}}{\overline{f}^k \cdot \overline{f}^{-k}} & \text{if } \theta_i^k \theta_j^{-k} > V^C(\theta_i^k) \text{ and } \theta_i^k \theta_j^{-k} > V^C(\theta_j^{-k}) \end{cases}$$

holds for all $(\theta_i^k, \theta_j^{-k}) \in \hat{\Theta}^k \times \hat{\Theta}^{-k}$.

B Appendix: Omitted Proofs

B.1 Multiple Types

Proof of Lemma 1 and 2 in the text.

Proof of Lemma 3. If $\eta_i^k < 1$ and $\Phi(\theta_i^k, \omega_i^k) \geq 0$ are optimal for any $\theta_i^k \in \Theta^k$, then $\eta_i^k = 1$ and $\Phi'(\theta_i^k, \omega_i^k,)$ are also optimal such that

$$\Phi(\theta_i^k, \theta_j^{-k}) = \Phi'(\theta_i^k, \theta_j^{-k}), \tag{1.31}$$

$$(1 - \eta_i^k) f(\theta_i^k) + \Phi(\theta_i^k, \omega_i^k) = \Phi'(\theta_i^k, \omega_i^k), \tag{1.32}$$

for all $\theta_i^k \in \Theta^k$ and $\theta_j^{-k} \in \Theta^{-k}$. For given $\eta_i^k < 1$ and matching rule $\mathcal{M}$, Equation 1.31 and 1.32 determine the new matching rules for $\eta_i^k = 1$.

Now consider the participation for type θ_i^k. In equilibrium, the participation constraint must be binding for agents to find it optimal to randomize in their participation decision. Suppose the participation constraint is binding, then it can be rewritten as

$$(1 - \delta) s_i^k = (1 - \delta) \sum_j \alpha(\theta_i^k, \theta_j^{-k}) \phi(\theta_j^{-k} | \theta_i^k) \left(\theta_i^k \theta_j^{-k} - \omega_i^k \right).$$

As the masses are the same by Equation 1.31, the total surplus extracted by the platform remains the same as optimality requires that the participation constraint continues to bind. Multiplying with the total mass of agents of type θ_i^k if $\eta_i^k < 1$ yields

$$(1-\delta)\eta_i^k f(\theta_i^k) s_i^k = \underbrace{(1-\delta)\sum_j \alpha(\theta_i^k, \theta_j^{-k})\Phi(\theta_j^{-k}|\theta_i^k)\left(\theta_i^k\theta_j^{-k} - \omega_i^k\right)}_{\text{Total Surplus}}.$$

Similarly, when multiplying with the total mass of agents of type θ_i^k if $\eta_i^k = 1$ yields

$$(1-\delta)f'(\theta_i^k)s_i^{k,\prime} = (1-\delta)\sum_j \alpha(\theta_i^k, \theta_j^{-k})\Phi(\theta_j^{-k}|\theta_i^k)\left(\theta_i^k\theta_j^{-k} - \omega_i^k\right).$$

Therefore, the total surplus extracted is the same in both cases by construction. Thus, if the platform charges a search fee both cases yield the same surplus.

In the case of advertising note that $f'(\theta_i^k)$ must increase if η_i^k increases, i.e. the steady-state mass increases if more agents participate everything else equal. Rewrite equation 1.31 as

$$\eta_i^k f(\theta_i^k)\phi(\theta_j^{-k}|\theta_i^k) = f'(\theta_i^k)\phi'(\theta_j^{-k}|\theta_i^k)$$

Therefore, to fulfill the equality in Equation 1.31 $\phi'(\theta_j^{-k}|\theta_i^k)$ must decrease to decrease the right-hand side. This implies that $s_i^{k,\prime} < s_i^k$ and therefore, the platform profit increases in the advertising case due to the concavity of $\nu(s_i^k)$. $\square$

Proof of Lemma 5 As defined in the Section 1.4.1, the set $\mathcal{G}$ is the set of profit levels following from all linear programs with $A \in \mathcal{A}^*$. I show that the set $\mathcal{G}$ is (i) non-empty with $\Pi(A) < \infty$ for all $A \in \mathcal{A}^*$ and $-\infty < \Pi(A)$ for at least one $A \in \mathcal{A}^*$ and (ii) finite.

To define set $\mathcal{G}$, recall the following definitions from the text. (i) Define a subset $\mathcal{A}^* \subset \mathcal{A}$, where $\mathcal{A}$ are the mutual acceptance matrices that can be implemented by a matching mechanism $\mathcal{M}$. Construct $\mathcal{A}^*$ through the following procedure: For every $A' \in \mathcal{A}$, construct a matrix A'' such that

$$\begin{aligned}
&\alpha'(\theta_i^k, \theta_j^{-k}) = \alpha''(\theta_i^k, \theta_j^{-k}) \text{ if } \alpha'(\theta_i^k, \theta_j^{-k}) \in \{0,1\},\\
&\alpha'(\theta_i^k, \theta_j^{-k}) = \alpha_{ij} \text{ otherwise,}
\end{aligned}$$

where α_{ij} is a variable in $[0,1]$. (ii) For each $A \in \mathcal{A}^*$, the linear program is given by Lemma 6. The value of the objective is given by $\Pi(A)$. Then, (iii) $\mathcal{G} = \bigcup_{A\in\mathcal{A}^*}\Pi(A)$.

(a) $\mathcal{G}$ is non-empty.

I will show that for any $A \in \mathcal{A}^*$, there exists an optimal value $\Pi(A) < \infty$ to the linear program. To do so, fix $A \in \mathcal{A}^*$ and consider the linear program as defined in Lemma 6 in Appendix A.1. To prove that an optimal solution exists, I show that: (i) the objective of the linear program is bounded, i.e., the linear program is not unbounded, and (ii) the feasible region of the variable vector, $\mathcal{P}$, is non-empty for a range of parameters. From both it follows that there exists an optimal solution by Dantzig (1963); Bertsimas and Tsitsiklis (1997).

(i) First, I show that the objective is bounded for all linear programs for fix $A \in \mathcal{A}^*$. For a maximization problem to be bounded there must exists a constant $C \in \mathbb{R}$ such that for all feasible $x \in \mathbb{R}^n$ $c^T x \leq C$ holds. The objective is bounded as

$$\sum_{k=A,B} \sum_{\theta_i^k \in \Theta^k} \frac{(1-\delta)s_i^k}{(1-\rho)} f(\theta_i^k) < \sum_{k=A,B} \sum_{\theta_i^k \in \Theta^k} \frac{(1-\delta)s_i^k}{(1-\rho)} \frac{\beta_i^k}{\delta} \equiv C. \tag{1.33}$$

This implies that $\Pi(A) < \infty$ for all $A \in \mathcal{A}^*$.

(ii) Second, I show that the feasible region is non-empty. The feasible region is defined by the set $\mathcal{P} = \{x \in \mathbb{R}^n : Hx \leq b\}$. For any $A \in \mathcal{A}^*$, there exists a matching rule under which the constraints are not inconsistent for a range of parameters. This follows from the fact that $\mathcal{A}^* \subset \mathcal{A}$ and the definition of $\mathcal{A}$ implies that $A \in \mathcal{A}$ if and only if there exists an exogenous matching rule for which an equilibrium with mutual acceptance matrix A exists. By Lemma 2 there exists at least one equilibrium that can be implemented by a matching mechanism, hence, $\mathcal{A}^*$ is non-empty. Therefore, the feasible region is non-empty for a range of parameters for each linear program for fix $A \in \mathcal{A}^*$.

Then, by strong duality (Dantzig, 1963), it follows that the linear program attains an optimal solution for any $A \in \mathcal{A}^*$. The optimal value to the linear program, $\Pi(A)$, is finite and $\mathcal{G}$ is non-empty.

(b) $\mathcal{G}$ is finite.

As $\mathcal{G} = \bigcup_{A \in \mathcal{A}^*} \Pi(A)$ and $\mathcal{A}^*$ is finite by construction, $\mathcal{G}$ is also finite as the profit level of a given linear program is a singleton. As each linear program for fix $A \in \mathcal{A}^*$ is bounded, the profit level takes on either a (finite) optimal value if an optimal solution exists or the value is undefined if the linear program is infeasible for given parameters. □

I prove Theorem 1 through a series of lemma. Recall that $s_i^k \in [0, \overline{u}] \equiv \mathcal{S}$ and denote the vector of search costs by $(s_1^k, ..., s_N^k)^{k=A,B} \equiv \mathbf{s}$.

Lemma 7. *Let the vector of search costs* $\mathbf{s}$ *be given. There exists an optimal solution with* $\Pi^* \equiv \max_{\mathbf{s}} \mathcal{G}(\mathbf{s})$.

Proof. By Lemma 5, the set $\mathcal{G}$ is finite and non-empty for any given $s_i^k \in \mathbb{R}_+$. Hence for given vector $\mathbf{s}$, $\mathcal{G}$ has a maximum element and $\Pi^* = \max \mathcal{G}$ is well-defined and has a finite value. $\square$

Now, let the platform choose the vector of search costs $\mathbf{s}$. I prove that there exists an optimal solution $\Pi^{*,s} \equiv \max_{\mathbf{s}} \Pi^*(\mathbf{s})$.

First, observe that if $s_i^k \geq \max_{\theta_j^{-k}}\{\theta_i^k \cdot \theta_j^{-k} - \omega_i^k\}$ for all $\theta_i^k \in \Theta^k$, no agent participates and the equilibrium profit is zero. Therefore, to make positive profits $s_i^k \leq \max_{\theta_j^{-k}}\{\theta_i^k \cdot \theta_j^{-k} - \omega_i^k\}$ for at least one $\theta_i^k \in \Theta^k$ such that the set of participating types $\hat{\Theta}^k$ is non-empty. Recall that $\mathcal{G}(\mathbf{s})$ is the set of profit levels induced through all linear programs that have a feasible solution for given $\mathbf{s}$. In slight abuse of notation, define $\mathcal{G}(\mathbf{s})$ as a correspondence from $\mathbf{s}$ to such profit levels $\Pi(\mathbf{s})$

$$\mathcal{G}(\mathbf{s}) : \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|} \rightrightarrows \mathbb{R}_0^+.$$

which assigns to each point $\mathbf{s}$ of $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$ a finite subset $\mathcal{G}(\mathbf{s})$ of $\mathbb{R}_0^+$. The correspondence is compact-valued as $\mathcal{G}(\mathbf{s})$ is a compact (finite) subset of $\mathcal{C}$ for all $\mathbf{s} \in \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}}$. In the following, I will show that the correspondence is upper hemicontinuous in $\mathbf{s}$ on $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}}$.

To do so, recall the matrix notation of the linear program in Appendix A.1:

$$\begin{aligned} &\max_{x \in X} xc^T \equiv \Pi_A(\mathbf{s}), \\ &s.t. H_A x \leq b_A, x \geq 0. \end{aligned}$$

Denote by subscript A, the profit level and constraint set of the linear program for given matrix $A \in \mathcal{A}^*$. In Lemma 5, I have shown that a linear program for a fixed $A \in \mathcal{A}^*$ has a solution for some $\mathbf{s} \in \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}}$. Additionally, whenever the linear program has a solution, it has an optimal solution. The value of the linear program, $\Pi_A(s)$, is thus finite on a set $\mathcal{J}_A \equiv \{\mathbf{s} \in \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|} | -\infty < \Pi_A(\mathbf{s}) < \infty\}$, where $\mathcal{J}_A \subseteq \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}}$. The set is compact.[25]

Lemma 8. *The value of the objective $\Pi_A(s)$ of a linear program for given matrix $A \in \mathcal{A}^*$ is upper hemicontinuous in $\mathbf{s}$ on $\mathcal{J}_A$.*

Proof. Fix $A \in \mathcal{A}^*$, and consider the associated linear program from Lemma 6. For given $A \in \mathcal{A}^*$, $\mathbf{s}$ changes vector c continuously, as each entry, $\nu(s_i^k)$ or 0, is continuous in

[25] The set $\mathcal{J}_A$ contains all $\mathbf{s} \in \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$ for which the value of the linear program is finite. In other words, the linear program must be bounded and feasible for those $\mathbf{s}$. By Lemma 5, the linear program is bounded. The linear program is feasible for some $\mathbf{s}$ if all constraints can be met, i.e. the feasible region $\mathcal{P}$ is non-empty. Suppose for contradiction that $\mathcal{J}_A$ is not compact. Now, take any sequence $\mathbf{s}_n \to \mathbf{s}$, for which the feasible region is non-empty for all $\mathbf{s}_n$. For the limit point $\mathbf{s}$ not to be in set $\mathcal{J}_A$, the feasible region must be empty for $\mathbf{s}$, and hence, at least one inequality must be violated strictly. But then, as the linear constraints are continuous in $\mathbf{s}$, the constraints must also be violated for $\mathbf{s}_n$ close enough to $\mathbf{s}$, a contradiction.

s_i^k. Furthermore $\mathbf{s}$ changes matrix H_A continuously as s_i^k linearly enters as a coefficient in the incentive and participation constraints. The optimal value of the linear program is given by

$$\Pi_A(\mathbf{s}) \equiv \sup_{x \in \mathbb{R}^n} \{c(\mathbf{s})x | H_A(\mathbf{s})x \leq b_A, x \geq 0\},$$

which is finite on $\mathcal{J}_A$. In slight abuse of notation, denote the correspondence from $\mathbf{s}$ to the optimal value of the linear program by

$$\Pi_A(\mathbf{s}) : \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|} \rightrightarrows \mathbb{R}_0^+.$$

Next, consider set of primal feasible solutions of the linear program that defines objective Π, which is given by the correspondence

$$\mathbf{s} \to P_A(\mathbf{s}) \equiv \{x | H_A(\mathbf{s})x \leq b, x \geq 0\}.$$

First, I show that the set of (primal) feasible solutions of the linear program is upper hemicontinuous in $\mathbf{s}$. Consider the following definition: $P_A(\mathbf{s})$ is upper hemicontinuous at $\mathbf{s}$ on $\mathcal{J}_A$ if

$$\mathbf{s} = \lim_{n \to \infty} \mathbf{s}_n, \; x_n \in P_A(\mathbf{s}_n), \text{ and } x = \lim_{n \to \infty} x_n,$$

implies that $x \in P_A(\mathbf{s})$.[26] To see that $P_A(\mathbf{s})$ is upper hemicontinuous, suppose that $\{\mathbf{s}_n\}_n \in \mathcal{J}_A$ and $\mathbf{s} = \lim_{n\to\infty} \mathbf{s}_n$. Let $\{x_n\}_n$ be a sequence such that for all n, $x_n \in P_A(\mathbf{s})$: $H_A(\mathbf{s}_n)x_n \leq b_A$, and $x = \lim_{n\to\infty} = x_n$. Since by the continuity of $H_A(\cdot)$ and independence of b_A in $\mathbf{s}$

$$||H_A(\mathbf{s}_n) - H_A(\mathbf{s})|| \to 0, \; ||x_n - x|| \to 0, \text{ and } ||b_A - b_A|| = 0,$$

it follows that $H_A x \leq b_A$ and $x \geq 0$, which yields $x \in P_A(\mathbf{s})$. This implies that $P_A(\mathbf{s})$ is in fact upper hemicontinuous in $\mathbf{s}$ on $\mathcal{J}_A$.

Next, I show that this implies that $\Pi_A(\mathbf{s}) = c(\mathbf{s})x$ is upper hemicontinuous in $\mathbf{s}$ on $\mathcal{J}_A$. Suppose that $\{\mathbf{s}_n\}_n \in \mathcal{J}_A$ and $\mathbf{s} = \lim_{n\to\infty} \mathbf{s}_n$. Let $\{\Pi_n\}_n$ be a sequence such that for all n, $\Pi_n \in \Pi_A(\mathbf{s})$, and $\Pi = \lim_{n\to\infty} \Pi_n$. Since by the continuity of $c(\cdot)$

$$||c(\mathbf{s}_n) - c(\mathbf{s})|| \to 0,$$

[26]This definition follows Wets (1985). Furthermore, let $||H|| = \sup_{x \in X} ||Hx||$

and the upper hemicontinuity of $P_A(\mathbf{s})$ on $\mathcal{J}_A$

$$||x_n - x|| \to 0$$

it follows that $\Pi \in \Pi_A(\mathbf{s})$. This implies that $\Pi_A(\mathbf{s})$ is in fact upper hemicontinuous in $\mathbf{s}$ on $\mathcal{J}_A$. $\square$

Recall that $\mathcal{G}(\mathbf{s}) = \bigcup_{A \in \mathcal{A}^*} \Pi_A(\mathbf{s})$ is the finite union over the equilibrium profit levels of each linear program.

Lemma 9. $\mathcal{G}(\mathbf{s})$ *is upper hemicontinuous in* $\mathbf{s}$ *on* $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$.

Proof. Recall that for each $\Pi_A(\mathbf{s})$ the value $\Pi_A(\mathbf{s})$ is finite on $\mathcal{J}_A$ and empty on $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|} \setminus \mathcal{J}_A$. I prove the lemma by induction over the equilibria associated with the finite set $\mathcal{A}^*$. Let there be $\overline{K}$ equilibria, which can be implemented by the linear programs and consider the correspondence $\mathcal{G}_K(\mathbf{s}) = \bigcup_{\{A_1,\ldots,A_K\}} \Pi_A(\mathbf{s})$ that includes K out of $\overline{K}$ equilibria. By induction, I will consider $\mathcal{G}_K$ to include increasingly more equilibria.

Base case: Let $\mathcal{G}_1$ be the correspondence that includes only the trivial equilibrium from Lemma 2 with $A_1 \in \mathcal{A}^*$. Note that $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|} = \mathcal{J}_{A_1}$ as the trivial equilibrium is a solution to the corresponding linear program for each $\mathbf{s} \in \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$. Hence, the statement follows from Lemma 8.

Induction step: The induction hypothesis states that $\mathcal{G}_K(\mathbf{s}) = \bigcup_{\{A_1,\ldots,A_K\}} \Pi_A(\mathbf{s})$ is upper hemicontinuous on $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$. Note that by the induction step, K includes the trivial equilibrium. It remains to show that $\mathcal{G}_K(\mathbf{s}) \cup \Pi_{A_{K+1}}(\mathbf{s})$ is upper hemicontinuous in $\mathbf{s}$ on $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$.

Recall that the correspondence $\mathcal{G}_K(\mathbf{s}) \cup \Pi_{A_{K+1}}(\mathbf{s})$ is upper hemicontinuous at $\mathbf{s}_0 \in \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$, if for any open set $V \subseteq \mathbb{R}_0^+$ with $\mathcal{G}_K(\mathbf{s}_0) \cup \Pi_{A_{K+1}}(\mathbf{s}_0) \subseteq V$, there exists an open neighborhood $U(\mathbf{s}_0) \subseteq \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$ such that if $\mathbf{s} \in U(\mathbf{s}_0)$, then $\mathcal{G}_K(\mathbf{s}) \cup \Pi_{A_{K+1}}(\mathbf{s}) \subseteq V$.

Let $\mathbf{s}_0 \in \mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$ and V be an open set with $\mathcal{G}_K(\mathbf{s}_0) \cup \Pi_{A_{K+1}}(\mathbf{s}_0) \subseteq V$. Suppose first that $\Pi_{A_{K+1}}$ is empty at $\mathbf{s}_0$. Since $\mathcal{G}_K(\mathbf{s}_0) \cup \Pi_{A_{K+1}}(\mathbf{s}_0) \subseteq V$, it follows that $\mathcal{G}_K(\mathbf{s}_0) \subseteq V$ and $\Pi_{A_{K+1}}(\mathbf{s}_0) \subseteq V$ by assumption (where V is the union of an open set and the empty set). By the upper hemicontinuity of $\mathcal{G}_K(\mathbf{s})$, there exists a neighborhood U_K of $\mathbf{s}_0$ such that $\mathcal{G}_K(\mathbf{s}_0) \subseteq V$ for all $\mathbf{s} \in U_K$. Additionally, there exists a neighborhood U_{K+1} of $\mathbf{s}_0$ such that $\Pi_{A_{K+1}}(\mathbf{s}_0) = \emptyset \subseteq V$ for all $\mathbf{s} \in U_{K+1}$ (by the compactness of $\mathcal{J}_{A_{K+1}}$. Let $U = U_K \cap U_{K+1}$. Then, for any $\mathbf{s} \in U$, both $\mathcal{G}_K(\mathbf{s}) \subseteq V$ and $\Pi_{A_{K+1}}(\mathbf{s}) \subseteq V$ such that $\mathcal{G}_K(\mathbf{s}) \cup \Pi_{A_{K+1}}(\mathbf{s}) \subseteq V$.

Let both $\mathcal{G}_K(\mathbf{s})$ and $\Pi_{A_{K+1}}(\mathbf{s})$ be non-empty at $\mathbf{s}_0$. Since $\mathcal{G}_K(\mathbf{s}_0) \cup \Pi_{A_{K+1}}(\mathbf{s}_0) \subseteq V$, it follows that $\mathcal{G}_K(\mathbf{s}_0) \subseteq V$ and $\Pi_{A_{K+1}}(\mathbf{s}_0) \subseteq V$. As both $\mathcal{G}_K(\mathbf{s}_0)$ and $\Pi_{A_{K+1}}(\mathbf{s}_0)$ are upper hemicontinuous for $\mathbf{s}_0$, it holds that: There exists a neighborhood U_K of $\mathbf{s}_0$ such that $\mathcal{G}_K(\mathbf{s}_0) \subseteq V$ for all $\mathbf{s} \in U_K$ and U_{K+1} of $\mathbf{s}_0$ such that $\Pi_{A_{K+1}}(\mathbf{s}_0) \subseteq V$ for all $\mathbf{s} \in U_{K+1}$. Then, for any $\mathbf{s} \in U$, both $\mathcal{G}_K(\mathbf{s}) \subseteq V$ and $\Pi_{A_{K+1}}(\mathbf{s}) \subseteq V$ such that $\mathcal{G}_K(\mathbf{s}) \cup \Pi_{A_{K+1}}(\mathbf{s}) \subseteq V$.

Therefore, $\mathcal{G}(\mathbf{s})$ is upper hemicontinuous in $\mathbf{s}$ on $\mathcal{S}^{|\Theta^k| \times |\Theta^{-k}|}$. $\square$

Lemma 10. *The function* $\Pi^*(\mathbf{s})$ *is upper semi-continuous in* $\mathbf{s}$ *on* $\mathcal{S}^{|\Theta^k|\times|\Theta^{-k}|}$.

Proof. The function Π^* is upper-semicontinuous if for every point $\mathbf{s} \in \mathcal{S}^{|\Theta^k|\times|\Theta^{-k}|}$, $\Pi(\mathbf{s}) \geq \limsup \Pi(\mathbf{s}_n)$ for every sequence $\{\mathbf{s}_n\}_n \subset \mathcal{S}^{|\Theta^k|\times|\Theta^{-k}|}$ satisfying $\lim_{n\to\infty} \mathbf{s}_n = \mathbf{s}$.

Let $\lim_{n\to\infty} \mathbf{s}_n = \mathbf{s}$, and define $\Pi^*_n = \max \mathcal{G}(\mathbf{s}_n)$, so that $\Pi^*_n \in \mathcal{G}(\mathbf{s}_n)$ for all n. Since for each $\mathbf{s}_n$ $\mathcal{G}(\cdot)$ is finite by Lemma 5 and the sequence $\{\Pi^*_n\}$ is bounded, it has a convergent subsequence by the Bolzano-Weierstrass theorem: $\Pi^*_{n_k} \to \Pi'$ for some $\Pi' \in \mathbb{R}^+_0$. Then, as $\Pi^*_{n_k} \in \mathcal{G}(\mathbf{s}_{n_k})$, $\mathbf{s}_{n_k} \to \mathbf{s}$, and $\Pi^*_{n_k} \to \Pi'$, the upper hemiconituity of $\mathcal{G}(\mathbf{s})$ implies that any limit point of $\Pi^*_{n_k}$ belongs to $\mathcal{G}(\mathbf{s})$, i.e. $\Pi' \in \mathcal{G}(\mathbf{s})$. Therefore, $\Pi' \leq \max G(\mathbf{s})$. Since $\Pi^*_{n_k} \to \Pi'$, this implies:

$$\lim_{n\to\infty} \sup \Pi_n = \lim_{n\to\infty} \sup \max \mathcal{G}(\mathbf{s}_n) \leq \max \mathcal{G}(\mathbf{s}).$$

□

Intuitively, Π^* is continuous in $\mathbf{s}$ except for jump points (discontinuities). At a jump point, the definition of upper semicontinuity requires that the function is only allowed to jump "up". By the upper hemicontinuity, I already know that the limit point — when taking a sequence of $\mathbf{s}$ — is still in $\max \mathcal{G}(\mathbf{s})$. Additionally, due to the definition of Π^* as $\Pi^* = \max \mathcal{G}(\mathbf{s})$, the limit point can only jump "up".

Proof of Theorem 1 By Lemma 10, $\max \mathcal{G} = \Pi^*(\mathbf{s})$ is upper semi-continuous in $\mathbf{s}$ and compact-valued. Thus, there exists a maximum by Weierstrass extreme value theorem on the compact set $\mathcal{S}^{|\Theta^k|\times|\Theta^{-k}|}$. □

Proof of Proposition 1 The proof proceeds by considering the cases where search costs are exogenous and where search costs are chosen as search fee or advertising.

Case 1: Exogenous Search Cost First, suppose search costs are exogenously given. Let the parameters be drawn uniformly from the following sets: $\theta^k_i \in \Theta^k = [\underline{\theta}, \overline{\theta}] \subseteq \mathbb{R}_+$, $\beta^k_i \in [0, \overline{\beta}]$, $\delta \in (0,1]$, $\omega^k_i \in \Omega = [0, \overline{\omega}]$, and $s^k_i \in [0, \overline{u}]$. An outcome is said to be generically suboptimal if the set of parameter values for which it is optimal has measure zero in the relevant parameter space.

For given $A \in \mathcal{A}^*$, an optimal solution is a matching rule for which the objective function of the linear program in Appendix A.1 attains its maximum value. Recall from Appendix A.1 that the platform solves

$$\max_{\{\Phi(\cdot), m(\cdot)\}^k_{ij}} \sum_{k=A,B} \sum_{\theta^k_i \in \Theta^k} \frac{(1-\delta)\nu(s^k_i)}{1-\rho} f(\theta^k_i),$$

subject to participation constraints (Equation 1.12), incentive constraints and 1.13), feasibility constraints (Equation 1.14) and steady-state constraints (Equation 1.15). Using the steady state conditions from Equation 1.15 to substitute for $f(\theta_i^k)$ yields

$$\max_{\{\Phi(\cdot),m(\cdot)\}_{ij}^k} \sum_{k=A,B} \sum_{\theta_i^k \in \Theta^k} \frac{(1-\delta)\nu(s_i^k)}{1-\rho} \frac{\beta_i^k - (1-\delta)\sum_j m(\theta_i^k, \theta_j^{-k})}{\delta}.$$

Both feasibility and steady-state constraints must be binding in the optimal solution. Additionally, at least one participation or incentive constraint must be binding in the optimal solution. Suppose otherwise, then the platform can decrease at least one $m(\theta_i^k, \theta_j^{-k})$ such that one constraint is binding and thereby increase its profits.

Recall that $\{m^{RM}(\theta_i^k, \theta_j^{-k})\}_{ij}^k$ is the vector of masses of matched pairs under random matching. Then, $m^{RM}(\theta_i^k, \theta_j^{-k}) = 0$ if $\alpha(\theta_i^k, \theta_j^{-k}) = 0$ and

$$m^{RM}(\theta_i^k, \theta_j^{-k}) = \frac{\alpha(\theta_i^k, \theta_j^{-k})\beta_i^k \mu(\theta_i^k, \omega_i^k)\beta_j^{-k}\mu(\theta_j^{-k}, \omega_j^{-k})}{\left(\sum_{\theta_i^k} \beta_i^k \mu(\theta_i^k, \omega_i^k)\right) \cdot \left(\sum_{\theta_j^{-k}} \beta_j^{-k} \mu(\theta_j^{-k}, \omega_j^{-k})\right)}$$

if $\alpha(\theta_i^k, \theta_j^{-k}) \in (0, 1]$ (see Appendix A.2).This is a function of the inflow vector $(\beta_1^k, ..., \beta_{N^k}^k)_k$, δ and the probability of a type θ_i^k being matched to their outside option ω_i^k $(\mu(\theta_i^k, \omega_i^k))$. Observe that for given $A \in \mathcal{A}^*$, $m^{RM}(\theta_i^k, \theta_j^{-k})$ is independent of s_i^k. Fix $A \in \mathcal{A}^*$. Given $\{m^{RM}(\theta_i^k, \theta_j^{-k})\}_{ij}^k$, I show that the participation and incentive constraints are generically non-binding. Rearranging and using the steady state condition yields the following constraints

$$\beta_i^k \left(\theta_i^k \theta_j^k - \omega_i^k + \frac{(1-\delta)}{\delta} s_i^k\right) \leq (1-\delta) \sum_j m^{RM}(\theta_i^k, \theta_j^{-k}) \left(\theta_i^k \theta_j^{-k} - \omega_i^k + \frac{(1-\delta)}{\delta} s_i^k\right), \tag{1.34}$$

$$\beta_i^k \frac{(1-\delta)}{\delta} s_i^k \leq (1-\delta) \sum_j m^{RM}(\theta_i^k, \theta_j^{-k}) \left(\theta_i^k \theta_j^{-k} - \omega_i^k + \frac{(1-\delta)}{\delta} s_i^k\right). \tag{1.35}$$

Suppose $\beta_i^k, \omega_i^k, \theta_i^k, \theta_j^{-k}$ and δ are drawn uniformly from their continuous intervals. Note that each constraint for a type θ_i^k is a linear equation in s_i^k. Hence, for given $\{m^{RM}(\theta_i^k, \theta_j^{-k})\}_{ij}^k$, there exists at most one s_i^k per participation or incentive constraint of type θ_i^k such that the constraint is binding. This implies that if s_i^k is drawn uniformly from a continuous interval, the set of parameters for which the constraint is binding has measure zero. Therefore integrating over the cases for which at least one constraint is binding, the corresponding set of parameters has measure zero as well. Hence, for each $A \in \mathcal{A}^*$, the constraints are generically non-binding. Lastly, since $\mathcal{A}^*$ is finite, this concludes the proof that random matching is generically suboptimal for exogenously given search costs. Next, consider the case where the platform chooses the vector of search costs.

Case 2: Endogenous Search Cost (Search Fee) Consider the case, in which the platform sets a linear search fee and earns $\nu(s_i^k) = s_i^k$ for all $\theta_i^k = \Theta^k, k = A, B$. The platform maximizes the total match surplus and fully extracts the surplus through the search fee. The platform solves the maximization problem in Equation 1.24 subject to Equation 1.25, 1.26, and 1.27 from Appendix A.2. Given the solution to this problem, $\{m^{PAM}(\theta_i^k, \theta_j^{-k})\}_{ij}^k$, the platform sets s_i^k to fully extract each type's surplus:

$$\beta_i^k \frac{(1-\delta)}{\delta} s_i^k = (1-\delta) \sum_j m^{PAM}(\theta_i^k, \theta_j^{-k}) \left(\theta_i^k \theta_j^{-k} - \omega_i^k + \frac{(1-\delta)}{\delta} s_i^k \right).$$

The optimal matching rule that maximizes total match surplus follows the procedure: Starting with the highest possible type on side A, each agent is matched to the highest possible type on side B. If there are not enough high types remaining on side B, the algorithm proceeds in descending order of type on side B until all agents of the highest possible type on side A are matched. The process continues in descending order with the next highest type on side A, each time matching to the next available remaining types on side B. Once all agents on B have been matched, any remaining agents on side A are assigned to their outside option. The optimal matching rule hence always differs from random matching. To see this, observe that higher types receive better recommendations under the above procedure, whereas each type receive the same recommendations under random matching.

Case 3: Endogenous Search Cost (Advertising) Now consider the case in which the platform earns a revenue of $\nu(s_i^k)$ when charging search costs s_i^k. I, again, examine random matching that satisfies the feasibility constraints outlined in Equation 1.14. For any $A \in \mathcal{A}^*$, and using the steady-state conditions to substitute for $f^{RM}(\theta_i^k)$, the platform's objective under random matching becomes the following maximization problem:

$$\max_{\mathbf{s}} \sum_{k=A,B} \sum_{\theta_i^k \in \Theta^k} \frac{(1-\delta)\nu(s_i^k)}{1-\rho} \underbrace{\frac{\beta_i^k - (1-\delta)\sum_j m^{RM}(\theta_i^k, \theta_j^{-k})}{\delta}}_{=f^{RM}(\theta_i^k)}. \tag{1.36}$$

subject to the participation and incentive constraints in Equation 1.34 and 1.35.

To maximize $\mathbf{s}$, observe first that $\nu'(s_i^k) > 0$ as $\nu(s_i^k)$ is strictly increasing in s_i^k. This implies that the platform has an incentive to increase the search costs as much as possible given the constraints. Therefore for $A \in \mathcal{A}^*$, the optimal solution is to choose s_i^k such that for each type $\theta_i^k \in \Theta^k, k = A, B$ either the participation or the relevant incentive constraint — induced by A — is binding. Note that the random matching vector satisfies the feasibility condition, and as random matching is independent of s_i^k feasibility remains to be satisfied.

Next, I show that the platform has an incentive to deviate from the above solution. First, suppose $A' \in \mathcal{A}^*$ consists only of entries equal to one, so all agents accept any match in the first period. Incentive constraints are slack, and the platform chooses $\mathbf{s}$ to make participation constraints binding. Under random matching, the platform can at most charge the expected value of a match. By deviating to positive assortative matching—the solution under a linear search fee—the platform can raise search costs and profits, since $\nu(s_i^k)$ is strictly increasing in s_i^k.

Second, consider any matrix in $A'' \in \mathcal{A}^* \setminus \{A'\}$. In this case, at least one type rejects another type with positive probability. As the match utility with the lowest type is the lowest, this implies that at least one type is willing to reject the lowest type. Consider the pair of types $(\theta_1^k, \theta_R^{-k})$ for which type $\theta_R^{-k} \in \Theta^{-k}$ is willing to reject the lowest type θ_1^k on the other market side ($\alpha(\theta_1^k, \theta_R^{-k}) = 0$). Recall that each type must be accepted by at least one other type on the opposite market side to be willing to participate, thus consider pairs $(\theta_1^k, \theta_A^{-k})$ and $(\theta_A^k, \theta_R^{-k})$ for which $\alpha(\theta_1^k, \theta_A^{-k}) = 1$ and $\alpha(\theta_A^k, \theta_R^{-k}) = 1$.

For fix $A'' \in \mathcal{A}^* \setminus \{A'\}$, I will show that the platform's profit can be improved by changing the matching rules for types $\theta_1^k, \theta_A^k, \theta_A^{-k}$ and θ_R^{-k} as well as adjusting their search costs. The platform will choose the mass of recommended pairs

$$\Phi'(\theta_1^k, \theta_R^{-k}), \Phi'(\theta_1^k, \theta_A^{-k}), \Phi'(\theta_A^k, \theta_R^{-k}), \Phi'(\theta_A^k, \theta_A^{-k}),$$

and the mass of matched pairs $m'(\cdot,\cdot) = \alpha(\cdot,\cdot)\Phi'(\cdot,\cdot)$ as detailed below. For all other types, the platform chooses the mass of recommended pairs such that

$$\begin{aligned}
\Phi'(\theta_i^k, \theta_j^{-k}) &= \Phi^{RM}(\theta_i^k, \theta_j^{-k}), \ \forall \theta_i^k \in \Theta^k \setminus \{\theta_1^k, \theta_A^k\}, \\
\Phi'(\theta_i^k, \theta_j^{-k}) &= \Phi^{RM}(\theta_i^k, \theta_j^{-k}), \ \forall \theta_j^{-k} \in \Theta^{-k} \setminus \{\theta_R^{-k}, \theta_A^{-k}\}.
\end{aligned}$$

Without loss of generality, suppose that the total mass of all types on market side A is smaller or equal than the total mass of all types on market side B: $\sum_{\theta_i^A \in \Theta^A} f^{RM}(\theta_i^A) \leq \sum_{\theta_i^B \in \Theta^B} f^{RM}(\theta_i^B)$. Then, for market side A, the platform chooses the mass of types that are recommended to their outside option such that

$$\Phi'(\theta_i^k, \omega_i^k) = \Phi^{RM}(\theta_i^k, \omega_i^k), \forall \theta_i^k \in \Theta^k \setminus \{\theta_1^k, \theta_A^k\}, k = A.$$

The mass of recommended pairs and the mutual acceptance probabilities remain the same as under random matching. As a result, the mass of matched pairs—defined as the product of these two terms—is also unchanged. Therefore, the participation and incentive constraints for all other types continue to hold. In addition, the feasibility constraint in Equation (1.14) and the steady-state constraint in Equation (1.15) are still satisfied.

For $\varepsilon \in [-\min\{\beta_i - m^{RM}(\cdot,\cdot)\}, \min\{m^{RM}(\cdot,\cdot)\}]$, the platform chooses $m^{RM}(\theta_1^k, \theta_A^{-k}) - m'(\theta_1^k, \theta_A^{-k}) = \varepsilon$ and $m^{RM}(\theta_A^k, \theta_R^{-k}) - m'(\theta_A^k, \theta_R^{-k}) = \varepsilon$, i.e. the platform changes the mass of the two matched pairs by ε. Substituting the change into the steady state condition (Equation 1.15) for type θ_1^k and type θ_R^{-k} yields

$$
\begin{aligned}
f^{RM}(\theta_1^k) + \frac{1-\delta}{\delta}\varepsilon =& \frac{1}{\delta}\left(\beta_1^k - (1-\delta)\left(-\varepsilon + \sum_{\Theta^{-k}} m^{RM}(\theta_1^k, \theta_j^{-k})\right)\right), \\
f^{RM}(\theta_R^{-k}) + \frac{1-\delta}{\delta}\varepsilon =& \frac{1}{\delta}\left(\beta_R^{-k} - (1-\delta)\left(-\varepsilon + \sum_{\Theta^{k}} m^{RM}(\theta_i^k, \theta_R^{-k})\right)\right).
\end{aligned}
$$

Therefore, by decreasing (increasing) the mass of the two matched pairs, increases (decreases) the steady state mass by $\frac{1-\delta}{\delta}\varepsilon$ compared to the steady state mass under random matching. Substituting $\Phi^{RM}(\theta_1^k, \theta_A^{-k}) - \Phi'(\theta_1^k, \theta_A^{-k}) = \varepsilon$ and $\Phi^{RM}(\theta_A^k, \theta_R^{-k}) - \Phi'(\theta_A^k, \theta_R^{-k}) = \varepsilon$ into the feasibility constraints of type θ_1^k and type θ_R^{-k} yields

$$
f^{RM}(\theta_1^k) + \frac{1-\delta}{\delta}\varepsilon = \Phi'(\theta_1^k, \theta_R^{-k}) - \varepsilon + \mathbf{1}_{k=A}\Phi^{RM}(\theta_1^k, \omega_1^k) + \sum_{\Theta^{-k}\setminus\{\theta_R^{-k}\}} \Phi^{RM}(\theta_1^k, \theta_j^{-k}), \tag{1.37}
$$

$$
f^{RM}(\theta_R^{-k}) + \frac{1-\delta}{\delta}\varepsilon = \Phi'(\theta_1^k, \theta_R^{-k}) - \varepsilon + \mathbf{1}_{k=A}\Phi^{RM}(\theta_R^{-k}, \omega_R^{-k}) + \sum_{\Theta^{k}\setminus\{\theta_1^{k}\}} \Phi^{RM}(\theta_i^k, \theta_R^{-k}), \tag{1.38}
$$

which implies that $\Phi'(\theta_1^k, \theta_R^{-k}) - \Phi^{RM}(\theta_1^k, \theta_R^{-k}) = \frac{1-\delta}{\delta}\varepsilon + \varepsilon = \varepsilon/\delta$. It remains to determine $\Phi'(\theta_A^k, \theta_A^{-k}) - \Phi^{RM}(\theta_A^k, \theta_A^{-k})$ and $m'(\theta_A^k, \theta_A^{-k}) - m^{RM}(\theta_A^k, \theta_A^{-k})$. To do so, consider two cases: either $\alpha(\theta_A^k, \theta_A^{-k}) = 0$ or $\alpha(\theta_A^k, \theta_A^{-k}) = 1$.

In the first case, $\alpha(\theta_A^k, \theta_A^{-k}) = 0$, I can exchange θ_A^k for θ_1^k and θ_A^{-k} for θ_R^{-k} in Equation 1.37 and 1.38 above. Then, it follows that $\Phi'(\theta_A^k, \theta_A^{-k}) - \Phi^{RM}(\theta_A^k, \theta_A^{-k}) = \varepsilon/\delta$ and $m'(\theta_A^k, \theta_A^{-k}) = 0$.

In the second case, $\alpha(\theta_A^k, \theta_A^{-k}) = 1$, the platform can set $\Phi'(\theta_A^k, \theta_A^{-k}) - \Phi^{RM}(\theta_A^k, \theta_A^{-k}) = m'(\theta_A^k, \theta_A^{-k}) - m^{RM}(\theta_A^k, \theta_A^{-k}) = \varepsilon$. Since the platform decreases (increases) the mass of the matched pair $(\theta_A^k, \theta_R^{-k})$ but increases (decreases) the mass of the matched pair $(\theta_A^k, \theta_A^{-k})$ by the same amount, this implies that the steady state mass of type θ_A^k is unchanged compared to the steady state masses under random matching. Additionally, feasibility continues to be satisfied as the platform shift mass ε from one recommended pair to the other. Similarly, the steady state mass of type θ_A^{-k} is the same as under random matching and the steady state constraint as well as feasibility constraint remain satisfied.

Next determine the change in search costs for types $\theta_1^k, \theta_A^k, \theta_A^{-k}$ and θ_R^{-k}. Note that for the newly chosen mass of recommended and matched pairs $(\Phi'(\cdot,\cdot), m'(\cdot,\cdot))$, the originally binding participation or incentive constraint is no longer binding. Since, however, the right-hand side of the participation or incentive constraints (see Equation 1.34 and 1.35) are linearly increasing in $m(\cdot,\cdot)$ and the left-hand side of the constraints are ordered due to the supermodularity of the match utility, the platform can choose a new search cost $\tilde{s}_i^k$ such that the constraint becomes binding again. Let the platform choose $\tilde{s}_1^k, \tilde{s}_A^k, \tilde{s}_R^{-k}, \tilde{s}_A^{-k}$ such that originally binding participation incentive constraint of each type is binding again.

Using Equations 1.34 and 1.35 and $m^{RM}(\theta_1^k, \theta_A^{-k}) - m'(\theta_i^k, \theta_i^{-k}) = \varepsilon$ and $m^{RM}(\theta_A^k, \theta_R^{-k}) - m'(\theta_A^k, \theta_R^{-k}) = \varepsilon$, the difference between $s_1^k - \tilde{s}_1^k$ and $s_R^{-k} - \tilde{s}_A^{-k}$ is given by

$$(1-\delta)\underbrace{\frac{\beta_1^k - (1-\delta)\sum_j m^{RM}(\theta_1^k, \theta_j^{-k})}{\delta}}_{f(\theta_1^k)}(s_1^k - \tilde{s}_1^k) = \varepsilon(\theta_1^k\theta_A^{-k} - \omega_1^k + \frac{1-\delta}{\delta}\tilde{s}_1^k). \tag{1.39}$$

Observe that the right-hand side is positive for $\varepsilon > 0$ as $\theta_1^k\theta_A^{-k} - \omega_1^k > 0$ due to the fact that both types mutually accept each other. Then, it follows that $\tilde{s}_1^k$ must be smaller than s_1^k for $\varepsilon > 0$. Additionally, by the steady state constraint, the factor on the left-hand side is equal to $(1-\delta)f^{RM}(\theta_1^k)$. Similarly, using $m^{RM}(\theta_A^k, \theta_R^{-k}) - m'(\theta_A^k, \theta_R^{-k}) = \varepsilon$ and taking the difference, $s_R^{-k} - \tilde{s}_A^{-k}$ is given by

$$(1-\delta)f^{RM}(\theta_R^{-k})(s_R^{-k} - \tilde{s}_R^{-k}) = \varepsilon(\theta_A^k\theta_R^{-k} - \omega_R^{-k} + \frac{1-\delta}{\delta}\tilde{s}_R^{-k}). \tag{1.40}$$

Observe that the right-hand side is again positive, so that $s_R^{-k} > \tilde{s}_A^{-k}$ for $\varepsilon > 0$.

Next, if $\alpha(\theta_A^k, \theta_A^{-k}) = 0$, the difference between the search costs for types θ_A^k and θ_A^{-k} can be derived as above

$$(1-\delta)f^{RM}(\theta_A^k)(s_A^k - \tilde{s}_A^k) = \varepsilon(\theta_A^k\theta_R^{-k} - \omega_A^k + \frac{1-\delta}{\delta}\tilde{s}_A^k), \tag{1.41}$$

$$(1-\delta)f^{RM}(\theta_A^{-k})(s_A^{-k} - \tilde{s}_A^{-k}) = \varepsilon(\theta_1^k\theta_A^{-k} - \omega_A^{-k} + \frac{1-\delta}{\delta}\tilde{s}_A^{-k}). \tag{1.42}$$

Again, it holds that $s_A^k > \tilde{s}_A^k$ and $s_A^{-k} > \tilde{s}_A^{-k}$ for $\varepsilon > 0$.

If $\alpha(\theta_A^k, \theta_A^{-k}) = 1$, recall that the platform sets: $m^{RM}(\theta_1^k, \theta_A^{-k}) - m'(\theta_i^k, \theta_i^{-k}) = \varepsilon$ $m^{RM}(\theta_A^k, \theta_R^{-k}) - m'(\theta_A^k, \theta_R^{-k}) = \varepsilon$ and $m^{RM}(\theta_A^k, \theta_A^{-k}) - m'(\theta_A^k, \theta_A^{-k}) = -\varepsilon$. Using Equations 1.34 and 1.35, the difference of $s_A^{-k} - \tilde{s}_A^{-k}$ is given by

$$(1-\delta)f^{RM}(\theta_A^{-k})(s_A^{-k} - \tilde{s}_A^{-k}) = \varepsilon(\theta_1^k\theta_A^{-k} - \omega_A^{-k} + \frac{1-\delta}{\delta}\tilde{s}_A^{-k}) - \varepsilon(\theta_A^k\theta_A^{-k} - \omega_A^{-k} + \frac{1-\delta}{\delta}\tilde{s}_A^{-k}). \tag{1.43}$$

Since $\theta_A^k > \theta_1^k$, the right-hand side is negative, so that $s_A^{-k} < \tilde{s}_A^{-k}$ for $\varepsilon > 0$. Similarly, the difference of $s_A^k - \tilde{s}_A^k$ is given by

$$(1-\delta) f^{RM}(\theta_A^k)(s_A^k - \tilde{s}_A^k) = \varepsilon(\theta_A^k \theta_R^{-k} - \theta_A^k \theta_A^{-k}), \tag{1.44}$$

where the right-hand side is non-negative if $\theta_R^{-k} \geq \theta_A^{-k}$.

To determine whether the deviation is profitable, consider the difference in profits between the deviation profits and random matching profits (Equation 1.36). In the first case, when $\alpha(\theta_A^k, \theta_A^{-k}) = 0$, the steady state mass of all four types $(\theta_1^k, \theta_A^k, \theta_R^{-k}, \theta_A^{-k})$ increases (decreases) by ${}^{(1-\delta)\varepsilon}/_{\delta}$ while their search costs decrease (increase). The difference in profits is therefore

$$\sum_{\theta_i^k \in \left\{\theta_1^k, \theta_A^k, \theta_R^{-k}, \theta_A^{-k}\right\}} \left[(\nu(\tilde{s}_i^k) - \nu(s_i^k)) f^{RM}(\theta_i^k) + \nu(\tilde{s}_i^k) \frac{(1-\delta)\varepsilon}{\delta} \right]. \tag{1.45}$$

Differentiating Equation 1.45 with respect to ε_i^k, and evaluating the condition at $\varepsilon = 0$, yields:

$$\sum_{\theta_i^k \in \left\{\theta_1^k, \theta_A^k, \theta_R^{-k}, \theta_A^{-k}\right\}} \nu'(\tilde{s}_i^k)|_{\tilde{s}_i^k = s_i^k} \frac{\partial \tilde{s}_i^k}{\partial \varepsilon}|_{\varepsilon=0} f^{RM}(\theta_i^k) + \frac{1-\delta}{\delta} \nu(\tilde{s}_i^k)|_{\tilde{s}_i^k = s_i^k}. \tag{1.46}$$

To determine the partial derivative of the search costs with respect to ε, I totally differentiate Equations 1.39, 1.40, 1.41, and 1.42. Evaluating the derivative at $\varepsilon = 0$ yields:

$$\begin{aligned}
\frac{\partial \tilde{s}_1^k}{\partial \varepsilon}|_{\varepsilon=0} &= -\frac{\theta_1^k \theta_A^{-k} - \omega_1^k + \frac{1-\delta}{\delta} s_1^k}{(1-\delta) f^{RM}(\theta_1^k)}, \\
\frac{\partial \tilde{s}_A^k}{\partial \varepsilon}|_{\varepsilon=0} &= -\frac{\theta_A^k \theta_R^{-k} - \omega_A^k + \frac{1-\delta}{\delta} s_A^k}{(1-\delta) f^{RM}(\theta_A^k)}, \\
\frac{\partial \tilde{s}_R^{-k}}{\partial \varepsilon}|_{\varepsilon=0} &= -\frac{\theta_A^k \theta_R^{-k} - \omega_R^{-k} + \frac{1-\delta}{\delta} s_R^{-k}}{(1-\delta) f^{RM}(\theta_R^{-k})}, \\
\frac{\partial \tilde{s}_A^{-k}}{\partial \varepsilon}|_{\varepsilon=0} &= -\frac{\theta_1^k \theta_A^{-k} - \omega_A^{-k} + \frac{1-\delta}{\delta} s_A^{-k}}{(1-\delta) f^{RM}(\theta_A^{-k})}.
\end{aligned}$$

For analytical convenience, consider the class of concave functions $\nu(s_i^k) = \kappa(s_i^k)^\alpha$ for $\kappa \in \mathbb{R}^+$ and $\alpha \in (0,1)$ from now on. Substituting $\nu(s_i^k) = \kappa(s_i^k)^\alpha$, $\nu'(s_i^k) = \kappa\alpha(s_i^k)^{\alpha-1}$, and the partial derivatives above into Equation 1.46 yields

$$
\begin{aligned}
D \equiv\; & \underbrace{\alpha(s_1^k)^{\alpha-1}\left(-\frac{\theta_1^k\theta_A^{-k} - \omega_1^k + \frac{1-\delta}{\delta}s_1^k}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_1^k)^\alpha}_{=d(\theta_1^k)} \\
& + \underbrace{\alpha(s_A^k)^{\alpha-1}\left(-\frac{\theta_A^k\theta_R^{-k} - \omega_A^k + \frac{1-\delta}{\delta}s_A^k}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_A^k)^\alpha}_{d(\theta_A^k)} \\
& + \underbrace{\alpha(s_A^{-k})^{\alpha-1}\left(-\frac{\theta_1^k\theta_A^{-k} - \omega_A^{-k} + \frac{1-\delta}{\delta}s_A^{-k}}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_A^{-k})^\alpha}_{=d(\theta_A^{-k})} \\
& + \underbrace{\alpha(s_R^{-k})^{\alpha-1}\left(-\frac{\theta_A^k\theta_R^{-k} - \omega_R^{-k} + \frac{1-\delta}{\delta}s_R^{-k}}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_R^{-k})^\alpha}_{=d(\theta_R^{-k})}.
\end{aligned}
$$

For the deviation to be profitable, the expression must be non-zero when being evaluated at $\varepsilon = 0$. First, observe that D is continuous in α and $D > 0$ if $\alpha = 0$. Second, I will argue that D has at most one root in α. To do so, examine the terms for θ_i^k:

$$
d(\theta_1^k) = \alpha(s_1^k)^{\alpha-1}\left(-\frac{\theta_1^k\theta_A^{-k} - \omega_1^k + \frac{1-\delta}{\delta}s_1^k}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_1^k)^\alpha.
$$

Differentiating with respect to α results in

$$
\begin{aligned}
\frac{\partial d(\theta_i^k)}{\partial\alpha} &= ((s_1^k)^{\alpha-1} + \alpha(s_1^k)^{\alpha-1}\ln s_1^k)\left(-\frac{\theta_1^k\theta_A^{-k} - \omega_1^k + \frac{1-\delta}{\delta}s_1^k}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_1^k)^\alpha \ln s_1^k, \\
&= (s_1^k)^{\alpha-1}\left(-\frac{\theta_1^k\theta_A^{-k} - \omega_1^k + \frac{1-\delta}{\delta}s_1^k}{(1-\delta)}(1+\alpha\ln(s_1^k)) + \frac{1-\delta}{\delta}s_i^k\ln(s_i^k)\right).
\end{aligned}
$$

Now, observe that $(s_1^k)^{\alpha-1}$ is strictly increasing in α, whereas the expression in brackets changes sign at most once since it is linear in α. This implies that $\frac{\partial d(\theta_i^k)}{\partial\alpha}$ changes sign at most once, in which case it is positive for some $\alpha < \alpha'$ and negative for $\alpha > \alpha'$. Similarly, this holds for the equivalent expressions, $d(\cdot)$, for each type $\theta_A^k, \theta_A^{-k}, \theta_R^{-k}$. Then, since the function D is continuous in α, $D > 0$ for $\alpha = 0$, and D is increasing in α for $\alpha < \alpha''$ and decreasing for $\alpha > \alpha''$, it follows that D has at most one root.

In the second case, when $\alpha(\theta_A^k, \theta_A^{-k}) = 1$, the steady state mass of types $(\theta_1^k, \theta_R^{-k})$ increases (decreases) by $\frac{(1-\delta)\varepsilon}{\delta}$ while the steady state mass of types θ_A^k and θ_A^{-k} remains unchanged. The difference in profits is therefore

$$\sum_{\theta_i^k \in \left\{\theta_1^k, \theta_R^{-k}\right\}} \left[(\nu(\tilde{s}_i^k) - \nu(s_i^k)) f^{RM}(\theta_i^k) + \nu(\tilde{s}_i^k)\frac{1-\delta}{\delta}\varepsilon\right] + \sum_{\theta_i^k \in \{\theta_A^k, \theta_A^{-k}\}} \left[(\nu(\tilde{s}_i^k) - \nu(s_i^k)) f^{RM}(\theta_i^k)\right]. \tag{1.47}$$

Differentiating Equation 1.47 with respect to ε and evaluating the condition at $\varepsilon = 0$, yields

$$\sum_{\theta_i^k \in \left\{\theta_1^k, \theta_R^{-k}\right\}} \left[\nu'(\tilde{s}_i^k)|_{\tilde{s}_i^k = s_i^k} \frac{\partial \tilde{s}_i^k}{\partial \varepsilon}|_{\varepsilon=0} f^{RM}(\theta_i^k) + \frac{1-\delta}{\delta}\nu(\tilde{s}_i^k)|_{\tilde{s}_i^k = s_i^k}\right] \tag{1.48}$$

$$+ \sum_{\theta_i^k \in \{\theta_A^k, \theta_A^{-k}\}} \left[\nu'(\tilde{s}_i^k)|_{\tilde{s}_i^k = s_i^k} \frac{\partial \tilde{s}_i^k}{\partial \varepsilon}|_{\varepsilon=0} f^{RM}(\theta_i^k)\right]. \tag{1.49}$$

Again, the expression must be non-zero for the deviation to be profitable. To determine the partial derivatives of $\tilde{s}_A^k$ and $\tilde{s}_A^{-k}$ with respect to ε, I totally differentiate Equations 1.43 and 1.44:

$$\frac{\partial \tilde{s}_A^k}{\partial \varepsilon}|_{\varepsilon=0} = \frac{\theta_A^k(\theta_A^{-k} - \theta_R^{-k})}{(1-\delta) f^{RM}(\theta_A^k)} > 0 \text{ if } \theta_A^{-k} > \theta_R^{-k},$$
$$\frac{\partial \tilde{s}_A^{-k}}{\partial \varepsilon}|_{\varepsilon=0} = \frac{\theta_A^{-k}(\theta_A^k - \theta_1^k)}{(1-\delta) f^{RM}(\theta_A^{-k})} > 0.$$

Substituting $\nu(s_i^k) = \kappa(s_i^k)^\alpha$, $\nu'(s_i^k) = \kappa\alpha(s_i^k)^{\alpha-1}$, and the partial derivatives above into Equation 1.49 yields

$$D_2 \equiv \alpha(s_1^k)^{\alpha-1}\left(-\frac{\theta_1^k\theta_A^{-k} - \omega_1^k + \frac{1-\delta}{\delta}s_1^k}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_1^k)^\alpha + \underbrace{\alpha(s_A^k)^{\alpha-1}\left(\frac{\theta_A^k(\theta_A^{-k} - \theta_R^{-k})}{(1-\delta)}\right)}_{d_2(\theta_A^k)}$$
$$+ \underbrace{\alpha(s_A^{-k})^{\alpha-1}\left(\frac{\theta_A^{-k}(\theta_A^k - \theta_1^k)}{(1-\delta)}\right)}_{=d_2(\theta_A^{-k})} + \alpha(s_R^{-k})^{\alpha-1}\left(-\frac{\theta_A^k\theta_R^{-k} - \omega_R^{-k} + \frac{1-\delta}{\delta}s_R^{-k}}{(1-\delta)}\right) + \frac{1-\delta}{\delta}(s_R^{-k})^\alpha.$$

Examining the two new terms shows that $d_2(\theta_A^{-k})$ is strictly increasing in α, and $d_2(\theta_A^k)$ is strictly increasing in α if $\theta_A^{-k} > \theta_R^{-k}$, and decreasing otherwise. Again, D_2 has at most one root.

Suppose $\beta_i^k, \omega_i^k, \theta_i^k, \theta_j^{-k}$ and δ are drawn uniformly from their continuous intervals. Then, there exists at most one α for which $D = 0$ (or $D_2 = 0$). Let α be drawn uniformly from $(0, 1)$, then random matching is generically suboptimal as such α is drawn with measure zero. $\square$

Proof of Proposition 2 Since market sides are fully symmetric, for brevity I drop the superscript k of each type θ_i^k. In this case, the positive assortative matching rule is defined as $\phi(\theta_i|\theta_j) = 1$ if and only if $i = j$ and results in steady-state mass $f(\theta_i) = \beta_i$ for any $\theta_i \in \Theta$.

Case 1: Search Fee

(a) "If" direction: PAM is optimal if the platform sets $s_i = \theta_i^2 - \omega_i$ for all $\theta_i \in \Theta$. As shown in Appendix A.2, PAM maximizes total match surplus across all agents. By choosing $s_i = \theta_i^2 - \omega_i$, the platform can extract each agent's match surplus as no agent is willing to pay more, thereby maximizing the platform's profit.

(b) "Only if" direction: Suppose, for contradiction, that PAM is profit-maximizing even if $s_i < \overline{s}_i$ for some $\theta_i \in \Theta \setminus \{\theta_1\}$ and

$$\overline{s}_i = \min\left\{\theta_i^2 - \omega_i, \theta_i^2 - \frac{\theta_i\theta_1 - \delta\omega_i}{1-\delta}\right\},$$

where the first entry is smaller than the second entry if $\omega_i > \theta_i\theta_1$. The platform's profit under PAM is strictly less than

$$\Pi^{PAM} < \frac{2(1-\delta)}{1-\rho}\left(\sum_{\theta_j \in \Theta\setminus\{\theta_i\}} \beta_j(\theta_j^2 - \omega_j) + \beta_i \overline{s}_i\right).$$

Next, observe that if the platform uses PAM with probability one in the next period, then type θ_i would reject the lowest type θ_1 in the (zero-probability) event they meet, since

$$\max\{\theta_i\theta_1, \omega_i\} < \delta\omega_i + (1-\delta)(\theta_i^2 - s_i). \tag{1.50}$$

Consider a deviation from PAM in which all types other than θ_1 and θ_i continue to only meet each other ($\Phi^D(\theta_j, \theta_j) = \beta_j$), but θ_1 and θ_i meet each other with mass $\epsilon \in (0, {}^{\min\{\beta_1,\beta_i\}}/_{\delta}]$ ($\Phi^D(\theta_i, \theta_1) = \epsilon$). Simultaneously, reduce the search fee of type θ_1 from $s_1 = \theta_1^2 - \omega_1$ to some s_1', which I will specify below. Then, I will show that there exists an $\epsilon > 0$ and a corresponding s_1' such that the resulting matching rule is feasible, incentive compatible, and strictly improves the platform's profit.

To check feasibility, substitute the steady state conditions in Equation 1.15 into the feasibility constraints in Equation 1.14:

$$\frac{\beta_i}{\delta + (1-\delta)\phi(\theta_i|\theta_i)} = \frac{\beta_i\phi(\theta_i|\theta_i)}{\delta + (1-\delta)\phi(\theta_i|\theta_i)} + \epsilon,$$
$$\frac{\beta_1}{\delta + (1-\delta)\phi(\theta_1|\theta_1)} = \frac{\beta_1\phi(\theta_1|\theta_1)}{\delta + (1-\delta)\phi(\theta_1|\theta_1)} + \epsilon.$$

Solving for the new (conditional) matching probabilities under the deviation, denoted ϕ^D, gives:

$$\phi^D(\theta_i|\theta_i) = \frac{\beta_i - \epsilon\delta}{\beta_i + (1-\delta)\epsilon}, \ \phi^D(\theta_1|\theta_1) = \frac{\beta_1 - \epsilon\delta}{\beta_1 + (1-\delta)\epsilon}. \tag{1.51}$$

For any $\min\{\beta_1,\beta_i\}/\delta > \epsilon > 0$, these probabilities are strictly less than one and larger than zero. Set the new search fee of type θ_1 to

$$s_1' = \frac{\beta_1 - \epsilon\delta}{\beta_1 + (1-\delta)\epsilon}(\theta_1^2 - \omega_1)$$

Next, I verify that the condition in Equation 1.50 for type θ_i remains satisfied. Under PAM, if $s_i < \overline{s}_i$, then the inequalities are slack. Since matching probabilities are continuous in ϵ, there exists a small $\epsilon > 0$ such that the condition remains non-binding or becomes just binding. Thus, type θ_i continues to reject matches with θ_1, and their search behavior does not change for sufficiently small ϵ.

Now consider the participation constraint of type θ_1 (Equation 1.12). Under PAM, its participation constraint is binding when $s_1 = \theta_1^2 - \omega_1$. Since θ_1 now meets type θ_i with positive probability, continuing to charge $s_1 = \theta_1^2 - \omega_1$ would violate the constraint. By lowering the search fee to s_1' as defined above, the constraint remains binding. The platform's profit given the new matching rule and search fees is

$$\begin{aligned}
\Pi^D &= \frac{2(1-\delta)}{1-\rho}\left(\frac{\beta_i s_i}{\delta + (1-\delta)\phi^D(\theta_i|\theta_i)} + \frac{\beta_1 s_1'}{\delta + (1-\delta)\phi^D(\theta_1|\theta_1)} + \sum_{j\neq 1,i}\beta_j(\theta_j^2 - \omega_j)\right), \\
&= \frac{2(1-\delta)}{1-\rho}\left((\beta_i + \epsilon)s_i + \frac{(\beta_1+\epsilon)(\beta_1 - \epsilon\delta)}{\beta_1 + (1-\delta)\epsilon}(\theta_1^2 - \omega_1) + \sum_{j\neq 1,i}\beta_j(\theta_j^2 - \omega_j)\right),
\end{aligned}$$

The deviation is profitable if $\Pi^D - \Pi^{PAM} > 0$, that is if

$$s_i - \frac{\varepsilon\delta}{\beta_1 + (1-\delta)\varepsilon}(\theta_1^2 - \omega_1) > 0,$$

which holds for $0 < \varepsilon < \frac{\beta_1 s_i}{\delta(\theta_1 - \omega_1) - (1-\delta)s_i}$.

Suppose for contradiction, that PAM is profit-maximizing even if $s_1 < \overline{s}_1$ for type θ_1

$$\begin{cases} \overline{s}_1 = \theta_1^2 - \omega_1 - \frac{\beta_i}{\beta_1}(\theta_i\theta_1 - \omega_i) & \text{if } \theta_i\theta_1 > \omega_i \\ \overline{s}_1 = \theta_1^2 - \omega_1 & \text{if } \theta_i\theta_1 \leq \omega_i. \end{cases}$$

Consider a deviation from PAM in which all types other than θ_1 and some type θ_i continue to only meet each other ($\Phi^D(\theta_j, \theta_j) = \beta_j$). For type θ_1 and θ_i, the platform chooses a new matching rule and search fee s_i' such that θ_i is indifferent between accepting and rejecting

θ_1 and θ_1 is indifferent between participating or not given $s_1 < \overline{s}_1$

$$\max\{\theta_i\theta_1, \omega_i\} = \frac{\delta\omega_i + (1-\delta)(-s_i + \phi(\theta_i|\theta_i)\theta_i^2)}{\delta + (1-\delta)\phi(\theta_i|\theta_i)},$$
$$\omega_1 = \frac{\delta\omega_1 + (1-\delta)(-s_1 + \phi(\theta_1|\theta_1)\theta_1^2)}{\delta + (1-\delta)\phi(\theta_1|\theta_1)}.$$

This results in matching rules $\phi(\theta_i|\theta_i) = \frac{(1-\delta)s_i + \delta(\theta_i\theta_1 - \omega_i)}{(1-\delta)\theta_i(\theta_i - \theta_1)}$ if $\theta_i\theta_1 > \omega_i$ or $\phi(\theta_i|\theta_i) = \frac{s_i}{\theta_i^2 - \omega_i}$ if $\theta_i\theta_1 \leq \omega_i$, and $\phi(\theta_1|\theta_1) = \frac{s_1}{\theta_1^2 - \omega_1}$. Given the matching rule and $s_1 < \theta_1^2 - \omega_1$, feasibility requires that s_i is chosen such that $f(\theta_i)(1-\phi(\theta_i|\theta_i)) = f(\theta_1)(1-\phi(\theta_1|\theta_1))$, where $f(\theta_k) = \frac{\beta_k}{\delta + (1-\delta)\phi(\theta_k|\theta_k)}$ for $k = 1, i$.

Then, the platform's deviation profit is larger than the profit under PAM when $s_1 < \overline{s}_1$. For $\delta \to 0$, it holds that

$$\beta_i \min\{\theta_i(\theta_i - \theta_1), \theta_i^2 - \omega_i\} + \beta_1(\theta_1^2 - \omega_1) \geq \beta_i(\theta_i^2 - \omega_i) + \beta_1 s_1,$$

where the inequality is strict if $\theta_i\theta_1 \leq \omega_i$ or if $\theta_i\theta_1 > \omega_i$ and $s_1 < \overline{s}_1$. This implies that there also exists a small movement to $\delta > 0$ where the inequality still holds.

Case 2: Advertisement

Suppose $\nu(\cdot)$ is strictly increasing, concave and fulfills the conditions in the proposition. Given PAM, the platform's profit is equal to

$$\Pi^{PAM} = \frac{2(1-\delta)}{1-\rho} \sum_{\theta_i \in \Theta} \nu(\theta_i^2 - \omega_i)\beta_i$$

where search costs are set to $s_i = \theta_i - \omega_i$ to maximize profits given PAM. Consider the a deviation in which all types other than type θ_1 and some type θ_i continue to meet each other:

$$\Phi'(\theta_j, \theta_j) = \Phi^{PAM}(\theta_j, \theta_j) = \beta_j, \forall \theta_j \in \Theta \setminus \{\theta_1, \theta_i\}.$$

For type θ_1 and θ_i choose the mass of recommended and matched pairs

$$\{\Phi'(\theta_1, \theta_1), \Phi'(\theta_1, \theta_i), \Phi'(\theta_i, \theta_i) m'(\theta_1, \theta_1), m'(\theta_1, \theta_i), m'(\theta_i, \theta_i)\}$$

such that $\beta_1 - \Phi'(\theta_1, \theta_1) = \beta_1 m'(\theta_1, \theta_1) = \varepsilon$ and $\beta_i - \Phi'(\theta_i, \theta_i) = \beta_i - m'(\theta_i, \theta_i) = \varepsilon$ for $\varepsilon \in (0, \min\{\beta_1, \beta_i\}]$. The new matching rule must satisfy the feasibility constraints

(Equation 1.14) and steady state conditions (Equation 1.15) below

$$\beta_1 + \frac{1-\delta}{\delta}\varepsilon = \frac{1}{\delta}\left(\beta_1 - (1-\delta)m'(\theta_1,\theta_1)\right), \tag{1.52}$$

$$\beta_1 + \frac{1-\delta}{\delta}\varepsilon = \Phi'(\theta_1,\theta_1) + \Phi'(\theta_1,\theta_i), \tag{1.53}$$

$$\beta_i + \frac{1-\delta}{\delta}\varepsilon = \frac{1}{\delta}\left(\beta_i - (1-\delta)m'(\theta_i,\theta_i)\right), \tag{1.54}$$

$$\beta_i + \frac{1-\delta}{\delta}\varepsilon = \Phi'(\theta_i,\theta_i) + \Phi'(\theta_1,\theta_i). \tag{1.55}$$

It follows that $\Phi'(\theta_1,\theta_i) = \frac{\varepsilon}{\delta}$. To ensure that type θ_i rejects type θ_1 under the new matching rule (so that $m'(\theta_1,\theta_i) = 0$), while type θ_1 participates, the platform chooses $(\tilde{s}_1, \tilde{s}_i)$ such that

$$\beta_i\left(\max\{0, \theta_i\theta_1 - \omega_i\} + \frac{(1-\delta)}{\delta}\tilde{s}_i\right) = (1-\delta)(\beta_i - \varepsilon)\left(\theta_i^2 - \omega_i + \frac{(1-\delta)}{\delta}\tilde{s}_i\right),$$

and

$$\beta_1\left(\frac{(1-\delta)}{\delta}\tilde{s}_1\right) = (1-\delta)(\beta_1 - \varepsilon)\left(\theta_1^2 - \omega_1 + \frac{(1-\delta)}{\delta}\tilde{s}_1\right),$$

hold, which results in

$$\tilde{s}_1 = \frac{(\beta_1-\varepsilon)(\theta_1^2-\omega_1)\delta}{\varepsilon+\delta(\beta_1-\varepsilon)};$$

$$\tilde{s}_i = \begin{cases} \frac{(\beta_i-\varepsilon)(\theta_i^2-\omega_i)\delta}{\varepsilon+\delta(\beta_i-\varepsilon)} & \text{if } \omega_i \geq \theta_i\theta_1 \\ \frac{\delta}{1-\delta}\frac{\beta_i\theta_i(\theta_i-\theta_1)-(\theta_i-\omega_i)(\varepsilon+\delta(\beta_i-\varepsilon)}{\varepsilon+\delta(\beta_i-\varepsilon)} & \text{otherwise }. \end{cases}$$

The deviation profit is therefore:

$$\Pi^D = \frac{2(1-\delta)}{1-\rho}\left(\nu(\tilde{s}_1)\left(\beta_1 + \frac{1-\delta}{\delta}\varepsilon\right) + \nu(\tilde{s}_i)\left(\beta_i + \frac{1-\delta}{\delta}\varepsilon\right) + \sum_{\theta_j\in\Theta\setminus\{\theta_1,\theta_i\}}\nu(\theta_j^2-\omega_j)\beta_j\right).$$

Then for $(s_1 = \theta_1^2 - \omega_1, s_i = \theta_i^2 - \omega_i)$, the deviation is profitable if $\Pi^D - \Pi^{PAM} > 0$:

$$D \equiv \nu(\tilde{s}_1)\left(\beta_1 + \frac{1-\delta}{\delta}\varepsilon\right) - \nu(s_1)\beta_1 + \nu(\tilde{s}_i)\left(\beta_i + \frac{1-\delta}{\delta}\varepsilon\right) - \nu(s_i)\beta_i \geq 0. \tag{1.56}$$

Rewriting the conditions yields

$$(1-\delta)(\nu(\tilde{s}_1)+\nu(\tilde{s}_i))\varepsilon \geq \delta[\beta_1(\nu(s_1)-\nu(\tilde{s}_1)) + \beta_i(\nu(s_i)-\nu(\tilde{s}_i))]$$

Rearranging for δ gives the following condition:

$$\delta \leq \frac{(\nu(\tilde{s}_1) + \nu(\tilde{s}_i))\varepsilon}{\beta_1(\nu(s_1) - \nu(\tilde{s}_1)) + \beta_i(\nu(s_i) - \nu(\tilde{s}_i)) + (\nu(\tilde{s}_1) + \nu(\tilde{s}_i))\varepsilon} \equiv \overline{\delta},$$

where $\overline{\delta} \in (0, 1)$ for $\varepsilon > 0$. $\square$

B.2 Binary Types

Lemma 11. *For $\delta \to 0$, the optimal matching rule that implements*
(a) A_{PAM} is

$$\begin{bmatrix} \frac{s}{\theta_h(\theta_h-\theta_l)} & 1-\frac{s}{\theta_h(\theta_h-\theta_l)} \\ 1-\frac{\beta_h(\theta_h(\theta_h-\theta_l)-s)}{\beta_h\theta_h(\theta_h-\theta_l)+(\beta_l-\beta_h)s} & \frac{\beta_h(\theta_h(\theta_h-\theta_l)-s)}{\beta_h\theta_h(\theta_h-\theta_l)+(\beta_l-\beta_h)s} \end{bmatrix}, \text{ if } \frac{\beta_h}{\beta_l} \le \frac{\theta_l^2-s}{\theta_h(\theta_h-\theta_l)-s}, \tag{1.57}$$

or otherwise,

$$\begin{bmatrix} \frac{\beta_h s}{\beta_l\theta_l^2+(\beta_h-\beta_l)s} & 1-\frac{\beta_h s}{\beta_l\theta_l^2+(\beta_h-\beta_l)s} \\ 1-\frac{s}{\theta_l^2} & \frac{s}{\theta_l^2} \end{bmatrix}, \text{ if } \frac{\beta_h}{\beta_l} \ge \frac{\theta_l^2-s}{\theta_h(\theta_h-\theta_l)-s}, \tag{1.58}$$

where at equality both matrices coincide. $\mathcal{O}(A_{PAM})$ is positive assortative.

(b) A_{WPAM} is

$$\begin{bmatrix} \frac{s}{\theta_h(\theta_h-\theta_l)} & 1-\frac{s}{\theta_h(\theta_h-\theta_l)} \\ 1-\frac{(\beta_l(\theta_h^2-s)-\beta_h(\theta_h(\theta_h-\theta_l)-s))s}{\theta_l(\theta_h-\theta_l)(\beta_h(\theta_h(\theta_h-\theta_l)-s)+\beta_l(\theta_h\theta_l-s))} & \frac{(\beta_l(\theta_h^2-s)-\beta_h(\theta_h(\theta_h-\theta_l)-s))s}{\theta_l(\theta_h-\theta_l)(\beta_h(\theta_h(\theta_h-\theta_l)-s)+\beta_l(\theta_h\theta_l-s))} \end{bmatrix}, \tag{1.59}$$

$$\text{if } \frac{(\theta_l^2-s)}{\theta_h(\theta_h-\theta_l)-s)} \le \frac{\beta_h}{\beta_l} \le \frac{(\theta_h^2-s)}{\theta_h(\theta_h-\theta_l)-s)}, \text{ and } \mathcal{O}(A_{WPAM}) \text{ is weakly assortative.}$$

A_{WPAM} is

$$\begin{bmatrix} \frac{\beta_h-\beta_l}{\beta_h} & 1-\frac{\beta_h-\beta_l}{\beta_h} \\ 1 & 0 \end{bmatrix}, \text{ if } \beta_h \ge \beta_l \text{ and } \frac{\beta_h-\beta_l}{\beta_h}\theta_h(\theta_h-\theta_l) \le s \le \theta_h\theta_l, \tag{1.60}$$

or

$$\begin{bmatrix} 0 & 1 \\ 1-\frac{\beta_h-\beta_l}{\beta_l} & \frac{\beta_h-\beta_l}{\beta_l} \end{bmatrix}, \text{ if } \beta_h \le \beta_l \text{ and } s \le \beta_l^2\theta_l^2+\beta_h\theta_l(\theta_h-\theta_l). \tag{1.61}$$

$\mathcal{O}(A_{WPAM})$ is non-assortative.

(c) A_{NAM} is

$$\begin{bmatrix} \frac{\beta_h-\beta_l}{\beta_h} & 1-\frac{\beta_h-\beta_l}{\beta_h} \\ \frac{s}{\theta_l(\theta_h-\theta_l)} & 1-\frac{s}{\theta_l(\theta_h-\theta_l)} \end{bmatrix}, \text{ if } 1 \le \frac{\beta_h}{\beta_l} \le \frac{\theta_h(\theta_h-\theta_l)}{\theta_h(\theta_h-\theta_l)-s}, \tag{1.62}$$

or

$$\begin{bmatrix} \frac{s}{\theta_h(\theta_h-\theta_l)} & 1-\frac{s}{\theta_h(\theta_h-\theta_l)} \\ 1-\frac{\beta_l(\theta_h^2-\theta_l^2-s)-\beta_h(\theta_h(\theta_h-\theta_l)-s)}{(\theta_h-\theta_l)\theta_l\beta_l} & \frac{\beta_l(\theta_h^2-\theta_l^2-s)-\beta_h(\theta_h(\theta_h-\theta_l)-s)}{(\theta_h-\theta_l)\theta_l\beta_l} \end{bmatrix}, \tag{1.63}$$

$$\text{if } \frac{\theta_h(\theta_h-\theta_l)}{\theta_h(\theta_h-\theta_l)-s} \le \frac{\beta_h}{\beta_l} \le \frac{\theta_h^2-\theta_l^2-s}{\theta_h(\theta_h-\theta_l)-s}, \text{ and } \mathcal{O}(A_{NAM}) \text{ is non-assortative.}$$

Proof of Lemma 11 The proof proceeds as follows. Fixing each matrix of mutual acceptance probabilities, I solve for the optimal matching rule by using the auxiliary problem from Appendix A.1. The linear program in the binary case is given by

$$\max \frac{2(1-\delta)s}{1-\rho}\left(f(\theta_h)+f(\theta_l)\right),$$

subject to feasibility and steady state conditions

$$f(\theta_h) = \Phi(\theta_h,\theta_h)+\Phi(\theta_h,\theta_l), \tag{1.64}$$
$$f(\theta_l) = \Phi(\theta_l,\theta_l)+\Phi(\theta_h,\theta_l), \tag{1.65}$$
$$\beta_h = f(\theta_h)\delta+(1-\delta)(\alpha(\theta_h,\theta_h)\Phi(\theta_h,\theta_h)+\alpha(\theta_h,\theta_l)\Phi(\theta_h,\theta_l)), \tag{1.66}$$
$$\beta_l = f(\theta_l)\delta+(1-\delta)(\alpha(\theta_l,\theta_l)\Phi(\theta_l,\theta_l)+\alpha(\theta_h,\theta_l)\Phi(\theta_h,\theta_l)), \tag{1.67}$$

as well as the respective participation and incentive constraints.

(a) A_{PAM}**:**

A_{PAM} induces the following constraints: A high type must be willing to continue searching after meeting a low type and the low type must be willing to participate. The transformed incentive and participation constraints take the following form

$$\beta_h(\delta\theta_h\theta_l+(1-\delta)s) \le (1-\delta)\Phi(\theta_h|\theta_h)(\delta\theta_h^2+(1-\delta)s), \tag{1.68}$$
$$\beta_l(1-\delta)s \le (1-\delta)\Phi(\theta_l|\theta_l)(\delta\theta_l^2+(1-\delta)s). \tag{1.69}$$

By Theorem 1 an optimal solution exists. In the binary case, the optimal solution can easily be checked. As the platform maximizes the steady state mass, it chooses $\Phi(\theta_h,\theta_h)$ and $\Phi(\theta_l,\theta_l)$ to be as small as possible without violating the constraints. Here, $\Phi(\theta_h,\theta_h)$ and $\Phi(\theta_l,\theta_l)$ are minimal when Equation 1.68 and Equation 1.69 bind resulting in

$$\Phi^{(a)}(\theta_h,\theta_h) = \frac{\beta_h((1-\delta)s+\delta\theta_h\theta_l)}{(1-\delta)((1-\delta)s+\delta\theta_h^2)},$$
$$\Phi^{(a)}(\theta_l,\theta_l) = \frac{\beta_l s}{(1-\delta)s+\delta\theta_l^2}.$$

Both the incentive and participation constraint, however, can only bind at the same time whenever

$$\left(\frac{\beta_h}{\beta_l}\right)^{(a)} = \frac{(1-\delta)(\theta_l^2-s)(s+\delta(\theta_h^2-s))}{(\theta_h(\theta_h-\theta_l)-s-\delta(\theta_h^2-s))(s+\delta(\theta_l^2-s))},$$

due to the feasibility constraints, Equation 1.64 and 1.65.

The steady state mass can be calculated by inserting $\Phi^{(a)}(\theta_h, \theta_h)$ and $\Phi^{(a)}(\theta_l, \theta_l)$ into

$$\begin{aligned} f(\theta_h) &= \frac{\beta_h - (1-\delta)\Phi(\theta_h, \theta_h)}{\delta}, \\ f(\theta_l) &= \frac{\beta_l - (1-\delta)\Phi(\theta_l, \theta_l)}{\delta}. \end{aligned}$$

The optimal matching rule is then given by $\phi(\theta_i|\theta_i) = \frac{\Phi(\theta_i,\theta_i)}{f(\theta_i)}$ for $i = h, l$.

If $\frac{\beta_h}{\beta_l} > (\frac{\beta_h}{\beta_l})^{(a)}$, only the participation constraint can be binding such that $\Phi(\theta_l, \theta_l) = \Phi^{(a)}(\theta_l, \theta_l)$. Inserting $\Phi(\theta_l, \theta_l) = \Phi^{(a)}(\theta_l, \theta_l)$ into the feasibility constraint of the low types yields $\Phi(\theta_h, \theta_l)$, which in turn determines $\Phi(\theta_h, \theta_h)$ by inserting it into the feasibility constraint of the high type. If $\frac{\beta_h}{\beta_l} < (\frac{\beta_h}{\beta_l})^{(a)}$, only the incentive constraint of the high type can be binding such that $\Phi(\theta_h, \theta_h) = \Phi(\theta_h, \theta_h)^{(a)}$ and the steps above can be repeated respectively.

(b) A_{WPAM}**:**

(b.1) A_{WPAM} induces the following constraints: A high type must be indifferent between searching and accepting low types

$$\beta_h(\delta\theta_h\theta_l + (1-\delta)s) = (1-\delta)\left(\Phi(\theta_h|\theta_h)(\delta\theta_h^2 + (1-\delta)s) + \alpha(\theta_h, \theta_l)\Phi(\theta_h, \theta_l)(\delta\theta_h\theta_l + (1-\delta)s)\right).$$

which holds for $\phi(\theta_h|\theta_h) = \frac{(1-\delta)s+\delta\theta_h\theta_l}{(1-\delta)\theta_h(\theta_h-\theta_l)}$. Additionally, low types must be willing to participate

$$\beta_l(1-\delta)s \leq (1-\delta)\left(\Phi(\theta_l, \theta_l)(\delta\theta_l^2 + (1-\delta)s) + \alpha(\theta_h, \theta_l)\Phi(\theta_h, \theta_l)(\delta\theta_h\theta_l + (1-\delta)s)\right).$$

From $\phi(\theta_h|\theta_h) = \frac{(1-\delta)s+\delta\theta_h\theta_l}{(1-\delta)\theta_h(\theta_h-\theta_l)}$ it follows

$$\begin{aligned} \Phi^{(b)}(\theta_h, \theta_h) &= \phi(\theta_h|\theta_h) \underbrace{\frac{\beta_h}{\delta + (1-\delta)(\phi(\theta_h|\theta_h) + \alpha(\theta_h, \theta_l)(1-\phi(\theta_h|\theta_h)))}}_{=f(\theta_h)} \\ &= \frac{\beta_h((1-\delta)s + \delta\theta_h\theta_l)}{(1-\delta)(\alpha(\theta_h, \theta_l)(\theta_h(\theta_h-\theta_l) - \delta\theta_h^2 - (1-\delta)s) + \delta\theta_h^2(1-\delta)s)} \end{aligned}$$

Then, $\Phi^{(b)}(\theta_h, \theta_l)$ follows by inserting $\Phi^{(b)}(\theta_h, \theta_h)$ in Equation 1.64, i.e.,

$$\frac{\beta_h\left(\theta_h\theta_l - (1-\delta)\theta_h^2 + (1-\delta)s\right)}{(1-\delta)(\alpha(\theta_h, \theta_l)(\delta\theta_h^2 - \delta s - \theta_h^2 + \theta_h\theta_l + s) - \delta\theta_h^2 + \delta s - s)}.$$

Furthermore, $\Phi^{(b)}(\theta_l, \theta_l)$ follows from feasibility of the low type by inserting $\Phi^{(b)}(\theta_h, \theta_l)$ into Equation 1.65.

The low type is indifferent between participating and not participating if

$$\alpha^{WPAM} \equiv$$
$$\left\{\alpha(\theta_h,\theta_l) : \beta_l s = \Phi^{(b)}(\theta_l,\theta_l)(\delta\theta_l^2 + (1-\delta)s) + \alpha(\theta_h,\theta_l)\Phi^{(b)}(\theta_h,\theta_l)(\delta\theta_h\theta_l + (1-\delta)s)\right\}.$$

For $\delta \to 0$, I get

$$\alpha^{WPAM} = \frac{s\left(\beta_h(\theta_h(\theta_h-\theta_l)-s) - \beta_l(\theta_l^2 - s)\right)}{(\theta_h(\theta_h-\theta_l)-s)\left(\beta_h\theta_l(\theta_h-\theta_l) + \beta_l\theta_l^2 + (\beta_h-\beta_l)s\right)}. \tag{1.70}$$

The mutual acceptance probability is then given by the above. For $\delta \to 0$, to ensure that $\alpha^{WPAM} \leq 1$ and $\phi(\theta_l|\theta_l) \geq 0$, the conditions in the lemma must hold.

(b.2) Additionally for $\beta_h \geq \beta_l$, the platform can implement A_{WPAM} by always matching low types with high types, i.e. $\phi(\theta_h|\theta_l) = 1$. This implies that low types search for only one period, such that $f(\theta_l) = \Phi(\theta_h,\theta_l) = \beta_l$. The high types' incentive constraint for $\alpha(\theta_h,\theta_l) = 1$ is

$$\beta_h(\delta\theta_h\theta_l + (1-\delta)s) \geq (1-\delta)\left(\Phi(\theta_h|\theta_h)(\delta\theta_h^2 + (1-\delta)s) + \Phi(\theta_h,\theta_l)(\delta\theta_h\theta_l + (1-\delta)s)\right),$$

and from the feasibility constraint (Equation 1.64), it follows that $\Phi(\theta_h,\theta_h) = \beta_h - \beta_l$. The incentive constraint of high types is satisfied if

$$s \geq \frac{\beta_h - (1-\delta)\beta_l}{(1-\delta)\beta_h}\theta_h(\theta_h-\theta_l) - \frac{\delta}{(1-\delta)}\theta_h^2$$

The participation constraint of low types is satisfied if $s \leq \theta_h\theta_l$:

$$\beta_l(1-\delta)s \leq (1-\delta)\beta_l(\delta\theta_h\theta_l + (1-\delta)s).$$

Lastly for $\beta_h \leq \beta_l$, the platform can implement A_{WPAM} by always matching high types to low types, i.e. $\phi(\theta_l|\theta_h) = 1$. This implies that high types search for only one period, such that $f(\theta_h) = \Phi(\theta_h,\theta_l) = \beta_h$. Low types must be willing to participate

$$\beta_l(1-\delta)s \leq (1-\delta)\left(\Phi(\theta_l,\theta_l)(\delta\theta_l^2 + (1-\delta)s) + \beta_h(\delta\theta_h\theta_l + (1-\delta)s)\right).$$

If the participation constraint is satisfied, low types also search for only one period, such that $f(\theta_l) = \beta_l$. Therefore, $\Phi(\theta_l,\theta_l) = \beta_l - \beta_h$. Thus, the participation constraint is satisfied if

$$s \leq \beta_l^2\theta_l^2 + \beta_h\theta_l(\theta_h-\theta_l),$$

and low types do not reject low types if

$$s \geq \frac{\delta\theta_l(\beta_h(1-\delta)\theta_h - \beta_l\theta_l)}{(1-\delta)(\beta_l - (1-\delta)\beta_h)},$$

which equals zero for $\delta \to 0$.

(c) A_{NAM}**:**

(c.1) A_{NAM} can be implemented if

$$\beta_h((1-\delta)s + \delta\theta_h\theta_l) \geq (1-\delta)\Phi(\theta_h|\theta_h)((1-\delta)s + \delta\theta_h^2) + (1-\delta)\Phi(\theta_h,\theta_l)((1-\delta)s + \delta\theta_h\theta_l),$$
$$\beta_l((1-\delta)s + \delta\theta_l^2) \leq (1-\delta)\Phi(\theta_h|\theta_l)((1-\delta)s + \delta\theta_h\theta_l).$$

As high types accept both high and low types and search for only one period, the steady state mass of high types is equal to their inflow: $f(\theta_h) = \beta_h$. The platform's profit from high types is, therefore, independent of the matching rule. To maximize profits, the platform minimizes $\Phi(\theta_h, \theta_l)$ such that

$$\Phi(\theta_h,\theta_l) = \frac{\beta_l((1-\delta)s + \delta\theta_l^2)}{(1-\delta)((1-\delta)s + \delta\theta_h\theta_l)},$$

and the incentive constraint of the low type binds. $\Phi(\theta_h,\theta_h) = \beta_h - \beta_l$ follows from the feasibility constraints (Equation 1.64), where $\Phi(\theta_h,\theta_h)$ and $\Phi(\theta_h,\theta_l)$ must be such that the incentive constraint of the high type is fulfilled, which is true if

$$1 \leq \frac{\beta_h}{\beta_l} \leq \frac{((1-\delta)s + \delta\theta_l^2)\,\theta_h(\theta_h - \theta_l)}{(\theta_h(\theta_h-\theta_l) - (1-\delta)s - \delta\theta_h^2)\,((1-\delta)s + \delta\theta_h\theta_l)}.$$

For $\delta \to 0$ this results in

$$1 \leq \frac{\beta_h}{\beta_l} \leq \frac{\theta_h(\theta_h-\theta_l)}{\theta_h(\theta_h-\theta_l) - s}.$$

(c.2) A_{NAM} can be implemented if a high type is indifferent between accepting and rejecting a low type, while a low type is willing to reject low types. Again as in part (b), $\phi(\theta_h|\theta_h) = \frac{(1-\delta)s+\delta\theta_h\theta_l}{(1-\delta)\theta_h(\theta_h-\theta_l)}$ must hold to ensure the indifference constraint of high types. Then for $\alpha(\theta_h,\theta_l) \in [0,1]$, $\Phi^{(c)}(\theta_h,\theta_h) = \Phi^{(b)}(\theta_h,\theta_h)$ and $\Phi^{(c)}(\theta_h,\theta_l) = \Phi^{(b)}(\theta_h,\theta_l)$. Inserting into the incentive constraint of the low type, the low type rejects low types if

$$\beta_l((1-\delta)s + \delta\theta_l^2) \leq (1-\delta)\alpha(\theta_h,\theta_l)\Phi^{(c)}(\theta_h,\theta_l)((1-\delta)s + \delta\theta_h\theta_l),$$

which holds with equality for

$$\alpha^{NAM} = \frac{\beta_l s}{(\beta_h - \beta_l)(\theta_h(\theta_h - \theta_l) - s)} \tag{1.71}$$

if $\delta \to 0$. It holds that $\alpha^{NAM} > 0$ generally, and $\alpha^{NAM} \leq 1$ if $\frac{\beta_h}{\beta_l} \geq \frac{\theta_h(\theta_h - \theta_l)}{\theta_h(\theta_h-\theta_l)-s}$. Additionally, $\phi(\theta_h|\theta_l) = \frac{(\beta_h-\beta_l)(\theta_h^2-\theta_l^2-s)}{\beta_l\theta_l(\theta_h-\theta_l)}$, which is larger than zero if $\beta_h \geq \beta_l$ and smaller than one if $\frac{\beta_h}{\beta_l} \leq \frac{\theta_h^2-\theta_l^2-s}{\theta_h(\theta_h-\theta_l)-s}$. $\square$

Proof of Proposition 3 Next, I determine the platform's preferred outcome. First, let $s \leq \theta_l^2$.
(i) For $\frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(a)}$, the profit when implementing A_{PAM} (Equation 1.57) is

$$\Pi^{(a.1)} = \frac{2\nu(s)(1-\delta)}{1-\rho}\left(\frac{2\beta_h\theta_h(\theta_h-\theta_l) + (\beta_l - \beta_h)(s + \delta(\theta_h^2 - s))}{s + \delta(\theta_h^2 - s)}\right).$$

For $\frac{\beta_h}{\beta_l} \geq \left(\frac{\beta_h}{\beta_l}\right)^{(a)}$, the platform can either implement A_{PAM} (Equation 1.58) or A_{WPAM} (Equation 1.59). The profits are

$$\Pi^{(a.2)} = \frac{2\nu(s)(1-\delta)}{1-\rho}\left(\frac{2\beta_l\theta_l^2 + (\beta_h - \beta_l)(s + \delta(\theta_l^2 - s))}{s + \delta(\theta_l^2 - s)}\right)$$

and

$$\Pi^{(b.1)} = \frac{2\nu(s)(1-\delta)}{1-\rho}\frac{(2\beta_h\theta_h^2\theta_l - (\beta_h - \beta_l)(2\theta_l^2 - s - \theta_l s) - \delta(\beta_h - \beta_l)(\theta_h - \theta_l)(s + \theta_h\theta_l))}{(\theta_h + \theta_l)(s + \delta(\theta_h\theta_l - s))},$$

where the difference is positive

$$\Pi^{(b.1)} - \Pi^{(a.1)} \geq 0.$$

Thus for $\left(\frac{\beta_h}{\beta_l}\right)^{(a)} \leq \frac{\beta_h}{\beta_l} \leq \frac{(1-\delta)\theta_h^2-s}{\theta_h(\theta_h-\theta_l)-s+\delta(\theta_h^2-s)}$ the platform implements A_{WPAM} and A_{PAM} if $\frac{\beta_h}{\beta_l} \geq \frac{(1-\delta)\theta_h^2-s}{\theta_h(\theta_h-\theta_l)-s+\delta(\theta_h^2-s)}$.

It remains to compare the profit in equilibrium (b) when implementing A_{WPAM} against the profit from equilibrium (c) when implementing A_{NAM}. Note that for $s \leq \theta_l^2$, the profit when implementing A_{WPAM} is maximized in (b.1) as agents in both equilibria in (b.2) only search for one period. The profit in (c) is

$$\Pi^{(c.1)} = \frac{2\nu(s)(1-\delta)}{1-\rho}\left(\beta_h + \frac{\beta_l\theta_l(\theta_h - \theta_l)}{s + \delta(\theta_h\theta_l - s)}\right)$$

or

$$\Pi^{(c.2)} = \frac{2\nu(s)(1-\delta)}{1-\rho}$$
$$\left(\frac{s(\beta_h-\beta_l)\theta_h(\theta_h-\theta_l)+s\beta_l\theta_l(\theta_h-\theta_l)+\delta(\beta_h\theta_h(\theta_h-\theta_l)(\theta_h\theta_l-s)+\beta_l(\theta_h-\theta_l)^2(s+\theta_h\theta_l)}{(s+\delta(\theta_h^2-s))(s+\delta(\theta_h\theta_l-s))}\right).$$

Then, it holds that $\Pi^{(b)} \geq \Pi^{(c.1)}, \Pi^{(c.2)}$.
(ii) Let $\theta_l^2 \leq s \leq \theta_h\theta_l$. Then, the platform can only implement A_{WPAM} or A_{NAM}. Alternatively, the platform can exclude low types from participating. Recall that $\Pi^{(c.1)}$ and $\Pi^{(c.2)}$ are strictly dominated by $\Pi^{(b.1)}$. Therefore, the platform implements either A_{WPAM} in Equation 1.59, 1.60, or 1.61. If $\beta_h \geq \beta_l$, the platform can either implement A_{WPAM} in Equation 1.59 or 1.61. If $\beta_h < \beta_l$, the platform can either implement A_{WPAM} in Equation 1.59 or 1.60. In this case, however, for too large s no low type is willing to participate such that the platform excludes low types. Note that at $s = \theta_h\theta_l$, the matching outcome is non-assortative if $\beta_h \geq \beta_l$, whereas only high types participate if $\beta_h < \beta_l$.
(iii) Let $\theta_h\theta_l \leq s \leq \theta_h^2$. If search costs are larger than $\theta_h\theta_l$, low types are no longer willing to participate. To maximize surplus from high types, the platform sets $\phi(\theta_h|\theta_h) = 1$. □

Proof of Proposition 4 (i) The platform implements A_{PAM} together with the matching rule as in Lemma 11 (a). As the positive assortative matching outcome maximizes match productivity, the welfare loss from mismatch is zero. For $\frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(a)}$, agents' expected search time is

$$\mathcal{T}(\theta_h) = \frac{\theta_h(\theta_h-\theta_l)}{s+\delta(\theta_h^2-s)},$$
$$\mathcal{T}(\theta_l) = \frac{\beta_h(\theta_h(\theta_h-\theta_l)+(\beta_l-\beta_h)s+\delta(\beta_l-\beta_h)(\theta_h^2-s)}{\beta_l(s+\delta(\theta_h^2-s))}.$$

Observe that $\mathcal{T}(\theta_h)$ is decreasing in s and δ. Differentiating $\mathcal{T}(\theta_l)$ with respect to s and δ yields

$$\frac{\partial\mathcal{T}(\theta_l)}{\partial s} = -\frac{\beta_h(1-\delta)\theta_h(\theta_h-\theta_l)}{\beta_l(s+\delta(\theta_h^2-s))^2} < 0,$$
$$\frac{\partial\mathcal{T}(\theta_l)}{\partial \delta} = -\frac{\beta_h\theta_h(\theta_h-\theta_l)(\theta_h^2-s)}{\beta_l(s+\delta(\theta_h^2-s))^2} < 0,$$

i.e. $\mathcal{T}(\theta_h)$ is decreasing in s and δ as well.

For $\frac{\beta_h}{\beta_l} \geq \frac{(1-\delta)\theta_h^2 - s}{\theta_h(\theta_h-\theta_l)-s+\delta(\theta_h^2-s)}$, agents' expected search time is

$$\mathcal{T}(\theta_h) = \frac{\beta_l\theta_l^2 + (\beta_h - \beta_l)(s + \delta(\theta_l^2 - s))}{\beta_h(s + \delta(\theta_l^2 - s)},$$
$$\mathcal{T}(\theta_l) = \frac{\theta_l^2}{s + \delta(\theta_l^2 - s)}.$$

Observe that $\mathcal{T}(\theta_l)$ is decreasing in s and δ. Differentiating $\mathcal{T}(\theta_h)$ with respect to s and δ yields

$$\frac{\partial \mathcal{T}(\theta_h)}{\partial s} = -\frac{\beta_l(1-\delta)\theta_l^2}{\beta_h(s + \delta(\theta_l^2 - s))^2} < 0,$$
$$\frac{\partial \mathcal{T}(\theta_h)}{\partial \delta} = -\frac{\beta_l(\theta_l^2 - s)\theta_l^2}{\beta_h(s + \delta(\theta_l^2 - s))^2} < 0,$$

i.e. $\mathcal{T}(\theta_h)$ is decreasing in s and δ as well.

(ii) The platform implements A_{WPAM} together with the matching rule as in Lemma 11 (b.1). The welfare loss from mismatches is

$$\mathcal{W}_{WPAM} = \alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l)(\theta_h - \theta_l)^2,$$

and agents' expected search time is

$$\mathcal{T}(\theta_h) = \frac{\beta_h - (1-\delta)(\Phi^{(b)}(\theta_h\theta_h) + \alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l))}{\beta_h\delta},$$
$$\mathcal{T}(\theta_l) = \frac{\beta_l - (1-\delta)(\Phi^{(b)}(\theta_l\theta_l) + \alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l))}{\beta_l\delta},$$

where

$$\Phi^{(b)}(\theta_h, \theta_h) = \frac{\beta_h\theta_h\theta_l - (\beta_h - \beta_l)(1-\delta)(\theta_l^2 - s)}{(1-\delta)(\theta_h^2 - \theta_l^2)},$$
$$\Phi^{(b)}(\theta_l, \theta_l) = \frac{\beta_h\theta_h\theta_l - (\beta_h - \beta_l)(1-\delta)(\theta_h^2 - s)}{(1-\delta)(\theta_h^2 - \theta_l^2)},$$

which are both increasing (decreasing) in s if $\beta_h > \beta_l$ ($\beta_h < \beta_l$) and increasing in δ. Note that $\Phi^{(b)}(\theta_h, \theta_l)$ followed from feasibility (see proof of Lemma 11) and α_{WPAM} is set to fulfill the low types' participation constraint. Using the implicit function theorem and differentiating the participation constraint with respect to s yields

$$\beta_l - (1-\delta)(\alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l) + \Phi^{(b)}(\theta_l, \theta_l)) =$$
$$\frac{\partial\Phi^{(b)}(\theta_l, \theta_l)}{\partial s}(\delta\theta_l^2 + (1-\delta)s) + \frac{\partial\alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l)}{\partial s}(\delta\theta_h^2 + (1-\delta)s),$$

where the left-hand side corresponds to $\delta f(\theta_l) > 0$ and $\Phi^{(b)}(\theta_l, \theta_l)$ is increasing in s if $\beta_h > \beta_l$ and decreasing otherwise. Thus, it follows that $\alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l)$ must be increasing in s if $\beta_l > \beta_h$ and either in-or decreasing for $\beta_l < \beta_h$ (depending on the parameter values). Using the implicit function theorem and differentiating the participation constraint with respect to δ yields

$$
\begin{aligned}
0 = & \frac{\partial \Phi^{(b)}(\theta_l, \theta_l)}{\partial \delta}(\delta\theta_l^2 + (1-\delta)s) + \Phi^{(b)}(\theta_l, \theta_l)(\theta_l^2 - s) \\
& + \frac{\partial \alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l)}{\partial \delta}(\delta\theta_h^2 + (1-\delta)s) + \alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l)(\theta_h\theta_l - s),
\end{aligned}
$$

As $\Phi^{(b)}(\theta_l, \theta_l)$ is increasing in δ, $\alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l)$ must be decreasing in δ for $s \leq \theta_l^2$. For δ for $s > \theta_l^2$, $\alpha_{WPAM}\Phi^{(b)}(\theta_h, \theta_l)$ can be either in- or decreasing in δ.

It follows that $\mathcal{W}_{WPAM}$ is increasing in s if $\beta_l > \beta_h$ and either in-or decreasing for $\beta_l < \beta_h$ (depending on the parameter values). Furthermore, $\mathcal{W}_{WPAM}$ is decreasing in δ for $s \leq \theta_l^2$ and either in- or decreasing for $s > \theta_l^2$.

Differentiating $\mathcal{T}(\cdot)$ with respect to s and δ yields

$$
\begin{aligned}
\frac{\partial \mathcal{T}(\theta_h)}{\partial s} &= -\frac{(1-\delta)\theta_h\theta_l\left(\beta_l\left((\theta_h-\theta_l)\delta+\theta_l\right)+(\theta_h-\theta_l)(1-\delta)\beta_h\right)}{\left((\theta_h\theta_l-s)\delta+s\right)^2\beta_h(\theta_h+\theta_l)} < 0, \\
\frac{\partial \mathcal{T}(\theta_l)}{\partial s} &= -\frac{(1-\delta)\theta_h\theta_l\left(\beta_h(\theta_h-\theta_l)(1-\delta)+\beta_l\left((\theta_h-\theta_l)\delta+\theta_l\right)\right)}{\left((\theta_h\theta_l-s)\delta+s\right)^2(\theta_h+\theta_l)} < 0, \\
\frac{\partial \mathcal{T}(\theta_h)}{\partial \delta} &= -\frac{\theta_h^2\theta_l\left(\beta_h\theta_l(\theta_h-\theta_l)+\beta_l(\theta_l^2-s)\right)}{(\delta\theta_h\theta_l+s(1-\delta))^2\beta_h(\theta_h+\theta_l)} < 0, \\
\frac{\partial \mathcal{T}(\theta_l)}{\partial \delta} &= -\frac{\theta_h^2\theta_l\left(\beta_h\theta_l(\theta_h-\theta_l)+\beta_l(\theta_l^2-s)\right)}{(\delta\theta_h\theta_l+s(1-\delta))^2(\theta_h+\theta_l)} < 0.
\end{aligned}
$$

That is, $\mathcal{T}(\cdot)$ is decreasing in s and δ.

(iii) The platform implements A_{WPAM} by the matching rule as in Lemma 11 (b.2). The matching outcome is non-assortative. The welfare loss from mismatch is

$$
\mathcal{W}_{NAM} = -2\beta_l(\theta_h - \theta_l)^2,
$$

and agents search for one period only. □

Proof of Proposition 5 The proof follows the structure of Proposition 3. Note all matching outcomes in Proposition 3 are implemented when choosing the search fee except the positive assortative matching outcome for $\frac{\beta_h}{\beta_l} \geq \left(\frac{\beta_h}{\beta_l}\right)^{(b)}$.

(i) Suppose the platform implements A_{PAM} for $\frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(a)}$. Recall that for $s = \theta_l^2$, $\left(\frac{\beta_h}{\beta_l}\right)^{(a)} = 0$ and thus A_{PAM} can never be implemented if there is a positive inflow of

both types. The platform maximizes its profit with respect to s under the constraint that $s \in [0, \theta_l^2]$ and the condition $\frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(a)}$ is still fulfilled.
(ii) Suppose the platform implements A_{WPAM} for $\left(\frac{\beta_h}{\beta_l}\right)^{(a)} \leq \frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(b)}$. There exists an $\theta_h(\theta_h - \theta_l) > \overline{s} > \theta_l^2$ such that if $s > \overline{s}$ A_{WPAM} can never be implemented if there is a positive inflow of both types. The platform maximizes its profit with respect to s under the constraint that $s \in [0, \overline{s}]$ and the condition $\left(\frac{\beta_h}{\beta_l}\right)^{(a)} \leq \frac{\beta_h}{\beta_l} \leq \left(\frac{\beta_h}{\beta_l}\right)^{(b)}$ is still fulfilled.
(iii) Suppose the platform implements A_{NAM} for $\beta_h > \beta_l$. Then, agents only search for one period. Therefore, the platform increases the search fee as much as possible. By Proposition 3, the upper limit is given by $s = \theta_h\theta_l$.
(iv) Lastly, the platform can exclude low types from participating. To maximize profits, the platform extract the surplus from high types by setting $s = \theta_h^2$. The platform does so for sufficiently high $\frac{\beta_h}{\beta_l}$.

The platform does not implement the positive assortative matching outcome for $\frac{\beta_h}{\beta_l} \geq \left(\frac{\beta_h}{\beta_l}\right)^{(b)}$. Recall that profits are

$$\frac{2(1-\delta)}{1-\rho}\left(\frac{2\beta_l\theta_l^2 s}{(1-\delta)s + \delta\theta_l^2} + (\beta_h - \beta_l)s\right).$$

Note that both terms are strictly increasing in s such that the platform would choose $s = \theta_l^2$ resulting in

$$\frac{2(1-\delta)}{1-\rho}\left(\beta_h\theta_l^2 + \beta_l\theta_l^2\right).$$

As A_{NAM} can be implemented for $\beta_h > \beta_l$ with $s = \theta_h\theta_l$, the profit from A_{NAM} is always strictly larger than the profit from A_{PAM} for $\frac{\beta_h}{\beta_l} \geq \left(\frac{\beta_h}{\beta_l}\right)^{(b)}$. $\square$

Proof of Proposition 6 For the proof of the first sentence, I first examine the first-order condition in Equation 1.9. Observe that $\frac{\nu(s)}{\nu'(s)}$ is strictly increasing in s due to the concavity of $\nu(s)$. Additionally, from Proposition 4 it follows that $\partial f(\theta_i^k)/\partial s < 0$, i.e. $f(\theta_i^k)$ is strictly decreasing in s. From the proof of Proposition 4, it follows that $\partial^2 f(\theta_i^k)/\partial s^2 > 0$. Thus, the right-hand side of Equation 1.9 decreases in s and the left-hand side increases in s which implies the statement.

To prove the second sentence, note that the profit for charging $s = \theta_h\theta_l$ and implementing $\mathcal{O}_N AM$ is

$$\frac{2(1-\delta)}{(1-\rho)}\nu(\theta_h\theta_l)(\beta_h + \beta_l),$$

and for charging $s = \theta_h^2$ and excluding low types is

$$\frac{2(1-\delta)}{(1-\rho)}\nu(\theta_h^2)(\beta_h).$$

Then, if $\nu(\cdot)$ is sufficiently concave, the platform prefers to charge a lower price for a discrete increase in demand:

$$\frac{\nu(\theta_h\theta_l)}{\nu(\theta_h^2)} \geq \frac{\beta_h}{\beta_h+\beta_l}(< 1).$$

□

Proof of Proposition 7 To characterize the profit-maximizing solution with overconfident users, note first that it is optimal for the platform to have all three types participate. Otherwise, the platform can always increase profits by including the formerly excludes type by charging a positive fee and matching them to each other. Consider the feasible mutual acceptance matrices of the form

$$\begin{bmatrix} \alpha(\theta_h,\theta_h) & \alpha(\theta_h,\theta_l) & \alpha(\theta_h,\hat{\theta}_l) \\ \alpha(\theta_l,\theta_h) & \alpha(\theta_l,\theta_l) & \alpha(\theta_l,\hat{\theta}_l) \\ \alpha(\hat{\theta}_l,\theta_h) & \alpha(\hat{\theta}_l,\theta_l) & \alpha(\hat{\theta}_l,\hat{\theta}_l). \end{bmatrix}$$

As overconfident users perceive to have the same continuation value as high types, $V^C(\theta_h)$, they follow the same acceptance strategy. That is, overconfident users accept high types with probability one and low types with probability $\alpha \in [0,1]$ if and only if high types do. Furthermore, overconfident users are accepted by high (low) types with positive probability if and only if high (low) types accept low types with positive probability.

The feasible mutual acceptance matrices are

$$A_1 = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}, A_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, A_3 = \begin{bmatrix} 1 & \alpha' & \alpha' \\ \alpha' & 1 & \alpha' \\ \alpha' & \alpha' & (\alpha')^2 \end{bmatrix}, A_4 = \begin{bmatrix} 1 & \alpha'' & \alpha'' \\ \alpha'' & 0 & 0 \\ \alpha'' & 0 & (\alpha'')^2 \end{bmatrix},$$

for $\alpha' \in (0,1]$ and $\alpha'' \in (0,1]$. The profit from implementing A_1 is given in the text preceding Proposition 7. Observe that implementing $A_2 - A_4$ can induce search for more than one period for at least one type.

To implement A_2–A_4, the incentive constraint ensuring that high types reject low types with positive probability must hold. The platform then maximizes revenue from both high and low types by maximizing match surplus and extracting it through the search fee conditional on leaving a rent of $\theta_h\theta_l$ to high types. From Appendix A.2, match surplus is maximized under positive assortative matching—that is, when the platform implements A_2. Moreover, agents must search for only one period; otherwise, surplus is lost due to $\delta > 0$.

Revenue from overconfident agents is maximized under $A_2 - A_4$ when they search for $\frac{1}{\delta}$ periods—i.e., no one they match with accepts them, and s_h is maximized. Under A_2, this is exactly the case: overconfident types are rejected, they search for $\frac{1}{\delta}$ periods, and the platform captures the match surplus from high types through s_h, i.e. s_h is maximal. Thus, it follows that the relevant constraints are given by

$$\theta_h\theta_l \leq \frac{(1-\delta)(-s+\phi(\theta_h|\theta_h)\theta_h^2)}{\delta+(1-\delta)\phi(\theta_h|\theta_h)}, \tag{IC-θ_h}$$

$$\theta_h\theta_l \leq \frac{(1-\delta)(-s+\phi(\theta_h|\theta_h)\theta_h^2)}{\delta+(1-\delta)\phi(\theta_h|\theta_h)}, \tag{PIC-$\hat{\theta}_l$}$$

$$0 \leq \frac{(1-\delta)(-s+\phi(\theta_l|\theta_l)\theta_l^2)}{\delta+(1-\delta)\phi(\theta_l|\theta_l)}. \tag{PC-θ_l}$$

From the steady state constraints, Equation 1.15, the platform's profit maximization problem can be written as

$$\frac{\beta_h s_h}{\delta+(1-\delta)\phi(\theta_h|\theta_h)} + \frac{\beta_l(1-\lambda)}{\delta+(1-\delta)\phi(\theta_l|\theta_l)} + \frac{\beta_l\lambda}{\delta},$$

subject to feasibility constraints, Equation 1.14, and the three incentive and participation constraints above. It can easily be verified that $s_l = \theta_l^2$ and $s_h = \theta_h(\theta_h - \theta_l) - {}^{d}/{}_{1-\delta}\theta_h\theta_l$ and $\phi(\theta_h|\theta_h) = 1, \phi(\theta_l|\theta_l) = 1$ and $\phi(\hat{\theta}_l|\hat{\theta}_l) = 1$ maximize the platform's profit and satisfy all constraints with equality. The platform's profit, Π_S^{OC}, is given in Proposition 7, where λ^* is derived by setting $\Pi_S^{OC} = \Pi_{PAM}^{OC}$ and solving for λ. □

C Appendix: Tables

Application	App Price	Subscriptions	One-Time Purchases
Tinder	Free	Tinder Gold (1 Week): $13.99 – 18.99 Tinder Gold (1 Month): $14.99 – 24.99 Tinder Plus (1 Month): $9.99	1 Boost: $3.99 – 7.99 3 Super Likes: $9.99 5 Super Likes: $4.99
Bumble	Free	Bumble Premium (1 Year) $129.99 – 169.99	5 Spotlights + Compliments $24.99 – 29.99 15 Spotlights + Compliments $44.99 – 59.99 30 Spotlights + Compliments $79.99 – 99.99
Hinge	Free	Hinge+ Subscription (1 Week): $16.99 Hinge Subscription (1 Month): $29.99 – 34.99 Membership (1 Month): $19.99 Hinge Subscription (1 Week): $14.99 HingeX Subscription: $24.99	Bundle of three Roses: $9.99 Bundle of twelve Roses: $29.99 Boost: $9.99 – 19.99
Match	Free	Match (1 Month): $19.99 – 42.99 Match (3 Months): $74.99 Match (6 Months): $129.99 Standard (1 Month): $44.99 Basic (1 Months): $44.99 Platinum (1 Week): $29.99	1 Top Spot: $2.99 Top Spot 10-Pack: $19.99 Boost 1-Pack: $5.99
Hily	Free	Hily Premium (1 Week): $14.99 Profile boost (1 Week): $5.99 – 9.99 Premium+ (1 Week): $24.99 Hily Elixir (1 Week): $19.99	1 Unblur: $4.99 5 Unblur: $12.99
Plenty of Fish	Free	Upgrade (1 Month): $19.99 Upgrade (3 Months): $38.99 Premium Membership (1 Month): $29.99	1 Token: $1.99 5 Tokens: $8.99 10 Tokens: $17.99
Badoo	Free	Badoo Premium (1 Week): $5.99 – 8.99 Super Powers (1 Week): $2.99 Super Powers (1 Months): $11.99	Pack of 100 Credits: $1.99 – 3.99
Coffee Meets Bagel	Free	Premium (1 Month): $14.99 – 34.99 Premium (3 Months): $74.99 Premium (6 Months): $71.99 Platinum (1 Month): $46.99 Platinum (3 Month): $99.99	200 Coffee Beans: $2.99 400 Coffee Beans: $4.99 3000 Coffee Beans: $24.99
Raya	Free	Membership (1 Month): $24.99 Membership (6 Month): $113.99 Raya+ Membership: $49.99	30 Extra Likes: $10.99 Skip the Wait: $7.99 5 Skip the Waits: $29.99 1 Direct Request: $4.99 3 Direct Requests: $12.99 12 Direct Requests: $49.99
MeetMe	Free	MeetMe (1 Month): $7.99 MeetMe (3 Months): $17.99 MeetMe+ (1 Month): $7.99	Pack of 200 Credits: $1.99 Pack of 500 Credits: $1.99 – 4.99 Pack of 1800 Credits: $14.99 Pack of 14500 Credits: $99.99 Pack of 3200 Credits: $24.99

Table A.1: A Selection of Dating Apps in the US Apple Store

Application	App Price	Subscriptions	One-Time Purchases
Tinder	Free	Tinder Gold (1 Week): 13,99 € Tinder Gold (1 Month): 8,99 – 27,49 € Tinder Platinum (1 Month): 32,99 €	1 Boost: 7,99 – 9,99 € 3 Super-Likes: 11,99 € 5 Super-Likes: 5,99 – 9,99 €
Bumble	Free	Bumble Premium (1 Week): 14,99 – 19,99 € Bumble Boost (1 Week): 5,99 – 6,99 € Bumble Premium (1 Month): 34,99 €	
Hinge	Free	Hinge+ Sub (1 Week): 14,99 € Hinge+ Sub (1 Month): 24,99 € HingeX Sub (1 Week): 24,99 €	Bundle of twelve Roses: 24,99 € Bundle of three Roses: 7,99 € One Superboost: 14,99 € One Boost: 7,99 €
LOVOO	Free	Lovoo Premium (1 Month): 11,99 – 24,99 €	300 Credits: 5,99 € 500 Credits: 4,99 € 550 Credits: 7,99 € 3000 Credits: 19,99£ € 5 Icebreaker: 5,99 € Unbegrenzte Likes: 1,19 €
Badoo	Free	Badoo Premium (1 Week): 5,99 – 7,99 € Badoo Premium (1 Month): 19,99 €	100 Badoo Credits: 1,99 – 4,99 € 550 Badoo Credits: 12,99 € Super Powers (1 Woche): 2,99 € Super Powers (1 Monat): 8,99 € Super Powers (1 Woche): 2,99 €
Parship	Free	Premium lite (6 Month): 209,99 – 229,99 € Premium classic (1 Year): 224,99 – 499,99 € Premium Comfort: 249,99 €	Parship Premium: 9,99 €
OkCupid	Free	OkCupid Premium (1 Month): 15,99 – 32,99 € OKCupid Premium (3 Month): 65,99 €	1 Boost: 1,99 – 7,99 € 2 Superlikes: 7,99 €
Raya	Free	Membership (1 Month): 18,99 € Membership (6 Month): 83,99 € Raya+ Membership (1 Month): 44,99 €	Skip the Wait 7,99 € 3 Direct Requests 12,99 € 1 Direct Request 4,99 € 30 Extra Likes 10,99 € 5 Skip the Waits 29,99 € 1 Travel Plan 9,99 €
LoveScout24	Free	Lovescout24 (1 Month): 39,99 € Mobile Plus (1 Month): 9,99 € Mobile Plus (1 Week): 4,99 € Lovescout24 (1 Week): 9,99 € Lovescout24 (3 Month): 89,99 €	1 Booster: 1,99 € Wer sucht mich?: 1,99 € Boost: 1,99 € Dateroulette: 2,99 € Favouriten-Funk: 1,99 €
ElitePartner	Free	ElitePartner Premium Go: 3,99 – 19,99 € Premium plus (1 Year): 399,99 € Premium basic (6 Months): 279,99 € Premium comfort (2 years) : 599,99 €	

Table A.2: A Selection of Dating Apps in the German Apple Store

App Name	Price	Contains Ads	Prices of In-App Purchases	Number of Installations
German Store				
happn	Free	Yes	0.59 – 274.99 €	100M+
Badoo	Free	Yes	0.39 – 244.99 €	100M+
Tinder	Free	Yes	0.29 – 324.99 €	100M+
SweetMeet	Free	Yes	0.59 – 219.99 €	50M+
Bumble	Free	No	0.29 – 314.99 €	50M+
BLOOM	Free	Yes	1.49 – 299.00 €	10M+
OkCupid	Free	Yes	0.71 – 194.99 €	10M+
Zoosk	Free	Yes	0.50 – 434.99 €	10M+
Mamba	Free	Yes	0.50 – 294.99 €	10M+
Boo	Free	Yes	0.46 – 218.85 €	10M+
US Store				
happn	Free	Yes	\$0.49 – 224.99	100M+
Badoo	Free	Yes	\$0.49 – 239.99	100M+
Tinder	Free	Yes	\$0.49 – 299.99	100M+
SweetMeet	Free	Yes	\$0.99 – 199.99	50M+
Bumble	Free	No	\$0.49 – 259.99	50M+
BLOOM	Free	Yes	\$1.99 – 349.00	10M+
OkCupid	Free	Yes	\$0.99 – 179.99	10M+
Zoosk	Free	Yes	\$0.49 – 399.99	10M+
Mamba	Free	Yes	\$0.99 – 264.99	10M+
Boo	Free	Yes	\$1.00 – 269.99	10M+

Table A.3: Most Popular Dating Apps in the German and US Google Play Store

App Name	App Price	In-App Purchases	Price	Adds	In-App Purchases	No. of Downloads
US Apple Store			US Android Store			
LinkedIn	Free	Career (1 Month): \$29, 99 – 39, 99 Business (1 Month): \$69, 99	Free	Yes	\$7.49 – 839.88	1B+
Indeed	Free	None	Free	Yes	none	100M+
Glassdoor	Free	None	Free	No	none	10M+
ZipRecruiter	Free	None	Free	No	none	10M+
Monster	Free	None	Free	No	none	5M+
German Apple Store			German Android Store			
LinkedIn	Free	Essentials (1 Month): 9, 99 € Career (1 Month): 29, 99 – 39, 99 € Business (1 Month): 69, 99 €	Free	Yes	7, 00 – 839, 88 €	1B+
Indeed	Free	None	Free	Yes	none	100M+
Glassdoor	Free	None	Free	Yes	none	10M+
Stepstone	Free	None	Free	Yes	none	10M+
Monster	Free	None	Free	Yes	none	5M+
Costs for Recruiters						
LinkedIn	The standard account is free. Premium accounts cost between 40 – 125 €/\$ (See above)					
Indeed	There is an option for free listings. Costly adds are charged per click, with a minimum of 5 €/\$ per day					
Glassdoor	No information					
Monster	Two Options: Monster+ Standard: Pay per Click and Monster+ Pro: 749€/\$299 per month					
Stepstone	Multiple tiers: “Campus” 199 €,“Select”: 329€, “Pro”: 1399 €, Pro Plus: 1699 €, Pro Ultimate: 2399 €					
Zip Recruiter	Pricing depends on the number of job ads. Ads are charged per day and per add:“Standard”: \$16, “Premium”: \$24 Plans are charged per ad and per month: “Standard”: \$299, “Premium”: \$419, “Pro”: \$719					

Table A.4: Selected Job Platforms in the German and US App Store

Chapter 2

Faking Network Size

2.1 Introduction

Joining a new platform often involves uncertainty for an individual user who cannot be sure of its overall popularity. With uncertainty about a platform's network size, users need to form beliefs about the participation decisions of others. Based on the data platforms collect and the in-depth knowledge of their business model, platforms naturally have an informational advantage relative to individual users. Moreover, they have an incentive to let consumers believe that joining will enable them to reap large network effects. Given these incentives, this paper investigates a platform's ability to use "fake profiles" to its advantage both when users are sophisticated about and when they are unaware of the platform's ability to do so.

The use of fake profiles is a common phenomenon on digital platforms. Prominent cases are those of Dating platforms. For example, the Federal Trade Commission (FTC) sued the Match Group for using fake profiles to persuade users to upgrade to a paid subscription. Other dating platforms use company-created fake profiles to interact with users on the platform, giving them the impression of a real contact with users often being unaware of this practice.[1] Further examples include cryptocurrency exchange platforms, which are under investigation by the SEC for engaging in trading financial assets themselves to artificially inflate the trading volume (so-called wash trading). Recent studies show that about 70% of unregulated trades are subject to wash trading (Cong et al., 2023). The economic costs to users and platforms are substantial. If fake profiles induce users to hold incorrect belief about the platform, they may make inefficient participation decisions. Furthermore, creating fake profiles is costly to the platform without generating additional value.

Formally, I investigate how a monopoly platform uses multiple signals to convince users of its network size. In particular, users can learn from the price they observe, a (cheap-talk) message, and the network size. Users are uncertain about the distribution of stand-alone values provided by the platform, while the platform has private information about this fundamental. Given the information asymmetry, suboptimal membership fees and fake profiles set by the platform are both costly signals about the fundamental. Fake profiles can increase the perceived network size but do not generate network effects ex post.

Users observe the membership fee first and then decide whether to join the platform. Thereafter, users who joined observe the perceived network size and decide whether to exit

[1] See `https://www.ftc.gov/news-events/press-releases/2019/09/ftc-sues-owner-online-dating-service-matchcom-using-fake-love`, last visited 01.09.2020); `https://www.verbraucherzentrale.de/wissen/digitale-welt/onlinedienste/onlinedating-auf-diesen-portalen-flirten-fakeprofile-21848`, last visited 01.09.2020 or `https://www.faz.net/aktuell/wirtschaft/unternehmen/straftaten-schiessen-wegen-datingplattform-in-die-hoehe-18792428.html`.

prior to paying the membership fee (following a "free-trial period"). In contrast to most of the signaling literature, I analyze a game with multiple signaling instruments, continuous signals and a continuum of receivers, where receivers (users) care not only about the sender's (platform's) action, but also about the action of other receivers (users).

Absent fake profiles, the platform's only signaling instrument about its size is the price. If users cannot observe demand and pay the price upfront, prices must be distorted upwards above full information prices to credibly signal a high fundamental. Only in higher states can a platform set inefficiently high prices to optimally separate from lower states. Including the possibility of creating fake profiles, users' understanding thereof is crucial when evaluating the market outcome. Before users pay the price, they observe the perceived network size without being able to distinguish between real and fake profiles. Sophisticated users, however, are fully aware of the platform's practice, whereas naive users are unaware of the possible use of fake profiles or believe that fake profiles are forbidden and hence not used.

Sophisticated users anticipate the platform's incentives correctly and hence, discount the perceived network size by the expected amount of fake profiles. In that case, both fake profiles and high prices are costly in that they reduce profits taken demand as given, and hence are substitutes for signaling a high fundamental. Abstracting from existence issues, the platform can always fully differentiate itself from those with less users through costly signaling based either on inefficiently high prices or the use of costly fake profiles. I identify parameter conditions such that the latter separating equilibrium exists, whereas the former always exists. Given its existence, in the unique separating equilibrium the platform with the lowest fundamental sets its full information price and all other platforms need to create fake profiles and distort their prices. Otherwise, the unique equilibrium has the same properties as the equilibrium absent fake profiles.

In contrast, if users dogmatically believe that every profile is real, i.e. they never considered the possibility of creating fake profiles, the platform uses the cheap-talk message to communicate the expected network size upfront. Users blindly believe this message upfront and the corresponding network size later on. The platform can exploit this misperception by using fake profiles to signal an unrealistic high network size, and thus value, to the users. In equilibrium, the platform always prefers to lie and deceive users by pretending that their network size is larger than it actually is. With a bounded state space, the platform with the highest fundamental, however, cannot induce unrealistic high user beliefs. Hence, depending on the costs of fake profiles, platforms below a threshold lie by creating fake profiles and set a higher price, and those above induce the highest possible belief about the state. This results in pooling on the observable instruments, but differentiation on the unobservable instrument.

The results imply that a platform would like to commit to refrain from using fake profiles with sophisticated users. Fake profiles are a wasteful investment for the platform used for separation. Rather than observing that platforms commit on not using fake profiles, it can be observed that platforms actually hide the use of fake profiles in their terms and conditions. This, however, likely indicates that users are mainly naive in these markets, which renders fake profiles profitable. With sophisticated users the platform would benefit from a regulation that provides commitment for the platform's claim that the observed network size is the true network size. This would result in full information prices being incentive-compatible as lying by platforms with a low fundamental is detectable and punished by exiting and non-paying users. If platforms are not able to credibly commit, however, they can profit from fake profiles with sophisticated users: signaling via fake profiles can lower the overall signaling costs when compared to signaling via distortionary prices only.

Methodologically, I apply an adjusted version of the D1 criterion developed by Banks and Sobel (1987) to refine the set of Perfect Bayesian Equilibria. It is well known that in certain classes of games, the D1 criterion selects a unique equilibrium outcome, which is separating, whenever there is a single receiver (or multiple receivers whose decisions are strategically independent). I extend this result to strategically interdependent receivers by imposing a restriction on the coordination problem of users' entry decision off the equilibrium path.

The remainder of the paper proceeds as follows. The related literature is discussed in Section 2.1.1. Section 2.2 describes the model and discusses potential applications. Section 2.3 analyzes the model when users are sophisticated, whereas Section 2.4 provides the analysis for naive users. Section 2.5 discusses common cases of fake profiles and Section 2.6 concludes. All omitted proofs are in Appendix B.

2.1.1 Related Literature

This paper is the first to introduce signaling into a model of platform adoption. As such it is related to models of platforms when there is incomplete information. Technically, it is related to the literature on signaling with multiple instruments. Since I allow for users who have incorrect beliefs, it is also related to papers on misleading consumers. I discuss these related papers below.

Platform Markets This paper belongs to the relatively sparse literature incorporating issues of incomplete information and asymmetric information (Halaburda et al., 2018b; Jullien and Pavan, 2019; Ke and Zhu, 2021; Kang and Muir, 2022) on platforms. Most models in the literature on platforms and two-sided markets assume complete information (Caillaud and Jullien, 2003; Rochet and Tirole, 2003, 2006; Armstrong, 2006; Halaburda et al., 2018a; Gal-Or, 2020). To the best of my knowledge, no paper has investigated asymmetric information between the platform and its users where the platform holds private information. The closest paper with respect to modeling the incomplete information is Jullien and Pavan (2019) who consider a platform market in which both users and platforms face uncertainty about participation decisions due to dispersion of information about their preferences. Especially given the growing importance of big data, I consider the more realistic case in which the platform has (superior) private information regarding its desirability to potential users. This, however, implies that the platform's choices act as a signal. Contrary to Jullien and Pavan (2019), I consider a platform with only one market side, or a platform on which both sides of the market are identical.

Signaling and Advertising As the monopoly platform has private information in my model, it is closely related to the literature on signaling (Kreps and Sobel, 1994), where signalling games with a continuum of states are studied by Mailath (1987) and in particular Mailath and von Thadden (2013). Fake profiles have not been studied in this context. Papers such as Kihlstrom and Riordan (1984), Milgrom and Roberts (1986), and Bagwell and Ramey (1988) that study the use of costly advertisement in combination with prices are conceptually closely related although they do not incorporate network effects. Fake profiles resemble persuasive advertisement, which is assumed to shift the willingness to pay of users (see Bagwell (2007) for an overview on advertisement). In a signaling model, Rhodes and Wilson (2018) analyzed false advertising used by firms to overstate the value of their products. False advertisement is only costly whenever it is punishable by a third-party. Buyers, nevertheless, may be affected by false advertisement in equilibrium. A key difference is how the amount of fake profiles is determined in my model in which not only the cost function but also equilibrium prices and demand determine the amount of fake profiles. Furthermore, the fact that fake profiles might not be observable to users influence their equilibrium amount.

The paper adds to the literature on signaling by identifying a novel channel — network effects — that makes signaling via price or fake profiles credible. Main channels in the literature on signaling are: 1) repeated purchases, 2) cost differences between qualities, and 3) information differences between users. Although learning by users bears similarities to repeated purchases, price signaling in my model even works absent learning. Due to the presence of network effects an increase in the price has two effects. First, the price has a direct effect on users' utility lowering their willingness to participate on the platform. It follows that additionally, the price also has an indirect effect on users' utility through the reduced participation decision of others, which further reduces their willingness to participate. Without network effects, users would not care about the state as their participation decision would be independent of those of other users.

Due to equilibrium multiplicity in signaling games, a wide range of papers focuses on appropriate equilibrium refinements (Cho and Kreps, 1987; Banks and Sobel, 1987; Cho and Sobel, 1990). As users exert positive externalities on each other, in my model the most prominent refinements in the literature fail to select a unique equilibrium. Therefore, I adopt a version of the D1 criterion for a continuum of states as in (Ramey, 1996) and impose a further (weak) restriction on the receivers' strategies off-path: that they are rationalizable (Bernheim, 1984; Pearce, 1984) given a common belief.

Consumer Naïveté The model investigates the effects of different types of user sophistication on the market outcome when users face fake profiles — thereby adding to the literature on consumer naïveté. See Heidhues and Kőszegi (2018) for a survey on the growing literature on how consumer naïveté affects market outcomes. While recent policy papers suggest (Crémer et al., 2019; Scott Morton et al., 2019; Fletcher et al., 2021), behavioral effects are particularly important in digital settings, academic research on this topic is scarce. My paper is among the first formal models to introduce consumer naïveté in platform markets. Others include Johnen and Somogyi (2024) who analyze and compare the sellers' and the platforms' incentive to hide parts of the price from naive consumers. They find that a platform has strong incentives to shroud additional fees if it increases perceived consumer surplus. Conceptually, my paper is closely related to work on consumer naivete in cheap talk models (Ottaviani and Squintani, 2006; Kartik et al., 2007; Chen, 2011) that analyze the impact of naive or credulous consumers who blindly believe the sender's message. In contrast to these papers, creating fake profiles is costly giving rise to signaling issues.

Manipulating Consumer Expectations More broadly, the paper is connected to the literature that studies the manipulation of consumer expectations, especially in network markets. Early contributions focus on the expectations of early adopters of a network good. More recently, the emergence of fake reviews for products on platforms is studied. Evidence of fake reviews on for example Amazon.com is provided by He et al. (2022) or Expedia.com and TripAdvisor.com by Mayzlin et al. (2014). Theoretic treatments of fake reviews can be found in Glazer et al. (2021) and Yasui (2020), where the most closely related paper is Knapp (2022). The author analyzes a cheap-talk game in which a reviewer of a good may create a truthful or fake review but abstracts from the platform setting with network effects. Similar to my paper, consumers differ in their understanding of the possibility of fake reviews (naive or sophisticated).

2.2 Model

I analyze a sender-receiver model with two types of players: a platform (sender) and a group of potential users (receivers) of mass one. The platform has private information about a fundamental $\theta \in [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$ that determines the users' distribution of stand-alone values r on the platform $F(r|\theta) \equiv F_\theta(r)$. The common prior about the fundamental is $\mu_0(\theta)$, a continuous probability distribution, and has full support on $\Theta \equiv [\underline{\theta}, \overline{\theta}]$.

The platform and users engage in the following four-period game. First, nature draws the fundamental $\theta \in \Theta$. The platform observes θ and sets a price p and message m, where the message space is restricted to the type space. Additionally, the platform may invest in fake profiles, where the amount of fake profiles is given by ξ. Then, upon observing the platform's message and price (p, m), as well as their own stand-alone value r_i, users decide whether or not to enter. Users who joined observe the perceived number of users and decide whether to exit. The perceived number will depend on the actual mass of users and fake profiles, in a way detailed below. Lastly, the platform collects fees.

Users: Payoff Users have a common outside option normalized to zero. They vary, however, in the utility they receive from joining the platform — their stand-alone value r_i. User i obtains utility

$$v_i = r_i + \beta n - p,$$

where the distribution of stand-alone values r, F_θ, is continuous with full and strictly positive support. Users benefit from positive network effects, β, and from the mass of users that stay on the platform, n, but pay price p.

Users: Actions and Beliefs Upon having learned about the true fundamental, the platform sets a price, sends a message, and determines the number of fake profiles. First, after observing a price-message pair, users update about the fundamental and form a belief $\mu(\theta|p, m)$ and, then after learning their individual stand-alone value, form a belief $\mu_1(\theta|p, m, r_i)$. Second, after joining the platform, users observe the perceived mass of users, which is a function of the mass of users who have joined and the mass of fake profiles. The corresponding belief is denoted by $\mu_2(\theta|r, p, m, \mathbb{I})$, where $\mathbb{I}$ denotes the information structure. Depending on the users' ability to learn about or observe fake profiles, they update about the fundamental based on the information structure $\mathbb{I} = \{\emptyset, \{[0,1], \mathbb{R}_0^+\}, [0,1] + \mathbb{R}_0^+\}$. Users may either not observe the network size at all, $\mathbb{I} = \emptyset$, observe the true network size and fake profiles separately, $\mathbb{I} = \{[0,1], \mathbb{R}_0^+\}$, or observe the sum of both, $\mathbb{I} = [0,1] + \mathbb{R}_0^+$, which may include fake profiles. Among these users, sophisticated users are aware of the possibility of fake profiles, while naive users blindly believe the message sent and take the network size at face value for values below or equal to one.[2]

[2] As the true network size is at most equal to a mass of one when all users enter, naive users take the network size at face value as long as it does not exceed a value of one. For values above one, naive users are free to hold any belief about the state, where I restrict attention to naive users holding the most pessimistic belief. Imposing any other belief such as the most optimistic belief, however, does not affect

User i's entry strategy in the first period is a mapping $\sigma_i^1 : \mathbb{R} \times M \to [0,1]$ from prices and messages to entry. For given price p and belief $\mu_1(\theta|\cdot)$, a user enters if their expected utility from entering is higher than their outside option. The aggregate entry decision of users depends on the distribution of r in society. Following entry, users update their beliefs to $\mu_2(\theta|\cdot)$ and decide whether to exit the platform; formally, their exit strategy in period two — given that the user entered in period one — is given by $\sigma_i^2 : \mathbb{R} \times M \times \mathbb{I} \to [0,1]$.

Platform: Payoff and Actions The platform is a monopolist that chooses a price-message pair (p, m) with $p \in \mathbb{R}_+$ and $m \in M = \Theta$ and a number of fake profiles $\xi \in \mathbb{R}_0^+$. The platform's strategy maps the state space into prices, messages, and fake profiles $\sigma^P : \Theta \to \mathbb{R} \times \Theta \times \mathbb{R}_0^+$. The platform maximizes its profit with respect to prices, messages, and fake profiles

$$\max_{p,\xi} (p - c)n(\theta, \mu, p) - \gamma\xi,$$

where $n(\theta, \mu, p)$ is the mass of users that stay on the platform given the true fundamental, their belief about it, and price p. Let c denote the marginal cost of the platform to serve one user and γ the marginal cost of creating a fake profile.

Equilibrium Concept The equilibrium concept is Perfect Bayesian Equilibrium (PBE) if all users are sophisticated. Strategies are optimal given beliefs at every information set. Beliefs of sophisticated users are updated via Bayes' rule whenever possible. At each information node, users optimize given their beliefs (sequentially rationality).

If users are naive, I use a Perception-Perfect Equilibrium (PPE). Naive users form their beliefs through the following rule. In the first period, naive users blindly believe in the fundamental stated by the observed message, and thus hold point beliefs. Following entry, users take the network size at face value (for a mass below or equal to one). If the observed network size confirms the expected network size given first period belief, μ_1^N, and price, p, the belief remains the same. If the observed network size, $\hat{n}$, does not match the expected network size, naive users revise their belief. The naive users' new belief must satisfy the following condition

$$\mu_2^N \equiv \{\theta' \in \Theta | n(\theta', \mu_1^N, p) = \hat{n}\}.$$

Naive users maximize expected utility given their beliefs.

As a tie-breaking rule, I impose that users enter only if they expect to stay: Whenever a user is indifferent between not joining the platform or joining the platform but leaving in period 3, I assume that the user does not enter.

the equilibrium as long as the message space is restricted to the type space and users understand that the type space is bounded.

Equilibrium Refinement For sophisticated users, off-path the set of equilibria is refined by adapting the D1 Criterion of Ramey (1996) to a signaling game in which the receivers strategically interact. Intuitively, under D1 users' out-of equilibrium beliefs put positive mass only to the types that are most likely to profit from a deviation from equilibrium. As users' participation decisions depend on the decision of other users and are thus not strategically independent, I impose a restriction on the coordination aspect of users' entry decision. For a given price, I suppose this induces a common receiver belief (as it does on the path of the play). Consumers take this common belief as given, and then resolve the coordination problem among themselves in the same way as they would if this belief was common knowledge. With common knowledge of the state, there is a unique rationalizable entry decision suggesting that the coordination problem should be resolved in exactly that way. The precise definition is given by Definition 4 in Appendix A. For naive users, I do not use an equilibrium refinement. Beliefs (on and off-path) are naively given by the simple rule specified above.

Assumptions To analyze the game, I impose regularity conditions on the family of distributions $F_\theta(r)$ and the strength of the network effect β.

Assumption 2. *The distributions $F_\theta(r)$ are*

1. *twice differentiable in r and θ with density $f_\theta(r)$,*
2. *where the corresponding densities $f_\theta(r)$ are single-peaked in r, and*
3. *the distributions have a (weakly) increasing hazard rate $\lambda(r;\theta)$ in r, and*
4. *common support.*

The above assumption on F_θ ensures that the optimization problem of the platform is well-behaved under complete information and that there exists a unique (monopoly) price. The assumption on common support and single-peakedness can be relaxed to allow for the family of uniform distributions as well.[3]

Assumption 3. *(MLRP) For $\theta > \theta'$, f_θ likelihood dominates $f_{\theta'}$: $\frac{f_\theta(r)}{f_{\theta'}(r)}$ is an increasing function.*

The monotone likelihood ratio property implies that F_θ first-order stochastically dominates $F_{\theta'}$ and hazard rate $\lambda(p,n;\theta)$ is strictly monotonically decreasing in θ. First-order

[3] Assuming non-common support has the following implication for the analysis. After observing the stand-alone value, users with a high stand-alone value form a belief that puts zero probability on states that are not possible. This does not affect the analysis of separating equilibria on-path, but plays a role for incentive-compatibility as a deviation to a lower state cannot be credible to those users. During most of the analysis, however, the relevant incentive compatibility is for a low type to mimic a high type. In pooling equilibria beliefs are dispersed and the analysis remains unchanged.

stochastic dominance implies that higher states lead to higher demand. The latter yields that the price elasticity increases with increasing θ holding participation constant. Hence, a higher state induces higher monopoly prices c.p. Lastly, to exclude multiplicity of (continuation) equilibria network effects β cannot be too strong.

Assumption 4. *(Network effects)* $\beta \in \{\beta \in \mathbb{R}_+ : 1/2 - \beta \max_{\theta,r} f_\theta(r) > 0,\ \theta \in \Theta\}$.

Networks effects must be small enough to avoid multiplicity of continuation equilibria, and hence, guarantees uniqueness of continuation equilibria.

2.2.1 Preliminaries

Under the assumptions made for any price the platform has set, there exists a unique cutoff strategy for users, even if the information is incomplete. Each user has private information about their own reservation value. All users with reservation values above the cutoff participate in the platform, while users below the cutoff do not. The first lemma defines the cutoff.

Lemma 1. *(Unique cutoff) In any equilibrium in which users hold a common belief upon observing (p, m), users use a cutoff strategy. The unique cutoff is given by*

$$r_c = p - \beta \int_\Theta (1 - F_\theta(r_c))\mu(\theta|r_c, p)d\theta, \tag{2.1}$$

which results in $n(\theta, \mu, p) = 1 - F_\theta(r_c)$ agents.

This lemma implies that users also follow a unique cutoff strategy in an equilibrium in which some types pool, i.e. when there is incomplete separation. Users' beliefs are dispersed as although all users have a common prior, they draw inferences from their own r. As a result, after observing a price and their own reservation value, users hold different beliefs. To establish the lemma, however, it is sufficient to suppose that upon observing the price but not yet their standalone value, users hold a common belief. On path this must be fulfilled because all users rely on Bayes rule, whereas off-path the common belief assumption is imposed.

As a benchmark, the next lemma characterizes the full information benchmark which corresponds to the first-best solution in prices and user participation.

Lemma 2. *(First-best) Under full information, there exists a unique equilibrium. In this equilibrium, the platform's profit maximizing price $p^{FI}(\theta)$ satisfies*

$$p^{FI} - c = \frac{1 - F_\theta(r^{FI})}{f_\theta(r^{FI})}(1 - \beta f_\theta(r^{FI})), \tag{2.2}$$

where r^{FI} denotes the equilibrium cutoff given p^{FI}. The full information price is strictly monotonically increasing in θ if the density $f_\theta(r^{FI})$ is strictly decreasing in θ.

It follows that the mark-up is always positive and hence, the price is always above marginal cost. If $F_\theta(\cdot)$ is the exponential distribution with scale parameter θ, which satisfies the MLRP, then indeed $f^{1,\theta}(r^{FI}) < 0$ and thus, the full information price increases in θ. Lastly, Assumptions 2-4 imply that the platform's profit for $\mu = \hat{\theta}$ fulfills the (strict) single-crossing property.

Lemma 3. *The platform's profit function, $\pi(\theta, \hat{\theta}, p)$, satisfies the single-crossing property, namely*

$$\frac{\partial}{\partial \theta}\left(\frac{\frac{\partial \pi(\theta,\hat{\theta},p)}{\partial p}}{\frac{\partial \pi(\theta,\hat{\theta},p)}{\partial \hat{\theta}}}\right) > 0.$$

2.3 Price and Message as Signals

In this section, I will discuss a benchmark for analyzing the effectiveness of signaling on platforms in which the presence of fake profiles does not impact demand. The benchmark assumes that users enter the platform and decide whether to stay without being able to observe the network size. In other words, the timing is as if user pay for their membership before joining the platform. Users cannot learn from the network size before their purchase decision, making fake profiles irrelevant. Hence, the participation and purchasing decision happen simultaneously represented by the information structure $\mathbb{I} = \emptyset$.

2.3.1 Sophisticated Users

Sophisticated users are rational and fully understand the signaling game. For those users, the price is the only credible signal and the message is ignored. Hence, I will suppress the message in the section below. For ease of exposition, the main part will focus on the construction of separating equilibria. The platform uses a one-to-one strategy $\tau : \Theta \to \mathbb{R}$ that maps the state to its chosen price and therefore, users hold a common belief on the path of play. I will focus on differentiable separating strategies τ.

Definition 1. A separating equilibrium consist of the platform's strategy τ, users' strategy σ_i and beliefs, μ, such that:

1. For any $p \in \tau(\Theta), \mu(p) = \tau^{-1}(p)$,
2. For any $\theta \in \Theta, \tau(\theta) \in \arg\max_{p\in\mathbb{R}_+} \pi(\theta, \mu(p), p)$ (Incentive Compatibility).

The platform maximizes its profit with respect to the price given that users form their beliefs according to $\mu(\theta|p, m, r) = \tau^{-1}(p)$. With a slight abuse of notation $n(\theta, \tau^{-1}(p), p)$ denotes the network size based on the true state θ, the belief $\tau^{-1}(p)$ which is a Dirac measure, and the price. Therefore, when the platform increases its price, the effects on profit are two-fold. The first effect is the direct price effect on the mark-up and demand, whereas the second effect is the belief effect, i.e. a higher price potentially signals a higher state. The platform's pricing strategy is determined by

$$\{\tau(\theta)\} \equiv \underset{p \in \tau \in ([\underline{\theta}, \overline{\theta}])}{\arg\max} (p-c) n(\theta, \tau^{-1}(p), p).$$

Assumptions 2-4 ensure that the profit is differentiable. In any separating equilibrium, rational users learn about the true state from the separating strategy. Focusing on differentiable separating strategies the first-order condition can be used. The first-order condition given that in equilibrium beliefs are correct, i.e., $\tau^{-1}(p) = \theta$ yields

$$n(\theta, \theta, p) + (p-c) \left. \frac{\partial n(\theta, \tau^{-1}(p), p)}{\partial p} \right|_{\tau^{-1}(p)=\theta} + (\tau^{-1}(p))'(p-c) \left. \frac{\partial n(\theta, \tau^{-1}(p), p)}{\partial \tau^{-1}(p)} \right|_{\tau^{-1}(p)=\theta} = 0.$$

The separating strategy $\tau(\theta)$ is given by the differential equation

$$\tau'(\theta) = -\frac{(\tau - c) n_2(\theta, \theta, \tau)}{n(\theta, \theta, \tau) + (\tau - c) n_3(\theta, \theta, \tau)}, \tag{2.3}$$

where $n_2(\cdot)$ and $n_3(\cdot)$ denote the partial derivatives with respect to the second and third arguments, respectively. Observe that setting $p = p^{FI}(\theta)$ from Equation 2.2 sets the denominator equal zero. Hence, setting the complete information prices for all types is not a solution. Prices must be distorted. Sequential rationality implies setting the initial value condition to $\tau(\underline{\theta}) = p^{FI}(\underline{\theta})$, i.e., the lowest type cannot do better than setting their first-best price. Given the initial value condition, there exists a unique solution to the differential equation that minimizes the level of costly signaling.

Proposition 1. *Suppose $\mathbb{I} = \emptyset$. Then under the equilibrium refinement in Definition 4, there exists a differentiable separating equilibrium outcome in which the equilibrium price $p^{S,*}$ is given by Equation 2.3 with $\tau(\underline{\theta}) = p^{FI}(\underline{\theta})$.*

Given differentiability and the initial value condition, the (differentiable) separating equilibrium is unique (Mailath, 1987). In the separating equilibrium, the platform "burns money" to credibly communicate its type to its users taking the form of distorted prices. The price as signaling device is feasible as the marginal cost of a price increase depends on the demand curvature, which in turn is influenced by the platform's true state. As shown in Lemma 3, the platform is more willing to trade-off and increase in price against an increase in demand. This link between the true fundamental and price is established by the network effects that arise on the platform. Thus, the incentive-compatible separating strategy must be increasing in the state as signed in the next corollary.

Corollary 4. *The equilibrium price $p^{S,*}$ is increasing in θ and is always greater than the full information price.*

To sign the pricing distortion, it is useful to recall that the full information price in Lemma 2 might be either increasing or decreasing in the state, but is always greater than marginal cost. In contrast, the equilibrium price under price signaling is always increasing in the state. Together with Equation 2.3 and the fact that profits increase in $\hat{\theta}$, this implies that the price must be larger than the full information price (i.e. at $p^{S,*}$ the denominator of Equation 2.3 must be negative). Hence, signaling always takes the form of inflated prices.

Additionally, I show that there exist no equilibria in which types partially pool on prices by applying the equilibrium refinement.

Proposition 2. *There exists no equilibrium in which the platform in more than one state θ sets a price $p(\theta) = p'$ under the equilibrium refinement.*

Applying the adjusted D1 Criterion in Definition 4 rules out any equilibrium in which types partially pool on prices. The highest type in the pool always has an incentive to deviate. The single-crossing property in prices implies that there exists a small increase in both price and demand for which the highest type prefers to deviate, while lower types do not. Since lower types would not choose such a price, D1 beliefs assign positive probability only to higher types. Then, since D1 beliefs assign higher probability to higher types following an off-path deviation, the user response must increase accordingly. This ensures that the highest type finds the deviation profitable, thereby breaking the pooling equilibrium.

The results in this section provide a novel rationale for platforms that charge high prices, namely to signal their high quality. Many platforms offering "premium" services only charge high prices to demonstrate that they can attract users with higher stand-alone values through their services.[4]

2.4 Price, Message and Fake Profiles as Signals

In this section, I consider a setting where users observe the network size after joining but cannot distinguish fake from real profiles, modeled by the information structure $\mathbb{I} = [0,1] + \mathbb{R}_0^+$. Due to the timing of the game, signals are observed sequentially and beliefs are updated twice. I analyze the following two cases: First, sophisticated users who are aware

[4]For example in the dating industry, platforms such as eHarmony.com or ElitePartner advertise their high quality services in comparison to other dating sites such as Match.com. ElitePartner, a dating site for academics, offers to create an account for free, but to take any action on the platform, users need to sign up for their membership which ranges between 70 Euro/month (6 months contract) to 35 Euro/month (24 month contract). To sign up on ElitePartner users need to certify their academic degrees.

of the possibility that the platform can create fake profiles but are unable to distinguish them ad hoc. Second, naive users who are unaware of the possibility of creating fake profiles, i.e., they simply believe that fake profiles are illegal or impossible to create.

2.4.1 Sophisticated Users

The analysis of sophisticated users, for whom fake profiles are unobservable but feasible, extends the analysis in Section 2.3.1 by introducing an additional signaling instrument. Throughout, I will continue to focus on differentiable separating equilibria as before.

Definition 2. A separating equilibrium consists of a platform's strategy that is a one-to-one mapping from the state to pairs of price and fake profiles $\rho : \Theta \to \mathbb{R}_+ \times \mathbb{R}_+,\ \theta \mapsto (p, \xi)$, users' strategy σ_i^1 and σ_i^2, and beliefs μ_1 and μ_2 such that:

1. For any $(p, \xi) \in \rho(\Theta), \mu_1(\cdot) = \rho^{-1}(p) = \mu_2(\cdot) = \rho^{-1}((p, \xi))$ (Belief Consistency).
2. For any $\theta, \theta' \in \Theta, \pi(\theta, \theta, \rho(\theta)) \geq \pi(\theta, \theta', \rho(\theta'))$ (Incentive Compatibility).

By construction, the platform faces a two-dimensional signaling problem as sophisticated users take both the price and the (expected) number of fake profiles as signal. The optimization problem can be formulated as the platform maximizing its profit given that users are able to infer the true state in the separating equilibrium:

$$\arg\max_{p,\xi} (p - c)n(\theta, \theta, p) - \gamma(\xi)$$

subject to incentive compatibility

$$\pi(\theta, \theta, \rho(\theta)) \geq \pi(\theta, \theta', \rho(\theta')), \forall \theta, \theta' \in \Theta,$$

given by

$$\begin{aligned} &(p(\theta) - c)n(\theta, \theta, p(\theta)) - \gamma(\xi(\theta)) \\ \geq\ &(p(\theta') - c)n(\theta, \theta', p(\theta')) - \gamma\left[\xi(\theta') + (n(\theta', \theta', p(\theta')) - n(\theta, \theta', p(\theta')))\right]. \end{aligned} \tag{IC}$$

A deviation in the equilibrium strategy $\rho(\theta)$ consists of a deviation in price $p(\theta)$ alongside a change in fake profiles $\xi(\theta)$. Although fake profiles bear similarities to the concept of advertising, note that the incentive constraints are different. As users only observe $n + \xi$, a deviation to price-fake profile pair (p', ξ') reveals information about the state. For a type θ to mimic a type θ', the platform instead needs to additionally adjust its fake profiles by the difference in demand when deviating to induce belief θ'.

Turning to analyzing the incentive constraint, note that IC must bind. Setting the IC slack would imply that the platform could decrease the difference between $\xi(\theta)$ and $\xi(\theta')$

and save costs. Rearranging yields

$$\gamma[\xi(\theta') - \xi(\theta)] =$$
$$(p(\theta') - c)n(\theta, \theta', p(\theta')) - (p(\theta) - c)n(\theta, \theta, p(\theta)) - \gamma \left[n(\theta', \theta', p(\theta')) - n(\theta, \theta', p(\theta'))\right]. \tag{IC*}$$

which pins down the fake profile strategy of type θ as a function of the price $p(\theta)$. Note that I restrict fake profiles to be non-negative throughout the analysis. Another possibility is to pin down the pricing strategy as a function of fake profiles. Hence, it is possible to construct a continuum of separating equilibria as separation can be achieved either via the price or fake profiles. Consider the following condition:

Condition 1.

$$(p - c)n_2(\theta, \theta, p) > \gamma n_1(\theta, \theta, p), \forall \theta \in \Theta.$$

Under Condition 1, the signaling benefit from an extra fake profile, the left-hand side of the inequality, outweighs the cost of an extra fake profile, the right-hand side of the inequality.

Proposition 3. *Suppose Condition 1 is fulfilled.*

(i) There always exists a separating equilibrium in which the platform sets zero fake profiles and prices are set to

$$\tau'(\theta) = -\frac{(\tau(\theta) - c)n_2(\theta, \theta, \tau(\theta))}{(\tau(\theta) - c)n_3(\theta, \theta, \tau(\theta)) + n(\theta, \theta, \tau(\theta))}, \tag{2.4}$$

with the initial value condition $p(\underline{\theta}) = p^{FI}(\underline{\theta})$.

(ii) There exists a separating equilibrium in which the platform sets a positive level of fake profiles given by

$$\gamma\xi = (p(\theta) - c)n(\theta, \theta, p(\theta)) - (p(\underline{\theta}) - c)n(\underline{\theta}, \underline{\theta}, p(\underline{\theta})) - \int_{\underline{\theta}}^{\theta} (p(t) - c - \gamma)n_1(t, t, p(t))dt, \tag{2.5}$$

and prices maximize equilibrium profits

$$p(\theta) - c - \gamma = -\frac{n_1(\theta, \theta, p)}{n_{13}(\theta, \theta, p)}, \text{ for } \theta > \underline{\theta}. \tag{2.6}$$

Suppose Condition 1 is violated. Then, there exists a separating equilibrium in which the platform sets zero fake profiles and full information prices in each state.

Suppose Condition 1 holds. Then, there always exists a separating equilibrium in which the platform uses only prices as a signal, as in Proposition 1. Since the equilibrium outcome in Proposition 1 satisfies the equilibrium refinement, so does the equilibrium outcome in Proposition 3. The proof of Proposition 1 shows that there is no price to which a type wants to deviate off-path, which also applies here. Additionally, independent of any belief, the platform does not want to deviate to positive fake profiles, as once users have entered — given any belief — no additional users can enter. Thus, fake profiles are costly but do not increase demand.

There can exist a second equilibrium under additional condition in which the platform creates a positive number of fake profiles. There are two main differences compared to the analysis of advertisement. First, in the IC fake profiles require that the additional term $\gamma\,[n(\theta,\theta',p(\theta'))-n(\theta',\theta',p(\theta'))]$ is present. This reduces, ceteris paribus, the slope of the fake profile function in equilibrium which must be created to ensure incentive-compatibility. Intuitively, when the platform with a low fundamental wants to mimic a platform with a high fundamental, the latter has an advantage of an already larger network size. Second, the price in a model with advertisement would be independent of the costs γ which appear as a mark-up on the right-hand side of Equation 2.4. Therefore, the platform can shift a part of the marginal costs of creating fake profiles to its users.

As in Section 2.3, I apply the equilibrium refinement to show that there exist no (partial) pooling equilibria.

Proposition 4. *There exists no equilibrium in which the platform sets a price p' and fake profiles $\xi(\theta)$ such that $n(\theta,\mu',p')+\xi(\theta)=n(\theta',\mu',p')+\xi(\theta')$ in more than one state.*

2.4.2 Naive Users

This section turns to the analysis of naive users. Users are assumed to be misspecified about fake profiles as they do not take the possibility into account that fake profiles can be created. Therefore, naive users take the network size on a platform at face value, i.e. suppose real profiles are equal to the profiles they see. As users cannot see the network size upfront, they observe the platform's message $m\in\Theta$ and take this as a literal statement about its network size. Given a price p, the platform's message $m\in[\underline{\theta},\overline{\theta}]$ can be interpreted as sending a message about the feasible network size $m\in[n(\underline{\theta},\underline{\theta},p),n(\overline{\theta},\overline{\theta},p)]$.

The notion of naive users is motivated by wrong legal beliefs, as users may believe that fake profiles are simply illegal or impossible to create in practice.[5] A majority of users that is new to those platforms is surprised afterwards about the use of fake profiles (see Section 2.5). In other cases, users form beliefs about these practices in traditional markets, where the use of fake profiles or similar practices is forbidden, and take over their beliefs to online markets or digital platforms. Fake profiles are often legal or the creation of fake profiles is legal as long as firms disclose their use in the terms and conditions. In

[5]See Armstrong and Vickers (2012) who make a similar argument towards naivete with respect to hidden prices.

the terms and conditions, however, the phrases are well-hidden (see Section 2.5). Hence, if users are naive with respect to the possibility of fake profiles, they take any demand that they deem feasible at face value. This approach is combined with the notion of credulity of users. The platform's message about the network size is technically cheap-talk, but users blindly believe in the message as they take the potential network size at face value. In online markets, this message could be the announcement of membership statistics or advertisement about the network size. Ottaviani and Squintani (2006) and Kartik et al. (2007) define the notion of credulous users for cheap-talk games. In the model at hand, the credulity stems from the naivete about the network size.

The equilibrium analysis is greatly simplified due to the fact that users put probability one onto the state that the platform announces (which corresponds to a mass of users on the platform). Hence, all subsequent beliefs in this section are point-beliefs, i.e. the Dirac measure on θ'. To analyze the equilibrium of the game with naive users, I need to define the platform's strategy and equilibrium concept.[6]

Definition 3. The platform's strategy ν is a LSHP (low types separate and high types pool) strategy if, for any price p, there exists a $\tilde{\theta} \in [\underline{\theta}, \overline{\theta}]$ such that:

1. For all $\theta < \tilde{\theta}, \nu(\theta) \in \{m(\theta) | m \in \Theta \setminus \{\overline{\theta}\}\},$ with $\nu(\theta) \neq \nu(\theta') \forall \theta \neq \theta'$.

2. For all $\theta \geq \tilde{\theta}, \nu(\theta) = \overline{\theta}$.

Given a price p, a perceived separating equilibrium consists of

1. A LSHP strategy on messages $\nu(\theta)$.

2. User beliefs $\mu(m) = m$.

3. A fake profile strategy $\xi(m) = n(m, m, p) - n(\theta, m, p)$.

While the platform may separate in some states but pool in others, naive users, however, hold separating beliefs and thus form a point belief after observing the message and perceived network size. The equilibrium is therefore termed a *perceived* separating equilibrium. Since *full* separation is not feasible under low costs, as specified below, there exists a cutoff state: the platform pools on the highest message if the fundamental is above the cutoff and separates if it is below. Note that the fake profile strategy mirrors the difference in real demand and forms part of the equilibrium. This definition is not restrictive, as the platform chooses the fake profile strategy optimally as shown in the following lemma.

Lemma 4. *The platform optimally sets fake profiles equal to* $\xi(\mu^*(m)) = n(\mu^*(m), \mu^*(m), p) - n(\theta^*, \mu(m), p)$ *such that* m *and* ξ *induce the same belief* $\mu^*(m)$.

[6] This definition is based on Kartik (2009) who defines a strategy about a message as a LSHP strategy in the context of a cheap-talk game.

Users hold two relevant beliefs for the platform, once during their entry decision $\mu(\theta|p, m, r) = m$ and once during their exit decision $\mu(\theta|p, m, r, n)$. If the announced message (or announced demand) differs from the actual observed demand on the platform, these two beliefs are not the same. Again, with a slight abuse of notation $n(\theta, m, p)$ denotes the true network size based on the true state θ, the belief $\mu(\cdot) = m$ which is a Dirac measure, and the price. Conversely, $n(m, m, p)$ denotes the believed network size of naive users given belief $\mu(\cdot) = m$. The platform's maximization problem is

$$\underset{\{m\in\Theta, p\in\mathbb{R}_+\}}{\arg\max} \; (p-c)n(\theta, m, p) - \gamma\left(n(m, m, p) - n(\theta, m, p)\right)$$

The first-order conditions with respect to p and m result in

$$p - c = -\frac{n(\theta, m, p)}{n_3(\theta, m, p)} + \gamma\left(\frac{n_3(m, m, p)}{n_3(\theta, m, p)} - 1\right) \tag{2.7}$$

$$(p - c + \gamma)n_2(\theta, m, p) = \gamma\left(n_1(m, m, p) + n_2(m, m, p)\right). \tag{2.8}$$

The first equation determines the optimal price given a chosen message m. If the message is equal to the true state, the optimal price is equal to the full information price. The second equation determines the choice of the optimal message. For a given price, the left-hand side is the marginal benefit of fake profiles, i.e. the increase in users' beliefs, and the right-hand side are the marginal costs. The first-order conditions are only applicable for $m \geq \theta$. If $m < \theta$, the platform does not set any fake profiles as the number of fake profiles is bounded away from zero.

Next, I determine the cutoff, i.e., the state at which the platform first chooses the highest message. To solve for the cutoff type $\tilde{\theta} < \overline{\theta}$, I examine the indifference condition. Let the profit of the platform in the indifferent state $\tilde{\theta}$ be

$$\pi(\tilde{\theta}, m, \xi(\tilde{\theta}, m), p) \equiv \pi(\tilde{\theta}, m),$$

which solves $\overline{\theta} = \arg\max_{m\in\Theta} \pi(\tilde{\theta}, m)$:

$$\left(-\frac{n(\tilde{\theta}, \overline{\theta}, p)}{n_3(\tilde{\theta}, \overline{\theta}, p)} + \gamma\frac{n_3(\overline{\theta}, \overline{\theta}, p)}{n_3(\tilde{\theta}, \overline{\theta}, p)}\right) n_2(\tilde{\theta}, \overline{\theta}, p) = \gamma\left(n_1(\overline{\theta}, \overline{\theta}, p) + n_2(\overline{\theta}, \overline{\theta}, p)\right). \tag{2.9}$$

Rewriting the equation gives the next lemma.

Lemma 5. *The indifferent type $\tilde{\theta}$ is the solution to*

$$\beta(1 - F_{\tilde{\theta}}(r(\overline{\theta}, p)) = \gamma, \tag{2.10}$$

which has a unique solution and solves for $\tilde{\theta} \in \Theta$ if

$$\underline{\gamma} \equiv \beta(1 - F_{\underline{\theta}}(r(\overline{\theta}, p)) \leq \gamma \leq \beta(1 - F_{\overline{\theta}}(r(\overline{\theta}, p)) \equiv \overline{\gamma}$$

The equilibrium is characterized in the following proposition.

Proposition 5. *The equilibrium with naive users is characterized as follows:*

(i) If $\underline{\gamma} \leq \gamma \leq \overline{\gamma}$. The indifferent type solves Equation 2.10. Types $\theta < \tilde{\theta}$ separate with $\nu(\theta) > \theta$ and types $\theta \geq \tilde{\theta}$ pool on $\nu(\theta) = \overline{\theta}$.

(ii) If $\gamma < \underline{\gamma}$, all types pool on $\nu(\theta) = \overline{\theta}$.

(iii) If $\gamma > \overline{\gamma}$, each type chooses $\nu(\theta) = \theta$.

Given message m, users believe $\mu(m) = m$, equilibrium prices are given by Equation 2.7 and the number of fake profiles is $\xi(m) = n(m, m, p) - n(\theta, m, p)$.

Suppose that γ lies within the specified upper and lower bounds. Then, by Lemma 5, Equation 2.10 has a unique solution. All types above the indifferent type pool on the highest message, while the types below choose m according to

$$\beta(1 - F_{\theta}(r(m, p)) = \gamma.$$

Setting $m = \theta$, i.e., every type reveals its true type to the naïve users, does not satisfy the equation, as the marginal benefit from lying upwards exceeds the cost. Hence, every type uses an inflated message $m > \theta$, such that complete separation (from the platform's perspective) is not possible. Due to the bounded state space, the highest type runs out of claims to make, and therefore higher types pool on the highest possible message. All types except the highest set a higher price than under full information, and even the lowest type creates fake profiles. The number of fake profiles increases up to the indifferent type and decreases afterward; only the highest type creates no fake profiles. To exploit consumers' naivete, a bound is imposed on the feasible strategies. In contrast to the fake profile strategy in Section 2.4.1, which was not bounded, consumer naivete makes the restriction of the state space binding, thereby influencing the strategy space.

Corollary 5. *Suppose $\underline{\gamma} \leq \gamma \leq \overline{\gamma}$. Then $\tilde{\theta}$, is increasing in γ and decreasing in β.*

In other words, the larger the network effects, the lower the indifferent type, as the benefit from creating fake profiles increases. Conversely, if the marginal costs increase, the indifferent type rises.

Suppose that $\gamma \leq \underline{\gamma}$. Then the lowest type already finds it optimal to send the message $m = \overline{\theta}$, and thus so do all types above. Therefore, all types pool on the message and the perceived network size. This implies that the lowest type creates fake profiles, and moreover, the number of fake profiles is decreasing in type. That is, the lowest type creates the most, and the highest creates none. Lastly, if $\gamma \geq \overline{\gamma}$. Then all types find it optimal to send a message equal to their true state. They create no fake profiles and set their full information prices.

From a welfare perspective, participation on the platform is distorted if $\gamma \leq \overline{\gamma}$. The new indifferent user does not benefit from the use of fake profiles, as they pay an inflated price for a non-existent network size. Users, however, who would have entered in the game without fake profiles and full information might benefit indirectly from fake profiles if they value network effects via β sufficiently. Due to excessive entry, the real network size on the platform increases, which may offset the higher prices for some users.

2.4.3 Regulation

This section analyzes possible remedies for and regulation to deal with the use of fake profiles. More specifically, I consider a ban of fake profiles, labeling fake profiles on the platform, mandatory disclosure of fake profiles upfront. It is assumed that the regulation is publicly known and users are educated and informed about the policy. The analysis considers the first three policies first and will show that all three will lead to the same unique market outcome. Lastly, it will be shown that educating users about the use of fake profiles is insufficient to prevent the use of fake profiles.

Banning Fake Profiles How does a ban of fake profiles impact the market outcome? Suppose the ban of fake profiles is public and users are informed about the policy. Sophisticated users will deduce that whenever they join a platform, they will observe the real network size. Hence, after joining sophisticated users are in a subgame of complete information in which the state is known. The unique equilibrium is summarized in Proposition 6.

Labeling Fake Profiles Suppose through labeling fake profiles, users can perfectly identify fake profiles and determine the real network size. In this case, labels must be clear, obvious and understood by users. Again, in the last period sophisticated users face a subgame of complete information (see Proposition 6). If fake profiles cannot serve as a signal due to perfect identification, no fake profiles are used by the platform. This follows directly from the fact that fake profiles are costly, but do not yield a positive benefit to the platform at this stage. Naive users behave as above.

Mandatory Disclosure Lastly, consider that platform must mandatory disclose their use of fake profiles upfront (either stating their use or abstention). Under signaling with price and fake profiles, or price only, the profit of a platform is always lower than the full information profit which is first-best as signaling is costly. In the presence of mandatory disclosure, the platform can choose to refrain from fake profiles credibly, which induces sophisticated users to deduce, again, that they will observe the real network size on the platform. As will be shown in the following proposition, this will enable the platform to achieve its full information profit.

Analysis All policies result in perfect knowledge of users about the real network size on the platform after joining, such that fake profiles cannot influence their perception. Users are sophisticated and take the price as the only costly signal. Hence, the platform maximizes its full information profit

$$\max_p (p-c)n(\theta,\theta,p), \text{ subject to } n(\theta,\mu_2(\cdot)=\theta,p) \leq n(\theta,\mu_1(\cdot)=\hat{\theta},p).$$

As users' participation decisions are made after observing the first-period price, the constraint imposes an upper bound on demand in the last period. The optimal prices are given by the first-order condition

$$n(\theta,\theta,p)+(p-c)\frac{\partial n(\theta,\theta,p)}{\partial p}=0 \tag{2.11}$$

resulting in $p = p^{FI}(\theta)$, the full information benchmark price given that $\theta \leq \hat{\theta}$.

To see that $p^{FI}(\theta)$ is an incentive-compatible separating strategy, suppose that the platform in state θ sets a price $p^{FI}(\theta') < p^{FI}(\theta)$ for $\theta' < \theta$. This influences demand in two ways: first, a price decrease leads to more demand holding all else constant and second, a price decrease influences the believed state $\hat{\theta}$ and leads to a lower expected state. This in turn, decreases demand all else constant. Suppose first that demand overall increases and more user enter than in equilibrium. Then, in the last period users observe the realized demand given price and belief $(p^{FI}(\theta'),\theta')$ and the true state θ. As too many users entered given belief θ', users exit again such that

$$n(\theta,\theta,p^{FI}(\theta') = 1 - F_\theta(p^{FI}(\theta') - \beta n(\theta,\theta,p^{FI}(\theta'))).$$

As the price and realized demand $n(\theta,\theta,p^{FI}(\theta'))$ are not profit maximizing in state θ, the platform does not face a profitable deviation. A similar argument can be constructed if demand overall decreases. Then, following the deviating price $p^{FI}(\theta')$ too few users join the platform than optimal.

Proposition 6. *Suppose the government regulates platforms by either banning fake profiles, forcing them to label fake profiles, or mandatory disclosing fake profiles. If users are aware of these policies, there exists a unique equilibrium that is separating and first-best. In equilibrium, the platform sets the full information price and zero fake profiles are used.*

This result stresses the importance of observing the network size before paying the membership fee. Compared to Section 2.3.1 in which users paid the membership fee without observing the network size, the platform can increase its profit by offering a free-trial period before collecting the (one-time) membership fee.[7] This "free-trial" period is desirable from a platform's perspective as full information prices are incentive compatible. The platform does not incur a loss in profit due to wasteful signaling. In equilibrium, both users and the platform are better off compared to the Section 2.3.1. In comparison to the signaling literature, the learning-from-demand-stage resembles repeat purchases for regular products or a full warranty/money-back guarantee.

The effect of educating naive users about the possibility of fake profiles are ambiguous. Users and the platform might be both worse off. If naive users are educated and no other action on fake profiles are taken, the platform might still use fake profiles for signaling as in Theorem 3. Due to the consumer sophistication, the platform still needs to engage in costly signaling. Comparing the equilibrium outcome in Theorem 3 and 5, the platform makes losses when moving from the latter to the first. With naive users, the platform makes higher profits than under full information, whereas with sophisticated users the platform makes lower profit than under full information. The effect on users depends on the prices and network effects. Naive users benefit from an increase in the real network size compared to the equilibrium in Theorem 3. Additionally, prices might also increase compared to the latter equilibrium.

2.5 Discussion

Convincing Users to Upgrade into a Premium Subscription The Dating platform "Match.com" presumably utilized third-party fake profiles to persuade (male) users into a paid subscription.[8] Following the model, users are able to sign up to the platform for free. Initially, undecided users, who did not pay for the membership, received emails from potentially interested users. In the stylized version, users are assumed to view the

[7] The model abstracts from discounting between periods. Otherwise, the firm must be sufficiently patient or the free-trial period must be sufficiently short.

[8] `https://www.ftc.gov/news-events/press-releases/2019/09/ftc-sues-owner-online-dating-service-matchcom-using-fake-love`

total perceived network size $\tilde{n}$. To interact with the other users, they needed to upgrade their free trial. The platform's network size included fake profiles. The platform allegedly used the third-party fake profiles ξ to direct messages towards non-paying users which lead them to upgrade to a premium membership and pay p. Due to the platform's intentional use of fake profiles, i.e. identifying the third-party fake profiles, directing those to non-paying users but keeping those away from paying members, it is plausible to assume that the platform incurred small (effort) costs γ.

Manipulating the Network Size: Wash Trades One of the largest cryptocurrency exchange platforms ("Binance") is under investigation by the SEC for "manipulative trading that artificially inflated the platform's trading volume". They engaged in so-called Wash trading. More precisely, another associated company ("Sigma Chain") owned by the same entity ("Zhao") as the crypto exchange platform manipulated the platform's trading volume by selling and buying the same financial assets, therefore artificially inflating the platform's volume.[9] Furthermore, the U.S. based affiliate of "Binance" called "BAM Trading Services" is accused of misleading investors about non-existent trading controls on Binance.US. Wash trading is prohibited in offline (financial) markets, e.g. in the US by the Commodity Exchange Act.[10] For example, the Intercontinental Exchange (ICE) takes measures to prevent self-trade to comply with regulations.[11]

In this application, the platform is a cryptocurrency trading platform and its users are potential investors both buying and selling assets on the platform. Network effects take the form of caring for liquidity. A platform with a large network, i.e. a high trading volume, has more liquid assets and is more credible. Fake profiles are financial assets that are self-traded by the platform and hence inflate the network size.

Manipulating the Network Size: Dating Platforms Other dating platforms use company-created fake profiles; a list of several dating sites using this practice has been published by the Verbraucherzentrale Bayern (Center for Consumer Advise Bavaria) in Germany. These platforms employ paid workers to create profiles, and interact with users on the platform, giving them the impression of a real contact.[12] It is not commonly known that platforms themselves create fake users to possibly stimulate demand, although it is legal to do so as long as it is mentioned in the terms and conditions. There are companies that specialize in providing employees as chat moderators to these platforms.[13] These

[9] `https://www.sec.gov/news/press-release/2023-101`

[10] `https://www.law.cornell.edu/uscode/text/7/chapter-1`

[11] `https://www.theice.com/publicdocs/futures/IFEU_Self_Trade_Prevention_FAQ.pdf`

[12] See `https://www.verbraucherzentrale.de/wissen/digitale-welt/onlinedienste/onlinedating-auf-diesen-portalen-flirten-fakeprofile-21848`, last visited 01.09.2020.

[13] For example, Cloudworkers or Agentur da Chatdeife are companies that employ freelancers to work for and on one or more social-community platforms. See also `https://www.spiegel.de/wirtschaft/service/singleboersen-ein-moderator-von-fake-profilen-spricht-ueber-`

chat moderators set up fake profiles and engage in conversations with the users of the platform pretending to be a real profile.

Furthermore, the UK Consumer and Markets Authority (CMA) confirms in its report about the online dating industry that dating platforms may use "pseudo profiles" or provider-generated profiles that could possibly mislead consumers. The CMA states that if these fake profiles are not disclosed as such, it may be in breach with the "Consumer Protection from Unfair Trading Regulations".[14] In another industry report issued by the Australian Competition and Consumer Commission (ACCC), the ACCC acknowledges that fake profiles generated by providers exist, but stress that this issue lies beyond the scope of their investigation mandate.[15] This shows that the use of fake profiles might be more common than initially expected and might not be restricted to the examples given above.

Evidence that chat bots might have been used by dating platforms exists for the dating site “Ashley Madison”. Ashley Madison was subject to a large data leak by hackers.[16] The dating site used “chat hostesses” before 2011 to engage men, which coincides with the notion of fake profiles in this context. After 2011, however, it is reported that they stopped employing “chat hostesses”. Instead, the dating platform allegedly used chat bots to deceive users to spend money on the platform. Although one might think that chats bots are easier to be identified, this might not have been the case. Users seem to have spent a reasonable amount of money on communicating with chat bots.

Lastly, there is evidence on dating platforms that use methods to create a similar effect as with platform-generated fake profiles. The CMA investigated the case of Venntro Media Group Ltd, a company that operates several dating sites. To inflate the network size on their dating sites, Venntro cross-registered their members on various sites and not only the site they originally signed up for.[17]

Launching Strategy for Start-Ups Upon launching a new platform, founders often generate artificial demand (or supply) to onboard producers (or consumers). This practice is documented in the business and management literature, e.g. Schirrmacher et al. (2017) or Reillier and Reillier (2017). Evans and Schmalensee (2016) describe the practice as

seinen-job-a-1113937.html. and https://www.ndr.de/fernsehen/sendungen/panorama_die_reporter/Undercover-als-Chatschreiberin-Abzocke-Flirtportal,sendung1098906.html for an interview (in German) and https://www.marieclaire.fr/,dating-assistant,750821.asp for an article (in French).

[14]See https://assets.publishing.service.gov.uk/media/5b114a8040f0b634abe911e7/compliance_statement.pdf.

[15]See https://www.accc.gov.au/system/files/927_ICPEN%20Dating%20Industry%20Report_D09.pdf.

[16]See https://financialpost.com/fp-tech-desk/inside-ashley-madison-calls-from-crying-spouses-fake-profiles-and-the-hack-that-changed-everything?__lsa=b245-a155.

[17]See https://www.gov.uk/government/news/online-dating-giant-vows-clearer-path-to-love.

"self-supply". In case studies by Schirrmacher et al. (2017) some platforms self-supplied in the beginning at launch to influence participants' beliefs, whereas one platform simulated fake demand.

2.6 Conclusion

Especially on Dating platforms the use of fake profiles is heavily relied upon. Suggestive evidence from a data leak of the platform "Ashley Madison" shows that fake profiles were used excessively. Most of the female users were in fact fake profiles. The data, however, included credit card transactions (mostly from men) indicating that many users spend a lot of money on the platform even though the chance of encountering a real women was surprisingly low.

Economic papers exploring the regulation of platform markets are scarce, although policy papers such as Fletcher et al. (2021) investigate common issues on platform that may need to be regulated. For fake profiles, one suggested policy is banning fake profiles. German cases suggest, however, that the disclosure cannot be hidden in terms and conditions. Similarly, one could consider mandatory disclosure policies. Voluntary and mandatory disclosure has been discussed by scholars such as Grossman (1981), Mathios (2000), or Fishman and Hagerty (2003).

In a classical model with rational users and voluntary disclosure all but the lowest type should disclose their type and state that they would not use fake profiles. In my model a platform would like to commit to refrain from using fake profiles with sophisticated users as they are costly, which would indicate that if users are sophisticated voluntary disclosure on fake profiles should be observed in online markets. Instead, their actual use is mentioned in the terms and conditions, and consumer protection and competition authorities try to inform unknowing consumers about these. In contrast, there is no evidence of information campaigns or initiatives of firms committing not to use fake profiles. Such voluntary disclosure might fail as the presence of naive users eliminates the incentives to voluntary disclose the own type.

Combining suggestive evidence and the failure to observe voluntary disclosure in these markets suggests that users are mainly naive. This speaks in favor of consumer protection policies against practices that influence network effects such as a ban of fake profiles or mandatory disclosure.

A Appendix: Definition of Equilibrium Refinement

Definition 4. Equilibrium Refinement for Sophisticated Users
Denote the equilibrium strategy profile by $\Sigma = ((p^*, m^*), r^*(p^*, m^*))$, where $r(\cdot)$ denotes the equilibrium cutoff mapping. The equilibrium profit of a platform of type θ is $\pi^*(\theta, \Sigma)$.

For a given price p, an arbitrary non-empty subset of sender type space $\tilde{\Theta} \subseteq \Theta$, and a non-empty subset of the other receivers action spaces $\tilde{\mathbf{Y}}_{-i}$ let

$$\mathrm{BR}_i(\tilde{\Theta}, \tilde{\mathbf{Y}}_{-i}) = \cup_{\rho_i \sim \Delta(\tilde{\Theta} \times \tilde{\mathbf{Y}}_{-i})} \arg\max_{y_i \in [0,1]} \mathbb{E}_{(\theta, \mathbf{y}_{-i}) \sim \rho_i}[u_i(\theta, y_i, \mathbf{y}_{-i}, p)] \forall i$$

be the set of user i's best responses to p for some belief ρ_i over sender type and the other receivers action pairs with support contained in $\tilde{\Theta} \times \tilde{\mathbf{Y}}_{-i}$. For an arbitrary non-empty subset of sender type space $\tilde{\Theta} \subseteq \Theta$ and $k \in \{0, 1, 2, ...\}$ let

$$Y_i^k(\tilde{\Theta}) = \mathrm{BR}_i(\tilde{\Theta}, \tilde{\mathbf{Y}}_{-i}^{k-1}(\tilde{\Theta})), \text{ and } Y_i^\infty(\tilde{\Theta}) = \cap_{k \in \{0,1,2,...\}} Y_i^k(\tilde{\Theta}) \forall i$$

be the set of rationalizable actions given $\tilde{\Theta}$ for receiver i. Denote by $\mathcal{R}^\infty(\tilde{\Theta}, p)$ the set of rationalizable receiver action profiles for given p and $\tilde{\Theta}$.

For a given out-of equilibrium price p and for each type θ, find all rationalizable action profiles $\alpha \in \mathcal{R}^\infty(\Theta, p)$ by users that would cause θ to deviate from equilibrium. For $\theta \in \Theta$, p, and equilibrium profile Σ,

$$D_\theta = \{\alpha \in \mathcal{R}^\infty(\Theta, p) : \pi^*(\theta, \Sigma) < \mathbb{E}_{r \sim \alpha} \pi(\theta, p, r)\},$$

is the set of receiver rationalizable actions for which type θ is strictly better-off deviating towards p, and

$$D_\theta^0 = \{\alpha \in \mathcal{R}^\infty(\Theta, p) : \pi^*(\theta, \Sigma) = \mathbb{E}_{r \sim \alpha} \pi(\theta, p, r)\},$$

is the set of receiver rationalizable actions for which type θ is indifferent between deviating towards p and setting equilibrium price p^*. If for some type θ there exists another type θ' such that

$$D_\theta \cup D_\theta^0 \subset D_{\theta'},$$

then (θ, p) may be pruned from the game. The set of types that cannot be deleted is denoted by $\Theta^*(p)$. A PBE violates D1 if there exists a type and action (θ, p) such that

$$\min_{\alpha \in \mathcal{R}^\infty(\Theta^*(p), p)} \pi(\theta, p, r) > \pi^*(\theta, \Sigma) \text{ for some } \theta \in \Theta^*(p). \tag{D1}$$

Discussion Ramey (1996) shows that under the following assumptions the unique D1 equilibrium is separating. The set of types is given by non-degenerate interval $[\underline{\theta}, \overline{\theta}]$, where prior beliefs are given by a continuous probability distribution $\mu(\theta)$ with full support. Signals are $p \in \mathbb{R}^k$ and the (single) receiver's response r is chosen from the real line. Payoff functions are continuously differentiable; the sender's payoff π increases in the receiver's response. The receiver's utility function u is strictly quasi-concave in its action r for each signal and type. The receiver's payoff is maximized in r by $r^*(\theta, p)$, which is strictly increasing in θ. Furthermore, $r^*(\theta, p)$ is uniformly bounded above and for $k = 1, .., n : \lim_{p_k \to +/-\infty} \pi(\theta, r, p) = -\infty$. Enger's Incentive Montonicity Condition holds for the k signals (weaker condition than the multi-dimensional single-crossing property).

In the model at hand the assumptions on the receiver's payoff functions and actions are not fulfilled. A receiver's response given price p is binary and depends on the belief over the other receivers' actions. The sender, however, does not care about the action of a single receiver, but cares about the aggregate action taken by the receivers. The receivers' payoffs are instead quasi-concave (linear) in the aggregate response. Additionally, the aggregate response $n^*(\theta, p)$ is strictly increasing in θ due to the assumptions on F_θ.

B Appendix: Omitted Proofs

Recall that stand-alone values r are distributed according to the cumulative distribution function $F_\theta(r)$, with associated density $f_\theta(r)$. Denote the derivative of the density with respect to r by $f'_\theta(r)$. Let $F^{1,\theta}(r)$ and $f^{1,\theta}(r)$ represent the first derivatives of the distribution function and the density, respectively, with respect to θ.

Proof of Lemma 1 First, consider the case of complete information, in which the assumption that users hold a common belief upon observing (p, m) is trivially fulfilled. Users play sequentially rationalizable strategies, i.e. they play a best-response to a symmetric cutoff of other. To establish the unique cutoff $r_c \in [\underline{r}, \overline{r}] \subseteq [-\infty, \infty]$, I will iterate on the best-responses of users once from above starting at $r_0 = \overline{r}$ and once from below starting at $r_0 = \underline{r}$.

Step (i) Iteration starting from $r_0 = \overline{r}$.

Consider the best response of an agent i given an arbitrary state θ, price and message pair (p, m) and the action profile of the other agents. The first iteration given the symmetric cutoff $r_0 = \overline{r}$ yields

$$BR_i^1(\{\theta\}, r_0 = \overline{r}) = \begin{cases} 1 & \text{if } r_i \geq p \\ 0 & \text{if } r_i < p. \end{cases}$$

In the first iteration, agents with a reservation value of $r_1 \equiv p$ or higher will always enter even if no one enters the platform (independent of their beliefs). Iterated elimination of not best responses yields a cutoff value of r_i in the $i+1$'th iteration given by

$$r_{i+1} = p - \beta(1 - F_\theta(r_i))$$

This sequence is bounded by $r \in [\beta - p, p]$ and strictly decreasing by the assumptions made in Section 2.2. Hence, the sequence converges to its limit

$$r_c = \lim_{i\to\infty} r_{i+1} = \lim_{i\to\infty} p - \beta(1 - F_\theta(r_i)) = p - \beta(1 - \lim_{i\to\infty} F_\theta(r_i)) = p - \beta(1 - F_\theta(r_c)).$$

The last inequality follows from the fact that the probability function is assumed to be continuous. Then, the condition

$$\overline{r}_c = p - \beta(1 - F_\theta(\overline{r}_c))$$

Step (ii) Iteration starting from $r_0 = \underline{r}$.

The first iteration given the symmetric cutoff $r_0 = \underline{r}$ yields

$$BR_i^1(\{\theta\}, r_0 = \underline{r}) = \begin{cases} 1 & \text{if } r_i \geq p - \beta \\ 0 & \text{if } r_i < p - \beta. \end{cases}$$

In the first iteration, agents with a stand-alone value below $r_1 \equiv p - \beta$ will never enter even if all others join the platform (independent of their beliefs). The cutoff value of r_i in the $i+1$'th iteration is given by

$$r_{i+1} = p - \beta(1 - F_\theta(r_i)),$$

as the sequence is bounded by $r \in [\beta - p, p]$ and strictly increasing,it converges to its limit

$$\underline{r}_c = p - \beta(1 - F_\theta(\overline{r}_c)).$$

Step (iii) Show that $\overline{r}_c$ and $\underline{r}_c$ coincide.

Given Assumption 3, for any $p \in \mathbb{R}_+$ there exists one and only one solution to the equation

$$r + \beta(1 - F_\theta(r)) = p, \tag{2.12}$$

as r is increasing in r with slope one, whereas $\beta(1 - F_\theta(r))$ is decreasing in r with slope smaller than one. Hence, the left-hand side is strictly increasing in r. Then, $r_c \equiv \{r \in$

$[\underline{r}, \overline{r}] : r+\beta(1-F_\theta(r)) = p\}$ characterizes the unique cutoff which is the unique sequentially rationalize user profile given the action-pair (p, m) by the platform.

Second, consider the case of incomplete information under the assumption that users hold a common belief upon observing (p, m). Given the equilibrium definition of PBE, there are two relevant types of equilibria — separating and pooling equilibria. In both equilibria, the common belief assumption is fulfilled again. In a separating equilibrium, users hold point-beliefs after observing (p, m), whereas users' Bayesian update in a pooling equilibrium after observing (p, m) is equal to their common (full support) prior. The proof proceeds as follows. First, I will show that if users play cutoff strategies, there exists a unique cutoff. Second, I will show that users will play a cutoff strategy in any equilibrium.

Step (i) Unique cutoff.

Suppose there exist two cutoffs defined by

$$\underline{r}_c + \beta \int_\Theta (1 - F_\theta(\underline{r}_c)) d\mu(\theta|p, m, \underline{r}_c) = p, \text{where} \tag{2.13}$$

$$\mu(\theta|p, m, \underline{r}_c) = \frac{\mu(\theta|p.m) f_\theta(\underline{r}_c|\theta)}{\int_{\tilde{\theta}\in\Theta} \mu(\tilde{\theta}|p, m) f_\theta(\underline{r}_c|\tilde{\theta}) d\tilde{\theta}}, \tag{2.14}$$

and

$$\overline{r}_c + \beta \int_\Theta (1 - F_\theta(\overline{r}_c)) d\mu(\theta|p, m, \overline{r}_c) = p, \text{where} \tag{2.15}$$

$$\mu(\theta|p, m, \overline{r}_c) = \frac{\mu(\theta|p.m) f_\theta(\overline{r}_c|\theta)}{\int_{\tilde{\theta}\in\Theta} \mu(\tilde{\theta}|p, m) f_\theta(\overline{r}_c|\tilde{\theta}) d\tilde{\theta}}. \tag{2.16}$$

For the sake of contradiction, suppose that the cutoff differ, e.g. $\underline{r}_c < \overline{r}_c$. Denote $X(r_1, r_2) = \int_\Theta (1 - F_\theta(r_1)) d\mu(\theta|r_2)$.

Lemma 6. *$X(r_1, r_2)$ is strictly decreasing in r_1 and (weakly) increasing in r_2.*

The first part follows directly from the fact that $1 - F_\theta(r_1)$ is decreasing in r_1 for all θ and hence, $\int_\Theta (1 - F_\theta(r_1)) d\mu(\theta|r_2)$ is decreasing in r_1 as well holding $\mu(\theta|r_2)$ fixed.

For the second part, note that for $r' > r''$, $\mu(\theta|p, m, r')$ has first-order stochastic dominance over $\mu(\theta|p, m, r'')$ due Assumption 3. Assumption 3 states that the family of densities $\{f_\theta(r) \equiv f(r|\theta)\}$ is assumed to have the monotone likelihood ratio property. By Milgrom (1981) (Proposition 2) a family of densities has the MLRP iff $r' > r''$ implies that r' is more favorable than r'' meaning that $\mu(\cdot|r')$ dominates $\mu(\cdot|r'')$.

Recall that $1 - F_\theta(r_1)$ is bounded by $[0, 1]$ and is strictly monotone increasing in θ by Assumption 2 and 3. Then, $\int_\Theta (1 - F_\theta(r_1)) d\mu(\theta|r') \geq \int_\Theta (1 - F_\theta(r_1)) d\mu(\theta|r'')$ holds as r' is more favorable than r'', where it holds with equality whenever $1 - F_\theta(r_1) \in \{0, 1\}$.

In equilibrium, both the left-hand side of Equation 2.13 and 2.15 must be equal to p. Therefore, $\underline{r}_c + \beta X(= \overline{r}_c + \beta \int_\Theta (1 - F_\theta(\overline{r}_c)) d\mu(\theta|p, m, \overline{r}_c)$ must hold with $\underline{r}_c < \overline{r}_c$. Note that in equilibrium, both the left-hand side of Equation 2.13 and 2.15 must be equal to p. Therefore, $\underline{r}_c + \beta X(= \overline{r}_c + \beta \int_\Theta (1 - F_\theta(\overline{r}_c))d\mu(\theta|p, m, \overline{r}_c)$ must hold with $\underline{r}_c < \overline{r}_c$. Note that In equilibrium, both the left-hand side of Equation 2.13 and 2.15 must be equal to p. Therefore, $\underline{r}_c + \beta X(\underline{r}_c, \underline{r}_c) = \overline{r}_c + \beta X(\overline{r}_c, \overline{r}_c)$ needs hold with $\underline{r}_c < \overline{r}_c$; however,

$$\underline{r}_c + \beta X(\underline{r}_c, \underline{r}_c) < \overline{r}_c + \beta X(\overline{r}_c, \underline{r}_c) \leq \overline{r}_c + \beta X(\overline{r}_c, \overline{r}_c),$$

where the first inequality follows from Assumption 4 for any given θ and the second inequality follows from the lemma above. This contradicts the initial assumption, thus, $\underline{r}_c = \overline{r}_c$.

Step (ii) Users play cutoff strategies in any equilibrium.

Define $r_{\inf} \equiv \inf\{r_i : u(r_i, p) \geq 0\}$ to be the user with the lowest r_i of the set of users that have a non-negative utility from joining the platform. Similarly, let $r_{\sup} \equiv \sup\{r_i : u(r_i, p) \leq 0\}$ be the users with the largest r_i of the set of users that have a negative or zero utility from joining the platform.

Therefore, it needs to hold that $r_{\sup} + \beta X(r_{\sup}, r_{\sup}) \leq p \leq r_{\inf} + \beta X(r_{\inf}, r_{\inf})$ for $r_{\inf} > r_{\sup}$ by definition. Imposing that $X(r_1, r_2)$ *strictly* increases in r_2, I can use the previous argument from Step (i) to show that there agents play cutoff strategies. As long as $r_{\inf}, r_{\sup} \in (\underline{r}, \overline{r})$ and hence $X(r_1, r_2)$ *strictly* increases in r_2, the following holds

$$r_{\sup} + \beta X(r_{\sup}, r_{\sup}) < r_{\sup} + \beta X(r_{\sup}, r_{\inf}) < r_{\inf} + \beta X(r_{\inf}, r_{\inf})$$

Therefore, $r_{\inf} = r_{\sup}$ and users follow a cutoff strategy.

Then, the condition

$$r_c = p - \beta \int_\Theta (1 - F_\theta(r_c))d\mu(\theta|p, m, r_c), \tag{2.17}$$

characterizes the unique cutoff. Note that if $\mu(\theta|r_c) = \delta_{\tilde{\theta}}$, i.e. the belief concentrates on state θ with probability 1 (a.s.), this condition is the same as under complete information. Observe that

Remark. If belief $\mu(\theta|p, m, r)$ has first-order stochastic dominance over a belief $\mu'(\theta|p, m, r)$, $\int_\Theta (1 - F_\theta(r_c))\mu(\theta|r_c)d\theta$ increases and therefore, the cutoff r_c decreases.

The result follows directly from the Equation 2.17 and Lemma 6. □

Proof Lemma 2 Under complete information, equilibrium demand is determined by the unique solution to

$$n^* = Pr(r + \beta n^* - p \geq 0) = 1 - F_\theta(p - \beta n^*),$$
$$\Leftrightarrow n^* = 1 - F_\theta(r^*)$$

given that Assumption 4. It is possible to rewrite this condition as $G(n^*; p) = 1 - F_\theta(p - \beta n^*) - n^*$. The implicit function theorem implies that a function g exists such that $n^* = g(p)$. Implicit differentiation yields

$$\begin{aligned} -f_\theta(p-\beta n^*) + \beta f_\theta(p-\beta n^*)\frac{\partial n}{\partial p} - \frac{\partial n}{\partial p} &= 0 \\ \frac{\partial n}{\partial p} &= \frac{-f_\theta(p-\beta n^*)}{1-\beta f_\theta(p-\beta n^*)}. \end{aligned}$$

The platform faces the optimization problem $\max_p (p-c)n(p)$ and yields

$$1 - F_\theta(r^*) + (p-c)\frac{-f_\theta(r^*)}{1-\beta f_\theta(r^*)} = 0,$$

which can be rewritten as

$$p - c = \underbrace{\frac{1-F_\theta(r^*)}{f_\theta(r^*)}}_{\eta(\theta,p)} \underbrace{(1-\beta f_\theta(r^*))}_{>0},$$

where $\eta(\theta, p)$ is the users price elasticity. Given Assumption 2 the hazard rate defined by $\lambda \equiv \frac{1}{\eta}$ is decreasing. Thus, the first-order condition solves for a unique price $p^*(\theta)$. The second-order condition is

$$-\left[2\frac{f_\theta(r^*)}{1-\beta f_\theta(r^*)} + (p-c)\frac{f'_\theta(r^*)}{(1-\beta f_\theta(r^*))^3}\right].$$

At p^* it holds that

$$\begin{aligned} &-\frac{1}{1-\beta f_\theta(r^*)}\left[2f_\theta(r^*) + f'_\theta(r^*)\frac{1-F_\theta(r^*)}{f_\theta(r^*)(1-\beta f_\theta(r^*))}\right] \\ \Leftrightarrow &-\frac{1}{1-\beta f_\theta(r^*)}\left[\frac{2(1-\beta f_\theta(r^*))[f_\theta(r^*)]^2 + f'_\theta(r^*)(1-F_\theta(r^*))}{f_\theta(r^*)}\right] < 0 \end{aligned}$$

Note that the term in rectangular brackets is positive if $2(1-\beta f_\theta(r^*)) \geq 1$ which holds if $1/2 \geq \beta f_\theta(r^*)$. The denominator is always positive, however, the numerator must also be positive due to the assumption that the hazard rate is increasing in r given $2(1-\beta f_\theta(r^*)) \geq 1$. To see this, take the first derivative of the hazard rate with respect to r

$$\lambda'(r) = \frac{[f_\theta(r^*)]^2 + f'_\theta(r^*)(1-F_\theta(r^*))}{[1-F_\theta(r^*)]^2} > 0, \text{ by Assumption 1.}$$

Lastly, to show that the equilibrium price can be increasing/decreasing or constant in state θ, note first that the hazard rate $\lambda(r, \theta)$ is strictly decreasing in θ by the MLRP property. It follows that

$$\frac{\partial \lambda(r,\theta)}{\partial \theta} = \frac{f^{1,\theta}(r)(1-F_\theta(r)) + f_\theta(r)F^{1,\theta}(r)}{(1-F_\theta)^2} < 0,$$

which allows to bound $f^{1,\theta}(r) < \frac{f_\theta(r)}{1-F_\theta(r)}(-F^{1,\theta}(r))$.

Taking the derivative of the first-order condition with respect to θ yields

$$\begin{aligned}\frac{\partial p}{\partial \theta} &= \frac{-\beta f^{1,\theta}(r) f_\theta(r)(1-F_\theta(r)) - (f^{1,\theta}(r)(1-F_\theta(r)) + f_\theta(r)F^{1,\theta}(r))(1-\beta f_\theta(r))}{f_\theta(r)^2} \\ &= \frac{-f^{1,\theta}(r)(1-F_\theta(r)) - f_\theta(r)F^{1,\theta}(r)(1-\beta f_\theta(r))}{f_\theta(r)^2} \\ &> 0\end{aligned}$$

if $f^{1,\theta}(r^{FI}) < 0$.

Proof of Lemma 3 To see that the single-crossing condition is fulfilled, denote profits by

$$\pi(\theta,\hat{\theta},p) = (p-c)(1-F_\theta(r_c)),$$

where $r_c = p - \beta(1-F_{\hat{\theta}}(r_c))$. Recall from Lemma 2 that the partial derivative of the cutoff with respect to p is

$$\frac{\partial r_c}{\partial p} = \frac{1}{1-\beta f_{\hat{\theta}}(r_c)} > 0.$$

The partial derivative of the cutoff with respect to $\hat{\theta}$ can be derived by totally differentiating the cutoff above:

$$\frac{\partial r_c}{\partial \hat{\theta}} = \frac{\beta F^{1,\hat{\theta}}}{1-\beta f_{\hat{\theta}}(r_c)} < 0.$$

The (strict) single-crossing property is satisfied if $\frac{\partial \pi(\theta,\hat{\theta},p)}{\partial p} / \frac{\partial \pi(\theta,\hat{\theta},p)}{\partial \hat{\theta}}$ is a strictly increasing function of θ. Taking the respective derivatives and rearranging, yields

$$\frac{\partial \pi(\theta,\hat{\theta},p)/\partial p}{\partial \pi(\theta,\hat{\theta},p)/\partial \hat{\theta}} = -\frac{(p-c)\left[\frac{-f_\theta(r)}{1-\beta f_{\hat{\theta}}(r_c)}\right] + 1 - F_\theta(r)}{(p-c)f_\theta(r)\left[\frac{-\beta F^{1,\hat{\theta}}(r)}{1-\beta f_{\hat{\theta}}(r_c)}\right]} \tag{2.18}$$

$$= \frac{1}{(-\beta F^{1,\hat{\theta}}(r))} + \frac{1-\beta f_{\hat{\theta}}(r_c)}{(p-c)(-\beta F^{1,\hat{\theta}}(r))}\frac{1-F_\theta(r)}{f_\theta(r)}, \tag{2.19}$$

where only the last term depends on θ. Recall that $^{1-F_\theta(r)}/_{f_\theta(r)}$ corresponds to the inverse hazard rate. Since the hazard rate is strictly decreasing in θ, the inverse hazard rate is strictly increasing in θ. Thus, the whole expression is strictly increasing in θ if $p > c$. $\square$

Lemma 7. *Choose any price p', users' response r', and type θ'. For any type $\theta < \theta'$, there exists $h \in \mathbb{R}_+$ and λ, such that $0 < \lambda < \varepsilon$ implies*

$$\pi(\theta', p_\lambda, r_\lambda) > \pi(\theta', p', r') \tag{A1}$$

$$\pi(\tilde{\theta}, p_\lambda, r_\lambda) < \pi(\tilde{\theta}, p', r'), \ \forall \tilde{\theta} \leq \theta, \tag{A2}$$

where $(p_\lambda, r_\lambda) = (p' + h\lambda, r' - \lambda)$.

Given the single-crossing property from Lemma 3, I show that there exists a small increase in price and increase in users' response (demand), i.e. a decrease in the cutoff r, that gives higher types scope to separate from lower types. Formally, the upcoming lemma provides a price-response pair for which a higher type would like to deviate whereas lower types do not. Note that both on-path and off-path user responses, i.e. the change in demand, are determined by the cutoff in Equation 2.1. Due to continuity, the cutoffs are $r \in [r(p, \delta_{\underline{\theta}}), r(p, \delta_{\overline{\theta}})]$ for given price p.

Proof of Lemma 7 The proof of this lemma follows Ramey (1996), but is simplified as there is only one signal p for which the one-dimensional single-crossing condition holds.

Take $\theta < \theta'$ and let $x \in \mathbb{R}$ be such that $x \geq MRS(\theta, p', r')$. Note that $x \neq MRS(\theta', p', r')$ and $\{MRS(\theta', p', r')\}$, $\{x\}$ are closed, convex sets as they are a singleton. Hence, it is possible to apply Minkowski's hyperplane separation theorem, which implies the existence of $h \in \mathbb{R}$, $h \neq 0$, such that

$$h \cdot MRS(\theta', p', r') < 1 < h \cdot x, \tag{2.20}$$

for some $h > 0$. Suppose $(p_\lambda, r_\lambda) = (p' + h\lambda, r' - \lambda)$ for $\lambda > 0$, i.e., a small increase in price and a small increase in demand (a small decrease in the cutoff). To determine whether $\pi(\tilde{\theta}, p_\lambda, r_\lambda) - \pi(\tilde{\theta}, p', r') < 0$, define

$$\zeta(\lambda, \tilde{\theta}) = \frac{-\pi(\tilde{\theta}, p_\lambda, r_\lambda) + \pi(\tilde{\theta}, p', r - \lambda)}{\pi(\tilde{\theta}, p', r - \lambda) - \pi(\tilde{\theta}, p', r')},$$

and then

$$\pi(\tilde{\theta}, p_\lambda, r_\lambda) - \pi(\tilde{\theta}, p', r') = \left[\pi(\tilde{\theta}, p', r - \lambda) - \pi(\tilde{\theta}, p', r')\right] \left[1 - \zeta(\lambda, \tilde{\theta})\right].$$

To determine the sign of $\zeta(\lambda, \tilde{\theta})$, observe that

$$\lim_{\lambda \to 0} \zeta(\lambda, \tilde{\theta}) = \lim_{\lambda \to 0} \frac{-\pi(\tilde{\theta}, p_\lambda, r_\lambda) + \pi(\tilde{\theta}, p', r - \lambda)}{\pi(\tilde{\theta}, p', r - \lambda) - \pi(\tilde{\theta}, p', r')} = -\frac{h\pi_p(\tilde{\theta}, p', r')}{\pi_r(\tilde{\theta}, p', r')} = h \cdot MRS(\tilde{\theta}, p', r'),$$

by l'Hospital rule and note that $\partial\pi(\cdot)/\partial r < 0$, so it is written as $-\pi_r(\cdot)$. Hence, it is possible to extend $\zeta(\cdot)$ continuously to $\lambda = 0$. Define

$$\zeta(\lambda, \tilde{\theta}) = \begin{cases} \frac{-\pi(\tilde{\theta}, p_\lambda, r_\lambda) + \pi(\tilde{\theta}, p', r-\lambda)}{\pi(\tilde{\theta}, p', r-\lambda) - \pi(\tilde{\theta}, p', r')} & \text{if } \lambda > 0 \\ -\frac{h\pi_p(\tilde{\theta}, p', r')}{\pi_r(\tilde{\theta}, p', r')} & \text{if } \lambda = 0 \end{cases}$$

In $\lambda \in \mathbb{R}_{>0}$ $\zeta(\lambda, \tilde{\theta})$ is differentiable as a composition of differentiable functions, however, the function is not differentiable in $\lambda = 0$ as $MRS(\tilde{\theta}, p', r') \neq 0$. For $\lambda > 0$ the function is strictly decreasing in λ

$$\frac{-h\pi_p(\tilde{\theta}, p', r')(\pi(\tilde{\theta}, p', r - \lambda) - \pi(\tilde{\theta}, p', r')) - (-\pi(\tilde{\theta}, p_\lambda, r_\lambda) + \pi(\tilde{\theta}, p', r - \lambda))(\pi_r(\tilde{\theta}, p', r'))}{(\pi(\tilde{\theta}, p', r - \lambda) - \pi(\tilde{\theta}, p', r'))^2} < 0$$

From Equation (2.20), it follows that $\zeta(\lambda = 0, \theta') < 1$ and $\zeta(\lambda, \theta') < 1$ as well, such that $\pi(\theta', p_\lambda, r_\lambda) - \pi(\theta', p', r') > 0$. Furthermore, $\zeta(\lambda = 0, \tilde{\theta}) > 1$ and hence $\zeta(\lambda, \tilde{\theta}) > 1$ for λ sufficiently small ($\lambda < \varepsilon$), such that $\pi(\tilde{\theta}, p_\lambda, r_\lambda) - \pi(\tilde{\theta}, p', r') < 0$, which needed to be shown. □

Proof of Proposition 1 By assumption $\tau(\theta)$ is a differentiable one-to-one strategy. Given that $\tau(\theta)$ is differentiable, it satisfies the differential equation in Equation 2.3 and hence, also the first-order condition implied by the incentive condition in Definition 1. Then, $\tau(\theta)$ satisfies the incentive condition if

$$\tau'(\theta)\pi_2(\theta, \hat{\theta}, p)\frac{d}{d\theta}\left\{\frac{\pi_3(\theta, \hat{\theta}, p)}{\pi_2(\theta, \hat{\theta}, p)}\right\} \geq 0, \tag{2.21}$$

which is proven in Theorem 6 of Mailath and von Thadden (2013). Note that the profit is monotonic in belief $\hat{\theta}$: $\pi_2(\theta, \hat{\theta}, p) = (p - c)n_2(\theta, \hat{\theta}, p) > 0$ for $p > c$. That is, the platform always has an incentive to manipulate beliefs in a way that users believe that the state is higher than it actually is. Taking the partial derivative, yields

$$\frac{\partial\pi(\theta, \hat{\theta}, p)}{\partial\hat{\theta}} = (p - c)f_\theta(r)\frac{-\beta F^{1,\hat{\theta}}(r)}{1 - \beta f_{\hat{\theta}}(r)} > 0.$$

Additionally, by Lemma 3 the last term is strictly positive if $p > c$. Thus, $\tau'(\theta)$ must be increasing for $p > c$ to fulfill Equation 2.21. Thus, for the platform to make positive profits, τ must be increasing in θ. From this, Corollary 4 follows.

Denote the equilibria inducing the equilibrium outcome in Proposition 1 by Γ_S, which only differ in their equilibrium messages. To show that the equilibrium outcome exists under the equilibrium refinement, suppose for contradiction that it fails to be one under the equilibrium refinement in Definition 4. Then there exists some price $p' \neq \{\tau(\theta)\}_{\theta \in \Theta}$ for which some type has a strict incentive to deviate when beliefs satisfy Definition 4. Define the cutoff r' under p' and set of types Θ' as follows

$$\begin{aligned} r' &\equiv \{r|\pi(\theta, r, p') = \pi^{*,S}(\theta) \text{ for some } \theta\} \\ \Theta' &\equiv \{\theta|\pi(\theta, r', p') = \pi^{*,S}(\theta)\}, \end{aligned}$$

where $\pi^{*,S}(\theta)$ are the equilibrium profits of type θ. Γ_S fails the equilibrium refinement if for any posterior satisfying $\mu(\theta|p') \subset \Theta'$ and any response $r(p')$, some type strictly prefers p' to the equilibrium action. That is, $r(p') < r'$ for some type to strictly prefer p' to the equilibrium action.

Then, fix $\theta' \in \Theta'$. By Lemma 7 for price p', cutoff r', and type θ', there exists a price-cutoff pair (p_λ, r_λ) such θ' can separate itself from lower types by choosing p_λ. Such separation ensures that no lower type would choose p_λ. By the definition of the equilibrium price, $\tau(\theta')$ is the least-cost signaling price, i.e. the smallest price for which type θ' can separate from lower types implying that $p_\lambda \geq \tau(\theta')$. Thus, type θ' has no incentive to deviate from the equilibrium price — a contradiction.

Proof of Proposition 2 I prove the following: In an equilibrium in which $p^*(\theta) = p'$ is set by more than one type, the highest type of the pool θ' can set price p_λ to break the equilibrium. For p_λ, there exists $r \in \mathcal{R}^\infty(\Theta^*(p_\lambda), p_\lambda)$ such that $\min_{r \in \mathcal{R}^\infty(\Theta^*(p_\lambda), p_\lambda)} \pi(\theta', r, p_\lambda) \leq \pi^*(\theta', \Sigma)$.

Consider an equilibrium candidate in which $p^*(\theta) = p'$ for more than one θ. Let $\theta' = sup\{\theta|p^*(\theta) = p'\}$ be the highest type in the pool and $r' = r(p')$ be the user response to observing price p' in equilibrium. Since $\{\theta|p^*(\theta) = p'\}$ is non-degenerate and $\mu_1(\theta|r, p')$ has full support on the closure of the set $\{\theta|p^*(\theta) = p'\}$, receivers place strictly positive probability on the set $cl\{\theta|p^*(\theta) = p'\} - \{\theta'\}$.

Given users use rationalizable strategies off-path, Lemma 1 provides a unique cutoff for given beliefs. Then, $n(\theta', p', r') < n(\theta', p', r^*)$, where the cutoff for the highest type in the pool is lower if users believe $\mu(\theta') = \delta_{\theta'}$ (r^*) than the cutoff r'

$$\begin{aligned} r' &= p' - \beta \int_\Theta (1 - F_\theta(r_c))\mu_1(\theta|r', p)d\theta, \\ r^* &= p' - \beta(1 - F_{\theta'}(r^*)), \end{aligned}$$

such that $r^* < r'$ as the users place strictly positive probability on lower types other than θ'.

The rest of the proof follows Ramey (1996). Take a type θ sufficiently close to the highest type of the pool θ' that yields $r' > r(p', \delta_\theta)$. Given θ, there exist small moves upwards in price and receiver response (p_λ, r_λ) supplied by Lemma 7 that satisfies Equation A1 and Equation A2. By taking λ sufficiently small $r_\lambda > r(p_\lambda, \theta)$.

For p_λ, the user response is either such that Equation A1 and Equation A2 are satisfied, or if $\pi(\tilde{\theta}, p_\lambda, r) \geq \pi(\tilde{\theta}, p', r')$ for $\tilde{\theta}$ resulting in $r > r_\lambda$, then because of the single-crossing property $\pi(\theta', p_\lambda, r) > \pi(\theta', p', r')$ such that θ' has a stricter incentive to deviate. In both cases, $D_{\tilde{\theta}} \cup D^0_{\tilde{\theta}} \subset D_{\theta'}$ holds, i.e. for types $\tilde{\theta}$ there are less rationalizable strategy profiles for which it can improve. Then, D1 criterion requires the support of $\mu^*(\theta|p_\lambda)$ to be in $[\theta, \overline{\theta}]$, i.e. $\Theta^*(p_\lambda) = [\theta, \overline{\theta}]$ with $\theta' \in [\theta, \overline{\theta}]$. By Equation 2.1, it must be that $r(p_\lambda, \delta_\theta) > r(p_\lambda, \mu(\Theta^*(p_\lambda))$, and Equation A1 implies that θ' has a profitable deviation breaking the equilibrium. $\square$

Proof of Proposition 3 Suppose Condition 1 is satisfied. I will construct the two types of separating equilibria in Theorem 3 (i) and (ii).

(i) Price Signaling Suppose first that zero fake profiles are used in equilibrium, i.e. $\xi(\theta) = 0$. The incentive constraints for $\theta' > \theta$ read

$$(p(\theta) - c)n(\theta, \theta, p(\theta)) \geq (p(\theta') - c)n(\theta, \theta', p(\theta')) - \gamma \left[n(\theta', \theta', p(\theta')) - n(\theta, \theta', p(\theta')) \right]$$
$$(p(\theta') - c)n(\theta', \theta', p(\theta')) \geq (p(\theta) - c)n(\theta', \theta, p(\theta)).$$

Observe that the incentive constraints for upward and downward deviations are asymmetric. If type θ' deviates downward to mimic type θ, it does not need to create fake profiles to match the lower demand. Instead, it would have to create negative fake profiles to reduce demand—something that is not feasible. At the entry stage, users hold the belief $\mu_1(\cdot) = \delta_\theta$ after observing the price $p(\theta)$, which leads to an entering mass of $n(\theta', \theta, p(\theta))$. This mass is smaller than the number of users who would have entered under the true type θ', so the platform is constrained in demand after entry, even if the true type θ' is (partially) revealed afterward.

Now observe that the second incentive constraint must be binding, whereas the first incentive constraint is slack. If the first IC were binding, the second IC would not be satisfied, and higher types would prefer to set the prices of lower types. Given differentiability, the resulting differential equation is

$$p'(\theta) = -\frac{(p(\theta) - c)n_2(\theta, \theta, p(\theta))}{(p(\theta) - c)n_3(\theta, \theta, p(\theta)) + n(\theta, \theta, p(\theta))}$$

The separating strategy is the same as in Proposition 1 and is indeed separating as $p'(\theta) \neq 0$ and $n(\theta, \theta, p(\theta)) \neq n(\theta', \theta', p(\theta')), \forall \theta \neq \theta'$ which can easily be verified.

Given that the equilibrium pricing strategy satisfies the differential equation and thus local incentive compatibility, global incentive compatibility is ensured by the single-crossing condition in Equation 2.19. For deviations from a higher type to a lower type's price, the argument follows as in Proposition 1. Global incentive compatibility for deviations from a lower type to a higher type's price is also satisfied—and even more straightforwardly. To see this, consider the following. If a deviation from $p(\theta)$ to $p(\theta')$, as well as from $p(\theta')$ to $p(\theta'')$ for $\theta'' > \theta' > \theta$, is unprofitable (due to local incentive compatibility), then, by the single-crossing property, a deviation of θ from $p(\theta)$ to $p(\theta'')$ is also unprofitable. Additionally, when deviating from $p(\theta)$ to $p(\theta')$ to induce belief θ', type θ needs to create additional fake profiles. A deviation from $p(\theta)$ to $p(\theta'')$ requires even more fake profiles because

$$n(\theta'', \theta'', p(\theta'')) - n(\theta, \theta'', p(\theta'')) > n(\theta', \theta', p(\theta')) - n(\theta, \theta', p(\theta'))$$

making a global deviation even less profitable.

(ii) Price and Fake Profile Signaling Suppose $\theta \in \hat{\Theta}$, where $\hat{\Theta} \subseteq [\underline{\theta}, \overline{\theta}]$. Fake profiles are given by

$$\gamma\xi = (p(\theta) - c)n(\theta, \theta, p(\theta)) - (p(\underline{\theta}) - c)n(\underline{\theta}, \underline{\theta}, p(\underline{\theta})) - \int_{\underline{\theta}}^{\theta} (p(t) - c - \gamma)n_1(t, t, p(t))dt,$$

and prices maximize equilibrium profits

$$\max_{p \in \mathbb{R}_+} (p(\underline{\theta}) - c)n(\underline{\theta}, \underline{\theta}, p(\underline{\theta})) + \int_{\underline{\theta}}^{\theta} (p(t) - c - \gamma)n_1(t, t, p(t))dt,$$

which results in

$$p^{**,S} \equiv p(\theta) = -\frac{n_1(\theta, \theta, p)}{n_{13}(\theta, \theta, p)} + c + \gamma. \tag{2.22}$$

The equilibrium exists under the following conditions: Let $\hat{\Theta}$ be such that $p^{**,S}(\theta) \leq p^{max}(\theta)$ holds for all $\theta \in \hat{\Theta}$, where $p^{max}(\theta)$ is given by the differential equation

$$p'(\theta) = -\frac{(p(\theta) - c)n_2(\theta, \theta, p(\theta)) - \gamma n_1(\theta, \theta, p(\theta))}{(p(\theta) - c)n_3(\theta, \theta, p(\theta)) + n(\theta, \theta, p(\theta))}.$$

For $p^{**,S}$ to maximize profits a necessary and sufficient condition is $f^{1,\theta}(r(p^{**,S})) > 0$. Lastly, γ must be sufficiently small

$$2\gamma \leq \left(-\frac{F^{1,\theta}(r)}{f^{1,\theta}(r)}(1 - \beta f_\theta(r))\right) \frac{-F^{1,\theta}(r)}{-F^{1,\theta}(r')} - \left(-\frac{F^{1,\theta'}(r')}{f^{1,\theta'}(r')}(1 - \beta f_{\theta'}(r'))\right). \tag{2.23}$$

I will construct the equilibrium outcome in the next steps. The incentive constraints for $\theta' > \theta$ are

$$\begin{aligned}&(p(\theta)-c)n(\theta,\theta,p(\theta))-\gamma\xi(\theta)\\&\geq (p(\theta')-c)n(\theta,\theta',p(\theta'))-\gamma\left[\xi(\theta')+(n(\theta',\theta',p(\theta'))-n(\theta,\theta',p(\theta')))\right],\end{aligned}$$

and

$$\begin{aligned}&(p(\theta')-c)n(\theta',\theta',p(\theta'))-\gamma\xi(\theta')\\&\geq (p(\theta)-c)n(\theta',\theta,p(\theta))-\gamma\left[\xi(\theta)-(n(\theta',\theta,p(\theta))-n(\theta,\theta,p(\theta)))\right],\end{aligned}$$

if $\gamma\left[\xi(\theta)-(n(\theta',\theta,p(\theta))-n(\theta,\theta,p(\theta)))\right] > 0$, or

$$(p(\theta')-c)n(\theta',\theta',p(\theta'))-\gamma\xi(\theta')\geq (p(\theta)-c)n(\theta',\theta,p(\theta)),$$

if $\gamma\left[\xi(\theta)-(n(\theta',\theta,p(\theta))-n(\theta,\theta,p(\theta)))\right] \leq 0$. I will show that it suffices to impose the incentive constraint only for nearby types. For close θ to θ', the incentive constraints are

$$\begin{aligned}&p'(\theta)[(p(\theta)-c)n_3(\theta,\theta,p(\theta))+n(\theta,\theta,p(\theta))]\\&+(p(\theta)-c)n_2(\theta,\theta,p(\theta))-\gamma\xi'(\theta)\leq\gamma n_1(\theta,\theta,p(\theta))\\&p'(\theta)[(p(\theta)-c)n_3(\theta,\theta,p(\theta))+n(\theta,\theta,p(\theta))]\\&+(p(\theta)-c)n_2(\theta,\theta,p(\theta))-\gamma\xi'(\theta)\geq\gamma n_1(\theta,\theta,p(\theta)).\end{aligned}$$

Setting the first constraint above to bind, implies that the second binds and vice versa. This results in a differential equation of the fake profile strategy as a function of prices:

$$\begin{aligned}&\gamma\xi'(\theta)=\\&p'(\theta)[(p(\theta)-c)n_3(\theta,\theta,p(\theta))+n(\theta,\theta,p(\theta))]+(p(\theta)-c)n_2(\theta,\theta,p(\theta))-\gamma n_1(\theta,\theta,p(\theta)).\end{aligned}$$

Sequential rationality implies that once the lowest type, $\underline{\theta}$, is identified as such, it cannot do better than setting zero fake profiles. This implies the following initial value condition for the differential equation: $\xi(\underline{\theta}) = 0$. For a given price $p(\theta)$, the differential equation can be solved via the Fourier method, yielding a unique solution (up to a constant):

$$\gamma\xi(\theta)=(p(\theta)-c)n(\theta,\theta,p(\theta))-(p(\underline{\theta})-c)n(\underline{\theta},\underline{\theta},p(\underline{\theta}))-\int_{\underline{\theta}}^{\theta}(p(t)-c-\gamma)n_1(t,t,p(t))dt.$$

Then, the equilibrium profit as a function of the price is given by

$$\begin{aligned}\Pi(\theta) &= (p(\theta)-c)n(\theta,\theta,p(\theta)) - \gamma\xi(\theta)\\ &= (p(\underline{\theta})-c)n(\underline{\theta},\underline{\theta},p(\underline{\theta})) + \int_{\underline{\theta}}^{\theta}(p(t)-c-\gamma)n_1(t,t,p(t))dt.\end{aligned}$$

and is maximized if prices are set to

$$\arg\max_{p(\theta)}(p(\underline{\theta})-c)n(\underline{\theta},\underline{\theta},p(\underline{\theta})) + \int_{\underline{\theta}}^{\theta}(p(t)-c-\gamma)n_1(t,t,p(t))dt\forall\theta\in\Theta.$$

Solving the maximization problem yields

$$p^{**,S} \equiv \begin{cases} p(\underline{\theta})-c = \frac{1-F_{\underline{\theta}}(r)}{f_{\underline{\theta}}(r)}(1-\beta f_{\underline{\theta}})(r) & \theta=\underline{\theta}\\ p(\theta)-c-\gamma = -\frac{F^{1,\theta}(r)}{f^{1,\theta}(r)}(1-\beta f_{\theta}(r)) & \theta\in(\underline{\theta},\overline{\theta}].\end{cases}$$

The necessary and sufficient condition for $p^{**,S}$ to maximize the platform's profits are given in the following. For $\theta=\underline{\theta}$, the profit function is concave under the assumptions (see Lemma 2), hence the only profit-maximizing price is the first-best price. For $\theta\in(\underline{\theta},\overline{\theta}]$, the profit function with respect to p is not necessarily single-peaked. The necessary and sufficient conditions are given by the first-order condition

$$(p-c-\gamma)n_{13}(\theta,\theta,p)+n_1(\theta,\theta,p)=0,$$

and the second-order condition

$$(p-c-\gamma)n_{133}(\theta,\theta,p)+2n_{13}(\theta,\theta,p).$$

For any price $p\in\mathbb{R}_+$, there exists a unique cutoff r by Lemma 1. If n_{13} changes sign (n_1 is strictly positive), there can be two price-response pairs solving the first-order condition. To do so, reformulate both conditions in terms of the underlying distribution

$$(p-c-\gamma)\frac{-f^{1,\theta}(r)}{1-\beta f_{\theta}(r)} - F^{1,\theta}(r)=0.$$

By Assumption 2, the density is single-peaked and hence $f^{1,\theta}(r)$ changes sign only once. If the density is right-skewed, $f^{1,\theta}(r)$ is strictly increasing in r. By the MLRP assumption, which implies that the hazard rate is monotonically decreasing in θ, it is possible to bound $f^{1,\theta}(r)$ by $f^{1,\theta}(r) < \frac{f_{\theta}(r)}{1-F_{\theta}(r)}(-F^{1,\theta}(r))$. Then, the first-order condition can exhibit (a) two solutions, or (b) only one solution.

(a) If the first-order solves for two solutions, one solution exists with $p-c-\gamma<0$ if $f^{1,\theta}<0$ and one solution exists with $p-c-\gamma>0$ if $f^{1,\theta}>0$. The second-order condition is

$$SOC := -\frac{F^{1,\theta}(r)}{f^{1,\theta}(r)} \frac{\partial f_\theta(r)/\partial r \partial\theta + f^{1,\theta}(r) f'(r)\beta/1-\beta f_\theta(r)}{1-\beta f_\theta(r)} + 2\frac{-f^{1,\theta}(r)}{1-\beta f_\theta(r)}.$$

If $f^{1,\theta} > 0$, then $SOC < 0$, and if $f^{1,\theta} < 0$, then $SOC > 0$. This implies that the profit function is well-behaved, so that $p - c - \gamma > 0$ is both the local and global maximum. To see that the interior maximum is also the global maximum, note that profits are zero when $p = 0$ and negative as $p \to \infty$.

(b) As $f^{1,\theta}(r)$ is bounded from above, only one solution at $p' - c - \gamma < 0$ with $f^{1,\theta} < 0$ can exist. In that case, the first-order condition is always larger than zero for prices higher that p'. The solution p' is a local minimum and the profit-maximizing prices are a corner solution. As the profit-maximizing solution must still adhere to incentive-compatibility, prices are chosen as high as possible given incentive-compatibility. The profit-maximizing solution is given by (i).

If the density is left-skewed, $f^{1,\theta}(r)$ is strictly decreasing in θ, i.e. $f^{2,\theta}(r) < 0$. The first-order condition solves for one solution with $p - c - \gamma > 0$ and $f^{1,\theta} > 0$. The price is again profit-maximizing as the second-order condition is negative at the solution.

Lastly, to check global incentive compatibility note that prices fulfill the single-crossing condition by Lemma 3. The single-crossing condition for fake profiles is given by

$$\frac{\frac{\partial \pi(\theta,\hat{\theta},p,\xi)}{\partial \xi}}{\frac{\partial \pi(\theta,\hat{\theta},p,\xi)}{\partial \hat{\theta}}} = \frac{-\gamma}{(p-c) f_\theta(r) \frac{(-\beta F^{1,\hat{\theta}}(r))}{1-\beta f_{\hat{\theta}}(r)}},$$

which is strictly increasing in θ if $f^{1,\theta}(r) > 0$. Hence, as long as $p(\theta) - c - \gamma = -\frac{F^{1,\theta}(r)}{f^{1,\theta}(r)}(1 - \beta f_\theta(r))$ is solved for a price-cutoff pair resulting in a positive mark-up ($p - c - \gamma > 0$) and a cutoff r such that $f^{1,\theta} > 0$, both signals meet the strict single-crossing condition.

Assume from now on $p - c - \gamma > 0$ and $f^{1,\theta} > 0$. As the fake profile strategy solves the differential equation, it thus, satisfies local incentive compatibility. To fulfill global incentive compatibility, for $\theta > \theta'$ (p, ξ) must satisfy

$$(p(\underline{\theta}) - c) n(\underline{\theta}, \underline{\theta}, p(\underline{\theta})) + \int_{\underline{\theta}}^{\theta} (p(t) - c - \gamma) n_1(t, t, p(t)) dt \geq (p(\theta') - c) n(\theta, \theta', p(\theta')) \quad (2.24)$$

$$- \left[(p(\theta') - c) n(\theta', \theta', p(\theta')) - (p(\underline{\theta}) - c) n(\underline{\theta}, \underline{\theta}, p(\underline{\theta})) - \int_{\underline{\theta}}^{\theta'} (p(t) - c - \gamma) n_1(t, t, p(t)) dt\right] \quad (2.25)$$

$$+ \gamma(n(\theta, \theta', p(\theta') - n(\theta', \theta', p(\theta')). \quad (2.26)$$

and for $\theta < \theta'$

$$(p(\underline{\theta}) - c)n(\underline{\theta}, \underline{\theta}, p(\underline{\theta})) + \int_{\underline{\theta}}^{\theta} (p(t) - c - \gamma)n_1(t, t, p(t))dt \geq (p(\theta') - c)n(\theta, \theta', p(\theta')) \quad (2.27)$$

$$- [(p(\theta') - c)n(\theta', \theta', p(\theta')) - (p(\underline{\theta}) - c)n(\underline{\theta}, \underline{\theta}, p(\underline{\theta})) - \int_{\underline{\theta}}^{\theta'} (p(t) - c - \gamma)n_1(t, t, p(t))dt] \quad (2.28)$$

$$- \gamma(n(\theta, \theta', p(\theta') - n(\theta', \theta', p(\theta')). \quad (2.29)$$

Observe that Equation 2.26 is harder to satisfy than 2.29 as the higher type must create less fake profiles to mimic the lower type. Hence, rewriting Equation 2.26 yields

$$\int_{\theta'}^{\theta} (p(t) - c - \gamma)n_1(t, t, p(t))dt \geq \int_{\theta'}^{\theta} (p(\theta') - c + \gamma)n_1(t, \theta', p(\theta'))dt.$$

A sufficient condition is

$$(\theta - \theta') [(p(\theta) - c - \gamma)n_1(\theta, \theta, p(\theta)) - (p(\theta') - c + \gamma)n_1(\theta, \theta', p(\theta'))] \geq 0,$$

such that the second term must be larger or equal to zero. Substituting the profit-maximizing prices $p^{**,S}$ into the condition and solving for γ yields

$$2\gamma \leq \left(-\frac{F^{1,\theta}(r)}{f^{1,\theta}(r)}(1 - \beta f_\theta(r))\right) \frac{-F^{1,\theta}(r)}{-F^{1,\theta}(r')} - \left(-\frac{F^{1,\theta'}(r')}{f^{1,\theta'}(r')}(1 - \beta f_{\theta'}(r'))\right). \quad (2.30)$$

Therefore, to fulfill global incentive compatibility, γ must be sufficiently small given by Equation 2.23.

Lastly, to guarantee that the non-negativity constraint on the fake profile strategy is not violated, the following condition must hold: For all $\theta \in \hat{\Theta}$,

$$p^{FP} \leq p^{max}(\theta), \quad (2.31)$$

where $p^{max}(\theta)$ is given by the differential equation $p'(\theta) = -\frac{(p(\theta)-c)n_2(\theta,\theta,p(\theta))-\gamma n_1(\theta,\theta,p(\theta))}{(p(\theta)-c)n_3(\theta,\theta,p(\theta))+n(\theta,\theta,p(\theta))}$. Note that $\gamma\xi'(\theta, p^{max}(\theta)) = 0$ and $\gamma\xi'(\theta, p^{FI}(\theta)) > 0$. Hence, the equilibrium with positive fake profiles exists for $\theta \in \hat{\Theta}$ for which the Inequality in Equation 2.31 holds.

Suppose Condition 1 is violated, then fake profiles are sufficiently costly for all types. As the cost of fake profiles are common knowledge, users anticipate that fake profiles cannot be used in equilibrium. Then, the equilibrium is the same as characterized in Proposition 6 (see proof of Proposition 6).

Proof of Proposition 4 The proof follows directly from the proof of Proposition 1 by considering the incentives for some type in the pool, $\theta' \in \Theta'$, to deviate from the pooled price p' in the first period. By Lemma 7, there continues to exist a price-cutoff pair (p_λ, r_λ) for type θ' that makes it possible to separate from lower types. This lemma is unaffected by the possibility of fake profiles as the strict single-crossing property holds for the price as a signal alone.[18] By the equilibrium refinement, users put strictly positive probability $\mu_1(\theta|p_\lambda)$ on a set Θ^* in support of strictly higher types than Θ'. Thus, in turn the cutoff decreases and the mass of entering users increases. As users cannot observe the number of fake profiles and do not know the true demand on the platform, the equilibrium refinement has no bite at this point and users do not revise their belief. Thus, $\mu_2(\theta|p_\lambda, n(\theta', \mu_1(\theta|p_\lambda), p_\lambda) + \xi)$ puts the same probability on set Θ^* and no user exits. Lastly, this implies that the profit of type θ' increases by the deviation, thereby breaking the pooling equilibrium.

Proof Lemma 4 Given the timing of the game, users "update" their beliefs twice. In the first period, users observe the message and hold beliefs $\mu_1(\theta|p, m, r) = m$. In the second period, users observe the perceived network size $\tilde{n} = n(\theta, \mu_1, p) + \xi$ and hold beliefs $\mu_2(\theta|p, m, r, \tilde{n})$. Note that $\pi_2(\cdot)$ is strictly increasing if the believed state increases.

Fix the price p and message m. After users' entry decisions, the true network size of users in equilibrium is bounded from above. Denote by $n(\theta, m, p)$ the number of users who enter given $\mu_1(\cdot) = m$, and by $n(\theta, \mu_2, p)$ the number of users on the platform given $\mu_2(\cdot)$. The number of users who remain on the platform are $\min\{n(\theta, m, p), n(\theta, \mu_2, p)\}$.

Suppose that $m \geq \theta$, where $m < \theta$ is irrelevant as profits increase in m; see above. Then to induce $\mu_2 \geq m$, the platform must create fake profiles such that:

$$\xi \geq n(m, m, p) - n(\theta, m, p). \tag{2.32}$$

Suppose the platform sets ξ such that $\mu_2(\cdot) > m$, i.e., the above inequality is strict. Then, the demand in equilibrium on the platform is $n(\theta, m, p) = min\{n(\theta, m, p), n(\theta, \mu_2, p)\}$. As demand is unaffected by the increase in fake profiles and profits decrease due to the cost of creating them, the platform does not set ξ such that $\mu_2(\cdot) > m$.

Suppose the platform sets ξ such that $\mu_2(\cdot) < m$, i.e, the inequality in Equation 2.32 is reversed. Then, demand in equilibrium on the platform is $n(\theta, \mu_2, p)$ and thus decreases. Since the platform has already chosen message m in the first period and found it optimal to do so, it will not set ξ such that $\mu_2(\cdot) < m$. This implies that the platform induces ξ such that $\mu_2(\cdot) = m$, i.e, the inequality in Equation 2.32 binds. $\square$

Proof of Lemma 5 First, I will show that the indifferent type is given as the unique solution to Equation 2.9. To get Equation 2.9, substitute the first-order condition in Equation 2.7 into the first-order condition in Equation 2.8. The indifferent type is given

[18] Ramey (1996) notes that "With multiple signals, such separating movements remain possible as long as the MRS of any one signal is strictly decreasing in type at every point of the space of signals and responses" (p.511).

by the following equation evaluated at $m = \overline{\theta}$:

$$\left(-\frac{n(\theta, m, p)}{n_3(\theta, m, p)} + \gamma \frac{n_3(m, m, p)}{n_3(\theta, m, p)}\right) n_2(\theta, m, p) - \gamma\left(n_1(m, m, p) + n_2(m, m, p)\right) = 0,$$

which is equal to Equation 2.9 when evaluated at $m = \overline{\theta}$. Note that $r(m, p) \equiv \overline{r}$ is given by

$$\overline{r} = p - \beta(1 - F_m(\overline{r}))|_{m=\overline{\theta}},$$

and the respective derivatives are

$$\begin{aligned}
\frac{\partial \overline{r}}{\partial p} &= \left.\frac{1}{1 - \beta f_m(\overline{r})}\right|_{m=\overline{\theta}}, \\
\frac{\partial \overline{r}}{\partial m} &= \left.\frac{(-\beta F^{1,m}(\overline{r}))}{1 - \beta f_m(\overline{r})}\right|_{m=\overline{\theta}}.
\end{aligned}$$

Substituting the respective derivatives into Equation 2.9 yields

$$\begin{aligned}
&\left(\frac{1 - F_\theta(\overline{r})}{f_\theta(\overline{r})}(1 - \beta f_{\overline{\theta}}(\overline{r})) + \gamma \frac{f_{\overline{\theta}}(\overline{r})}{f_\theta(\overline{r})} \frac{1 - \beta f_{\overline{\theta}}(\overline{r})}{1 - \beta f_{\overline{\theta}}(\overline{r})}\right) \cdot \frac{f_\theta(\overline{r}) \cdot (-\beta F^{1,\overline{\theta}}(\overline{r}))}{1 - \beta f_{\overline{\theta}}(\overline{r})} \\
&= \gamma\left(-F^{1,\overline{\theta}}(\overline{r}) + f_{\overline{\theta}}(\overline{r}) \cdot \frac{-\beta F^{1,\overline{\theta}}(\overline{r})}{1 - \beta f_{\overline{\theta}}(\overline{r})}\right).
\end{aligned}$$

Simplifying and setting $\theta = \tilde{\theta}$ yields Equation 2.10

$$\beta(1 - F_\theta(\overline{r})) = \gamma, \tag{2.33}$$

for $\overline{r} = p - \beta(1 - F_{\overline{\theta}}(\overline{r}))$, where p is given by Equation 2.7. Then the indifferent type $\tilde{\theta} \in \Theta$ is the solution to Equation 2.10. Since the left-hand side is a constant and the right-hand side is strictly increasing in θ, the equation has a unique solution. Define

$$\begin{aligned}
\overline{\gamma} &\equiv \beta(1 - F_{\overline{\theta}}(\overline{r})), \\
\underline{\gamma} &\equiv \beta(1 - F_{\underline{\theta}}(\overline{r})).
\end{aligned}$$

The solution to Equation 2.10 is unique and solves for a $\tilde{\theta} \in \Theta$ if and only if $\underline{\gamma} \leq \gamma \leq \overline{\gamma}$. □

Proof of Proposition 5 (i) Let $\underline{\gamma} \leq \gamma \leq \overline{\gamma}$. For $\theta < \tilde{\theta}$ the equilibrium strategy is given by Equation (2.8), which is separating. Note that the relevant separating strategy is defined with respect to m as users are unaware of $\xi(\theta)$ and ignore p as a signal. That is, $m \neq m'$ for $\theta \neq \theta'$. Fixing p, incentive compatibility is fulfilled as $\{p, m\} = \arg\max \pi(\theta, m, p)$, i.e., m is chosen to fulfill the first-order condition in Equation

2.8:

$$\beta[1 - F_\theta(r(m,p))] = \gamma.$$

Recall that users also believe in the message and the state conveyed by the network size off the equilibrium path. As the equilibrium outcome uniquely maximizes the platform's profit in each state, there exists no incentive to deviate from the equilibrium.

For $\theta > \tilde{\theta}$, types choose $m = \overline{\theta}$, resulting in pooling on the highest available message. These types have no incentive to deviate off-path. First, note that types cannot deviate upward in the message space: they would prefer to send a message above $\overline{\theta}$, but the type space is bounded above by $\overline{\theta}$, making such deviations infeasible. Second, consider off-path deviations in fake profiles for given $m = \overline{\theta}$. Since users also believe in m off-path, demand cannot exceed $n(\overline{\theta}, \overline{\theta}, p)$ after entry. Therefore, deviating in the fake profiles does not attract additional users and is only costly. Lastly, consider deviations in the price. Given that demand is $n(\overline{\theta}, \overline{\theta}, p)$, and that price p is optimally chosen by the platform conditional on the message $m = \overline{\theta}$, there is no profitable deviation in price either. Any deviation would reduce profits.

(ii) Let $\gamma < \underline{\theta}$, then for the lowest type it holds that

$$\beta(1 - F_{\underline{\theta}}(\overline{r})) > \gamma, \tag{2.34}$$

which implies that the lowest type sets $m(\underline{\theta}) = \overline{\theta}$. Since the left-hand side of the above equation is increasing in type, all types larger than the lowest type set $m(\underline{\theta}) = \overline{\theta}$ as well. There exists no incentive to deviate as prices are chosen optimally given message m.

(iii) Let $\gamma > \overline{\gamma}$, then for the highest type it holds that

$$\beta(1 - F_{\overline{\theta}}(\overline{r})) < \gamma, \tag{2.35}$$

which implies that the highest type sets $m(\overline{\theta}) = \overline{\theta}$. Since the left-hand side is increasing in type, all types smaller than the highest type set $m(\theta) = \theta$ as well. There exists no incentive to deviate as prices are equal to the full information prices. □

Proof Proposition 6 The platform maximizes its profit

$$\max_{p,\xi} (p - c)n(\theta, \theta, p) - \gamma\xi,$$

subject to incentive compatibility

$$\pi(\theta, \theta, p(\theta)) \geq \pi(\theta, \theta', p(\theta')).$$

Additionally, the platform's equilibrium demand is bounded above by the demand of users that enter: $n(\theta, \mu_1, p) \geq n(\theta, \mu_2, p)$.

First, suppose that the government bans fake profiles, i.e., $\xi = 0$. I will characterize the equilibrium outcome and show that if (i) the platform needs to label fake profiles, or (ii) must mandatorily disclose fake profiles, or (iii) Condition 1 is violated, this leads to the same equilibrium outcome.

Let $\xi = 0$. By Lemma 2, the unique solution to the first-order condition in Equation 2.11 is $p^{FI}(\theta)$. Suppose that the platform randomizes over messages in equilibrium.

Now, fix the equilibrium strategy $p^{FI}(\theta)$; this is separating as $\frac{\partial p^{FI}(\theta)}{\partial \theta} > 0$ if $f^{1,\theta}(r^{FI}) < 0$, and it is differentiable. To prove that this construction is a separating equilibrium, it must be shown that incentive compatibility is satisfied. As fake profiles are banned, the information structure is $\mathbb{I} = [0, 1]$; that is, users observe the true network size. This implies that given separating beliefs $\mu_1(p^{FI}(\theta)) = \theta$ in the first period, after observing network size n' in the second period, users hold beliefs $\mu_2(\cdot)$ according to

$$\mu_2(\cdot) = \left\{\theta' \in \Theta | n(\theta', \theta, p^{FI}(\theta)) = n'\right\}.$$

The equation solves for a unique θ', as sophisticated users can predict the cutoff r_c from $(\mu_1 = \theta, p^{FI}(\theta)$. Due to the first-order stochastic dominance of $F_{\theta'}(r)$ with respect to the true θ', for a given $r = r_c$, there exists only one θ' that solves the equation.

The incentive compatibility constraints are thus

$$\pi(\theta, \theta, p^{FI}(\theta)) \geq \pi(\theta, \theta, p^{FI}(\theta')). \tag{2.36}$$

By Lemma 2, $p^{FI}(\theta)$ uniquely maximizes the profit of θ and thus, the inequality is always satisfied. Similarly, for any out-of-equilibrium beliefs satisfying the equilibrium refinement, there exists no profitable deviation for any type. Similarly to the logic above, for a given deviation p' and out-of-equilibrium beliefs μ', users can predict r' such that they perfectly know the true state θ in the second period. Therefore, deviating from the full information price is never profitable.

Now, consider cases (i) to (iii). In case (i), the information structure is $\mathbb{I} = \{[0, 1], \mathbb{R}_+\}$, i.e. users can perfectly identify fake profiles. Setting $p^{FI}(\theta)$ as separating strategy implies that for any given amount of fake profiles, the incentive constraints are as in Equation 2.36. Since creating fake profiles is costly and the platform receives no benefit from doing so, it chooses $\xi = 0$.

In case (ii), the platform must mandatorily disclose its use of fake profiles. The equilibrium outcome with fake profiles is given by Theorem 3. The resulting profits, for given price and fake profile strategy $\rho(\tau) = (p, \xi)$, are

$$(p - c)n(\theta, \theta, p) - \gamma\xi,$$

for $p \neq p^{FI}(\theta)$ for $\theta \in \Theta \setminus \{\underline{\theta}\}$ and $\xi \geq 0$. This implies that profits are always smaller than under full information. Hence, disclosing $\xi = 0$ and setting full information prices dominates disclosing $\xi > 0$.

In case (iii), Condition 1 is violated. Then, there exists a separating equilibrium in which the platform sets the full information prices and zero fake profiles. Recall that the information structure is $\mathbb{I} = \{[0, 1] + \mathbb{R}_+\}$. The separating equilibrium with full information prices and zero fake profiles is incentive compatible if

$$(p^{FI}(\theta) - c)n(\theta, \theta, p^{FI}(\theta)) \geq (p^{FI}(\theta') - c)n(\theta, \theta', p^{FI}(\theta')) - \gamma\xi', \tag{2.37}$$

where $\xi' = n(\theta, \theta', p^{FI}(\theta')) - n(\theta, \theta, p^{FI}(\theta))$. Since Condition 1 is violated, γ is such that type θ is not willing to marginally or discretely increase ξ, thereby satisfying Equation 2.37. If $\xi = 0$, the incentive constraint reduces to Equation 2.36, which is again satisfied. $\square$

Chapter 3

Credibility of Secretly Colluding Manufactures in Retail Contracting

Co-Authored with Matthias Hunold, Johannes Muthers, and Alexander Rasch

3.1 Introduction

Manufacturers who secretly start an illegal cartel must increase their wholesale prices for initially often unsuspecting downstream retailers. These retailers may, however, refuse to accept higher prices for the fear of being outcompeted by fellow downstream firms that (continue to) receive better offers. There is evidence from cartel cases in which manufacturers have failed to achieve a retail price increase without explicit communication to their retailers. For example, in the German coffee cartel, the coffee roasters were only able to sustain higher wholesale prices after coordinating a retail price increase with their retailers (Holler and Rickert, 2022). Manufacturers coordinated prices with retailers in cartels involving Anheuser Busch (beer), Haribo (gummi bears), Ritter (chocolate), and Melitta (coffee). The same issue of convincing retailers to accept higher wholesale prices appears to be the underlying problem in a number of so-called hub-and-spoke cartels.[1] Although this obstacle to collusion has been widely discussed and documented in practice, it is hitherto unmodelled.

Conceptually, we argue that with strategic uncertainty of retailers, in the sense that, if retailers are unaware or uncertain about manufacturer collusion within a secret-contracting setting, the simple Nash-equilibrium logic is inadequate to study the initial formation of manufacturer collusion.[2] According to this logic, retailers know the manufacturers' strategies even off the equilibrium path, which is why retailers can perfectly predict whether manufacturers collude. We suggest that potential strategic misinterpretations and misunderstandings are important for the feasibility of collusion in vertically related markets.

In collusion models, the number of equilibria is usually infinite, creating scope for coordination problems. It is intuitive in the presence of competitive prices and low profits that manufacturers attempt to replace competition with collusion. To formalize this, we express retailers' incorrect expectations by allowing for potentially incorrect beliefs off the equilibrium path. In contrast to a Nash equilibrium in which downstream retailers anticipate the specific collusive strategy, retailers may incorrectly believe that collusion is unlikely or even infeasible. This is important not only when starting to collude, but also for the problem of opportunism that even a monopolist faces. For instance, if retailers are convinced of collusion, manufacturers may also exploit such beliefs. The equilibrium concept of weak Perfect Bayesian Equilibrium provides sufficient freedom off the equilibrium path to model beliefs that are not only shaped by the strategies, but also by exogenous events, such as the experience of a firm in a different market. To capture the problem of

[1] See, for example, Harrington (2018) for a description of the cheese cartel in the UK.

[2] Strategic uncertainty and collusion have been analyzed by Blume and Heidhues (2008); however, they consider strategic uncertainty among cartel members only. Instead, to capture the many settings described below in which retailers did not suspect the cartel's existence, we focus on strategic mistakes or uncertainty by retailers.

the initial formation of collusion we characterize the collusive equilibrium and ask under which conditions it can be reached.

Formally, we analyze an infinitely repeated pricing game that features private contracting within each of two exclusive manufacturer-retailer pairs. We focus on manufacturer collusion and retailers who maximize short-term profits. The equilibrium concept of (weak) perfect Bayesian equilibria requires that retailers' beliefs are correct on the equilibrium path. Following a deviation, however, the market outcome depends on the retailers' responses. Because these responses are shaped by the retailers' beliefs, it is useful to contemplate which beliefs retailers hold and act upon. We study how retailers' beliefs may react dynamically to the actions observed in previous periods. In our model, *within-period* passive beliefs capture that retailers view a deviation by one manufacturer as independent from the other manufacturer's behavior. Passive beliefs are also prominently considered in the literature on private contracts in one-shot games (see, for example, Segal, 1999 and Rey and Vergé, 2004).

In our analysis, we first consider beliefs that do not change in past behavior. For these beliefs, we find that manufacturers employing efficient non-linear tariffs and grim-trigger strategies are incapable of sustaining *any* price above the competitive price. Hence, such non-reactive beliefs are self-fulfilling, rendering self-sustaining collusion infeasible. Two forces are driving this result. First, a colluding manufacturer cannot commit that the other manufacturer will offer the same contract, such that its retailer may fear that the rival receives a better offer. Therefore, in contrast to a monopolist, it is not the lack of own commitment that drives the result, but the inability to credibly commit in the name of the other manufacturer. Second, in line with the above cited evidence from various collusion cases, the retailer lacks the (equilibrium) knowledge that collusion occurs. This effect is novel and practically important, but conceptually difficult to capture.

We then turn to beliefs that adapt to observed past behavior of the manufacturers. First, we consider beliefs that (correctly) anticipate the manufacturers' trigger strategies. The retailers thus implicitly understand not only that manufacturers collude, but also how they collude. For such beliefs that capture perfect anticipation of the collusive behavior, collusion at industry-profit-maximizing prices is feasible and may be stable with respect to unilateral deviations, contrasting the result with non-reactive beliefs. Note that there are indeed cases in which upstream firms were apparently able to establish collusion without having to resort to communication with downstream firms.[3] At the same time, however, collusion may not be *opportunism-proof* because the manufacturers may have a joint

[3]A number of cartels have been detected in the automotive industry, where suppliers coordinated on high prices at the expense of carmakers and eventually final consumers. The cases involved rolling bearings (Frankfurter Allgemeine Zeitung, "Schaeffler has to pay a cartel fine of 370 million euros", 03/19/2014), safety belts, airbags, and steering wheels (Tagesschau, "Millions in fines against auto suppliers", 03/05/2019), and doors and electric window lifts (European Commission, "Antitrust: Commission fines car parts suppliers of € 18 million in cartel settlement", 09/29/2020); last accessed 1/31/2025. Recently, authorities in several European countries have started an investigation into the market for fra-

incentive to lower prices, similarly to the monopolist in Hart et al. (1990). We show that the opportunism condition is more restrictive than the stability condition.

Opportunism-proofness relates to the concept of renegotiation-proofness introduced by Farrell and Maskin (1989) in the sense that opportunism-proofness is a necessary condition for renegotiation-proofness. We apply their concept to the coalition of colluding manufacturers and define opportunism in this context as a joint deviation from the collusive agreement. This reflects that manufacturers may want to renegotiate the collusive wholesale price if a joint deviation is profitable for given actions of the retailers. Moreover, to ensure that manufacturers do not want to renegotiate the wholesale price during the punishment phase, we additionally consider a condition on *credible punishment.* If both conditions on opportunism-proofness and credible punishment are met, the collusive equilibrium is renegotiation-proof.

Collusion with trigger strategies usually is not renegotiation-proof (see, for example, Bernheim and Ray, 1989 and Farrell and Maskin, 1989): firms have an incentive to jointly deviate in the punishment phases by renegotiating higher prices, which undermines the credibility of the necessary threat to punish. Although a few solutions for certain settings are known, such as asymmetric punishment, the problem of renegotiation appears to be simply ignored in the applied literature.[4] We find that if the colluding firms sell via retailers using private contracts, renegotiation-proof equilibria with the usual symmetric punishment exist. If retailers fully anticipate the collusive strategy, they expect a punishment phase once a manufacturer deviates by making an out-of-equilibrium offer. This expectation of punishment can render punishment credible because retailers are not willing to accept the original collusive contracts during the punishment phase. Hence, if the stricter opportunism-proofness condition holds, with beliefs that anticipate a collusive strategy, collusion can become a renegotiation-proof equilibrium of the game for sufficiently patient manufacturers.

The case descriptions mentioned in the first paragraph reveal that manufacturers occasionally struggle to establish collusion because retailers do not accept the increased wholesale prices. This raises the question under which condition retailers' responses allow for the formation of collusion. We formally define the *formability* of collusion that requires the existence of a potential path from a non-collusive history to collusion. For example, with beliefs that anticipate collusion with grim-trigger strategies, retailers' responses make collusion impossible after observing a single period of competition. Hence,

grances and fragrance ingredients involving the world's largest manufacturers. See, for example, Tagesschau, "DAX group Symrise under suspicion of collusion", 03/08/2023; last accessed 1/31/2025.

[4]An exception is McCutcheon (1997). We differ from McCutcheon (1997) who builds on the model of costly renegotiation proposed by Blume (1994). By contrast, following Farrell and Maskin (1989) we assume that renegotiation is costless and establish that the collusive equilibrium can nevertheless be renegotiation-proof when considering renegotiation between the collusive players.

pure grim-trigger beliefs and also history-independent beliefs do not support the formation of collusion.

We introduce an adaptive belief, whereby retailers expect collusion or competition depending on the past behavior of the manufacturers. These intuitive beliefs can describe settings in which players cannot perfectly anticipate the actions of other players in parts of the game but are able to learn and adapt to the equilibrium. We parameterize the adaption speed that specifies how fast a retailer that believes in competitive prices but receives collusive wholesale price offers switches to believing in collusive prices. The adaption speed may range from a single period to many periods. We find that collusion can be formable with these adaptive beliefs, but such adaptive beliefs perform poorly in the punishment phase in the sense that punishment is not renegotiation-proof. We find that a faster speed of adaption makes collusion easier to form but harder to sustain: collusion becomes less opportunism-proof.

In summary, our key contributions are as follows. First, from an application point of view, we show that, in a vertically related structure, manufacturer collusion may be difficult to establish. We even find that, for certain retailer beliefs and supply contracts, collusion cannot be enforced at all. Also, we demonstrate that the opportunism problem, which is well established in the monopoly context, has rich implications for the theory of cartel stability and formation because it may occur during both collusive and punishment periods. Furthermore, starting to collude can be complicated because retailers may not be willing to accept the new collusive contract. Starting to collude is in various instances more difficult than sustaining collusion. Second, from a conceptual point of view, we go beyond the stability condition, which is the main focus of the analysis in many theories on collusion, and model other factors that may hinder collusion.

The remainder of the article is structured as follows. Section 3.2 relates our model to the relevant literature. We set up the model in Section 3.3 and subsequently analyze beliefs that are independent of the history of the game in Section 3.4.1. In Section 3.4.2, we study trigger beliefs and adaptive beliefs. We review symmetric beliefs in Section 3.5.1 and analyze linear wholesale prices in Section 3.5.2. We provide a competition policy discussion in Section 3.6, where we relate our theory to business practices which manufacturers may use to deal with the problem of formation and opportunism. These practices include communicating non-binding price increase announcements, resale price maintenance, vertical integration, communication between retailers, downsizing of packs, and buyback policies. Section 3.7 concludes.

3.2 Related Literature

Our paper contributes to four aspects that have been analyzed in the literature: (i) manufacturer collusion, (ii) a monopolist's opportunism problem, (iii) downstream retailers' types of beliefs, and (iv) cartel formation.

Manufacturer collusion. Nocke and White (2007) and Normann (2009) analyze tacit collusion among manufacturers in vertical relationships but, in contrast to us, focus on whether vertical integration makes tacit collusion easier to sustain. Both articles consider perfect Bertrand competition among manufactures and compare an industry with no integration to one in which one pair of firms is vertically integrated. In Nocke and White (2007), manufacturers compete in two-part tariffs. The authors show that it is easier for manufacturers to sustain collusion in a scenario with vertical integration. Normann (2009) shows that this finding carries over to a situation in which manufacturers set linear prices, even though double marginalization leads to different collusive and deviation profits. Piccolo and Miklós-Thal (2012) find similar results for the case in which retailers have full bargaining power. Under public contracts, Schinkel et al. (2008) show that, when manufacturers have full bargaining power but need to make sure that retailers do not sue for private damages, upstream collusion requires low wholesale prices and possibly negative franchise fees. Piccolo and Reisinger (2011) analyze the impact of exclusive territories granted to retailers on the feasibility of manufacturer collusion. Under observable contracts, establishing exclusive territories has two opposing effects on collusive stability. Exclusive territories soften punishment but they also reduce deviation profits. The second effect is due to the fact that, when a manufacturer deviates, retailers of competing products adjust their prices, whereas retailers of the same product do not. Because the effect on deviation tends to dominate, exclusive territories tend to facilitate tacit collusion.

Our contribution to this literature is twofold: First, we study the opportunism problem that colluding manufacturers face with secret contracting. Second, we demonstrate the relevance of retailers' behavior under strategic uncertainty for establishing and maintaining collusion.[5]

Opportunism problem. We relate to the classic opportunism problem of a monopolist in a vertical structure with secret contracting and two-part tariffs (Hart et al., 1990; O'Brien and Shaffer, 1992; McAfee and Schwartz, 1994) as colluding manufacturers are

[5] The literature also contains explanations of how resale price maintenance (RPM) can facilitate manufacturer collusion but abstracts from the relevance of retailer beliefs. Jullien and Rey (2007) study how (RPM) affects collusion when only the retailers observe local shocks on demand. They assume that colluding manufacturers reveal all wholesale prices to all retailers. Hunold and Muthers (2020) show that, absent any uncertainty and asymmetric information, RPM can still facilitate manufacturer collusion when there is retail bargaining power. They focus on subgame perfect Nash equilibria.

similar to a monopolist in that they are competing with themselves when making coordinated offers to different retailers. In such a scenario, the upstream firm that deals with multiple competing downstream firms through bilateral contracts may – as discussed further below – encounter the following problem: The upstream firm is interested in maintaining high prices and profits but it cannot commit to refraining from opportunistic moves. Indeed, the upstream firm has an incentive to increase bilateral surplus with one downstream firm, which is anticipated by the other downstream firm(s). The existence of this opportunism problem has been evoked as an explanation for vertical mergers and various vertical restraints as measures aimed at restoring the upstream firm's market power (O'Brien and Shaffer, 1992; McAfee and Schwartz, 1994; Rey and Vergé, 2004). The restraints include exclusive dealing, non-discrimination clauses, and industry-wide RPM. We also relate to Gaudin (2019) who shows that, in the framework of Hart et al. (1990), linear contracting can lead to higher retail prices absent collusion as well. Do and Miklós-Thal (2021) explore shortcomings of the seminal papers by considering a version of sequential (re)contracting between upstream and downstream firms. The opportunism problem that we study in detail has been neglected in the collusion literature so far and we establish that it has important consequences. We complement the literature by introducing an opportunism-proofness condition that is related to concept of renegotiation proofness by Farrell and Maskin (1989).

Beliefs. The literature on secret contracting between manufacturers and retailers, which dates back to Hart et al. (1990), McAfee and Schwartz (1994) and Segal (1999), emphasizes the relevance of the retailers' beliefs. The literature mostly focuses on passive, symmetric, or wary beliefs in static settings. Whereas passive beliefs suppose that the agents treat unexpected offers as mistakes, symmetric beliefs could correspond to a rule of thumb, where agents conjecture that identical principals make identical offers (Pagnozzi and Piccolo, 2011). Wary beliefs – according to which a retailer anticipate that rivals get the offer that maximize the manufacturer's profits – are often used when passive beliefs are implausible or induce non-existence of equilibria (Rey and Vergé, 2004; Rey and Tirole, 2007; Miklós-Thal and Shaffer, 2016). Empirical evidence on passive and symmetric beliefs was tested by Zhang (2021) and Martin et al. (2001) in experiments. Moreover, Aoyagi et al. (2024) study beliefs in finitely and infinitely repeated games experimentally and find that the same history of play can lead to different beliefs, and the same belief can lead to different action choices. We contribute to this literature by providing an in-depth analysis of how retailers' beliefs affect manufacturer collusion in a dynamic setting with secret contracting. In our analysis, we show how the retailers' contract acceptance decisions change depending on their perception of the collusive strategy ranging from fully anticipating the collusive (trigger-) strategies to not expecting any collusion.

In repeated games, the notion of passive beliefs differs across papers as the exact notion from static settings cannot be applied. Within this literature, our paper is most closely related to Piccolo and Reisinger (2011), Gilo and Yehezkel (2020), and Liu and Thomes (2020). In an extension, Piccolo and Reisinger (2011) consider secret contacting to investigate strategic delegation in a model with colluding manufacturers. The authors consider a perfect Bayesian equilibrium with passive beliefs. Their use of passive beliefs requires that those are correct on and *off* the equilibrium path. Similar to their approach, Liu and Thomes (2020) use weak perfect Bayesian equilibrium and consider passive beliefs as well as symmetric beliefs. They do not define passive beliefs explicitly but employ them in the sense that out-of-equilibrium offers do not change the expectation about the other manufacturer (in deviation periods). In punishment periods, they assume that retailers understand that there is punishment going on, such that the outcome is the competitive outcome of the stage game. Thereby, they implicitly assume that the beliefs of retailers are correct in the punishment period, which is the case if retailers anticipate the complete strategy of the manufacturers. That is, their definition is equivalent to a definition of passive beliefs whereby retailer beliefs are consistent with the strategy profile of the manufacturers in any subgame.

This consistency does not follow from the definition of weak PBE, which is silent on beliefs in nodes that are not reached on the equilibrium path with positive probability. It presumes that retailers have a correct expectation even in subgames that they never observe, given the strategy of the manufacturers, and thus cannot have any experience with. The beliefs thus rely on information that retailers cannot obtain on the equilibrium path. In addition, in this game, a similar equilibrium path can be sustained by many different manufacturer strategies that imply different behaviors off the equilibrium path. Thus, it is also theoretically impossible for retailers to predict the off-equilibrium behavior with certainty. In this context, we consider different beliefs explicitly as they correspond to different understandings about how the market works.

Gilo and Yehezkel (2020) study secret contracts and vertical collusion involving retailers and a single manufacturer. Their approach mirrors ours in the sense that they analyze the role of the manufacturer in retailer collusion, whereas we look at the retailers' role in manufacturer collusion. The major difference is that we abstract from retailer collusion by assuming short-lived retailers, whereas repeated retailer interaction is at the heart of their analysis. Gilo and Yehezkel consider symmetric pure-strategy, perfect Bayesian equilibria. Whereas the asymmetry in information about the contracts is similar, the timing differs from ours because the retailers in their model are those who offer contracts. They use the concept of "rational beliefs," which they define as players anticipating the rational action of others in off-equilibrium situations. They argue that this implies that manufacturers and retailers both understand whether a period is collusive or not after a deviation.

Cartel formation. The literature on cartel membership formation has focused on a variety of different aspects. In this literature, the question of how to initiate cartels typically focuses, among other things, on contracts (Selten, 1973; d'Aspremont et al., 1983), on stochastic opportunities to form a cartel (Harrington and Chang, 2009), on the heterogeneity with regard to capacities and umbrella pricing (Bos and Harrington, 2010), on signals and beliefs of producers (Harrington and Zhao, 2012; Harrington, 2017), on antitrust as a facilitating factor (Andersson and Wengström, 2007; Bos et al., 2013, 2015), on quality differentiation (Bos et al., 2020), and the ability to overcome strategic uncertainty absent communication (Blume and Heidhues, 2008). Selten (1973) analyzes the case of quantity competition and assumes that cartels can be enforced via contracts, and that a cartel acts as a Stackelberg leader. He shows that a cartel is stable in the sense that outsiders do not want to be part of the cartel, and insiders cannot profitably leave the cartel as long as the number of cartel members is relatively small. d'Aspremont et al. (1983) obtain a similar result for the case of price leadership.

Harrington and Chang (2009) consider a set of heterogeneous industries in which stochastic opportunities to form a cartel arise to explain the birth and death of cartels and to inform antitrust authorities about the extent of cartels that have not been discovered. Bos and Harrington (2010) endogenize the composition of a cartel in an industry in which heterogeneous firms differ in their capacities. They show that non-all-inclusive cartels set umbrella prices, and that mergers involving moderate-sized firms may result in the most severe coordinated effects. Harrington and Zhao (2012) analyze whether different types of players (patient and impatient) manage to cooperate via grim-trigger strategies when players signal and coordinate through their actions. The authors show that there is always a positive chance of cooperation, but cooperation may fail altogether. Moreover, the longer cooperation does not occur, the less likely it is to occur in the next period. Harrington (2017) focuses on mutual beliefs to coordinate prices. In the context of price leadership, firms are assumed to commonly believe that price increases will be at least matched. The firms, however, lack any shared understanding about who will lead, when they will, and at what prices. Sufficient conditions are derived, which ensure that supra-competitive prices emerge, but price is bounded below the maximal equilibrium price.

In contrast to the literature on cartel membership formation, we address the question whether firms can transit to a collusive equilibrium once they have reached – possibly explicitly – a common understanding to collude in a currently non-collusive industry. We thereby focus on the process of establishing collusive outcomes that firms need to go through until they may reach a stable collusive equilibrium. At the core of our analysis is the response of retailers in the transition process.

3.3 Model

We study manufacturer collusion in an infinitely repeated stage game with time $t = 0, \ldots, \infty$.[6] There are two symmetric manufacturers, M_A and M_B, that compete by selling horizontally differentiated products to their exclusive retailers R_A and R_B. Each manufacturer makes an exclusive and secret two-part tariff offer with a unit wholesale price w_i and a franchise fee F_i to its retailer, with $i \in \{A, B\}$.[7]. In general, with more than two manufacturers, retailers may interact with only a subset of the manufacturers, yielding a similar setting as the one we study with two pairs of exclusive manufacturers and retailers.

3.3.1 Set-Up

Timing and information. In each period, the following stage game unfolds:

1. Each manufacturer makes a private contract offer to its retailer.
2. Each retailer decides whether to accept its offer. Post contract acceptance, the fixed fees are sunk.
3. The retailers simultaneously and non-cooperatively set their retail prices p_i, $i \in \{A, B\}$.

The manufacturers' contract offers and the retailers' contract acceptance decisions are secret. Hence, a retailer cannot observe the contract offered to the rival. Moreover, when competing in the downstream market, the retailers are unable to observe each others' input prices and, thus, are forming beliefs about their rival's contract.

At the end of each period, all actions are revealed to all players. All players thus know the complete history of the game at the end of a period. We focus on the secret contracting problem between manufacturers and retailers and avoid that manufacturer collusion is directly hampered by the long-term unobservability of manufacturers' actions to each other.[8]

The supergame is a game of complete information but unobservable actions. The manufacturers are long-lived and discount next-period profits with the common discount factor δ with $0 < \delta < 1$. The retailers have a discount factor of zero, such that they do not take future profits into account. This assumption ensures that the retailers cannot

[6] We use the terms manufacturer for upstream firms and retailer for downstream firms. Our arguments naturally extend to any stage in a supply chain.

[7] An extension to retailers serving the products of both manufacturers as common agents is interesting but not straightforward. In this case, a retailer would receive offers of both manufacturers but would remain uninformed about the offers that other retailers receive. This may lead to a variant of the classic opportunism problem (Hart et al., 1990) and to non-existence of equilibria (Rey and Vergé, 2010)

[8] Note that we do not consider a scenario in which the (colluding) manufacturers try to get retailers on board, that is, there is no communication between manufacturers and retailers on the issue of collusion. This appears to be in line with the cartel cases mentioned in the Introduction.

collude. One can similarly assume that the retailers are short-lived but one should keep in mind that the retailers do see the history of the game.[9]

Assumptions on costs, demand, and profits. All (fundamental) costs are zero. We assume that the outside option (opportunity cost) of each retailer is equal to zero. We consider general demand functions that fulfill the standard properties summarized in

Assumption 5. *Demand $D_i(p_i, p_{-i})$ (with $i \in \{A, B\}$)*

- decreases in the own price p_i $(\partial D_i(p_i, p_{-i})/\partial p_i < 0)$,
  - increases in the other product's price p_{-i} $(\partial D_i(p_i, p_{-i})/\partial p_{-i} > 0)$, and
  - the own-price effect dominates the cross-price effect $(|\partial D_i(p_i, p_{-i})/\partial p_i| > \partial D_i(p_i, p_{-i})/\partial p_{-i})$.

To ensure that there exists a unique and stable equilibrium in the downstream market, we assume that the Hessian matrix of $D_i(p_i, p_{-i})$ has a negative and dominant main diagonal. This results in well-behaved retail profits that are twice differentiable and concave. Note that this also implies that the retailers' reactions behave normally, such that $\partial p_i \left(w_{it}, p_{-i}\right) / \partial w_{it} > 0$ and, consequently, $\partial D_i \left(p_i, p_{-i}\right) / \partial p_i < 0$ hold.[10]

Our assumptions on retailer profits mostly carry over to manufacturer profits because manufacturers internalize retailer profits using two-part tariffs. In some cases, however, manufacturers' true actions and retailers' beliefs differ in such a way off the equilibrium path that the behavior of the manufacturers' profit is not identical to that of the retailers' profit. In those cases, we analogously assume that manufacturers' profits are well behaved, such that optimal behavior can be derived from the respective first-order conditions. We comment on these cases below.

Equilibrium concept. We consider (pure-strategy) weak perfect Bayesian equilibria (weak PBE) and focus on symmetric equilibria.[11] We use the formal definition of a weak perfect Bayesian equilibrium from Mas-Colell et al. (1995): A profile of strategies and a system of beliefs is a weak PBE in extensive form games if it has the following properties:

1. The strategy profile is sequentially rational for the given belief system (each player maximizes expected utility at each node).

[9] Related literature assumes that retailers are short-lived or only live for one period. See, for example, Piccolo and Reisinger, 2011 and Jullien and Rey, 2007. Such an assumption rules out hub-and-spoke collusion as well as vertical collusion Gilo and Yehezkel (2020), where retailers are a part of the collusion.

[10] This can be shown by applying the implicit function theorem on the retailer's first-order condition for optimal pricing.

[11] The game defined includes two groups of players, manufacturers and retailers. Within each group, the players are symmetric. The equilibria we study are strongly symmetric in the sense that the (symmetric) manufacturers use a common continuation strategy on and off the equilibrium path (Athey et al., 2004; Jullien and Rey, 2007). The retailers play symmetric equilibrium actions in the downstream market.

2. The system of beliefs is derived from the strategy profile through Bayes' rule whenever possible.

This implies that the retailers' beliefs are consistent with equilibrium strategies on the equilibrium path. However, when observing an out-of-equilibrium offer, the second condition does not apply as this information set is reached with zero probability, meaning that Bayes' rule does not restrict these beliefs. Throughout the analysis we focus on the case that the manufacturers correctly anticipate the retailer beliefs.

We impose an additional condition of "no-signaling-what-you-don't-know" for most of the analysis.

Condition 1. The belief of retailer i (about its rival retailer $-i$) at the beginning of period $t+1$ depends only on the history up to date t ($\mathcal{H}_t$), but not on the action of manufacturer i at date t ("no-signaling-what-you-don't-know").

This condition captures the idea that the deviation of one manufacturer should not signal (private) information about the other manufacturer. Hence, a retailer's belief should not change.[12] Condition 1 has an intuitive interpretation in our game: A deviation of one manufacturer should not change the retailer's belief about the other manufacturer's contract offer. The condition, then, is equivalent to within-period passive beliefs.

We use weak PBE in conjunction with Condition 1 to be explicit about our equilibrium definition as there is a lack of clarity in the literature about the definition of perfect Bayesian equilibrium (PBE). According to the definition in Fudenberg and Tirole (1991a), our equilibrium would be considered a PBE. Other definitions of PBE, e.g., in González-Díaz and Meléndez-Jiménez (2014), require subgame perfection, which we do not want to generally impose. In our setting, subgame perfection would require that the retailers correctly anticipate the manufacturers' actions in information sets that are reached with zero probability. We explicitly want to allow for retailers' incorrect judgment in these nodes. For instance, this allows retailers to misjudge which of the different trigger strategies manufacturers employ if they feature the same equilibrium path but different off-equilibrium punishment strategies.

[12]Condition 1 is derived from the definition of the PBE that Fudenberg and Tirole (1991a,b) provide. This definition is suited for games of incomplete information with independent types. Because in our game of complete information the private information of the retailers is generated by manufacturers' actions, the concept can only be applied analogously. Similarly to Pagnozzi and Piccolo (2011), we adapt condition B(iii) ("no-signaling-what-you-don't-know") on p. 332 in Fudenberg and Tirole (1991a) to our game. Originally, the condition means that, in a signaling game, the actions of other players with independent types have no effect on beliefs about a player's type if this player acts the same. In our setting, because retailers act simultaneously, one retailer cannot observe the other retailer's action such that the condition translates into: A retailer cannot infer anything about a rival's current pricing from the current action of its own manufacturer.

The condition is also in the spirit of sequential equilibrium, where (out-of equilibrium) beliefs are naturally passive within a period because deviations are the result of a random mistake ("tremble"). Note that we cannot apply sequential equilibrium because we have a continuous action space.

Our game is reduced-form but reflects a game that could be richer, for instance, by including a cheap-talk stage. If retailers consider it possible that manufacturers may communicate, manufacturers' actions might be correlated when a deviation occurs. Given this interpretation, Condition 1 does not naturally capture the spirit of sequential equilibrium in this "extended game."[13]

Each retailer forms a belief about the contract offer made to the rival and about how the rival reacts to its contract offer when accepting the contract in the second stage. Following Rey and Vergé (2004), we focus on retailers that form beliefs about the resulting retail price of the rival, which is the payoff-relevant information that retailers are lacking.[14] In our setting there is a unique mapping from the belief a retailer has about its rival's wholesale price to the expected retail price of the rival.[15]

3.3.2 Formability and Renegotiation-Proofness

We focus on the phenomenon of collusion among manufacturers and aim to address additional issues beyond the standard stability and incentive compatibility problem. In particular, we shed light on the possibility for manufacturers to form collusion, that is, the transition from a competitive to a collusive market outcome. Moreover, we highlight the robustness of collusion to the opportunity of manufacturers to renegotiate, both during the collusive phase and after a deviation. Both issues are theoretically and empirically appealing because collusion in real-world markets has various important aspects besides the stability problem. To characterize whether the collusive equilibrium allows for renegotiation or cartel formation, we define conditions that can be checked in a similar fashion to the traditional stability condition.

Definitions. First, we characterize how manufacturers can overcome the coordination problem by defining under which conditions collusion is formable in the sense that collusion can be started and maintained. For example, suppose that manufacturers have been competing for a large number of periods and want to start collusion with trigger strategies. Formability addresses the question whether the system of retailers' beliefs allows for a transition path to a collusive equilibrium. We define a weak PBE as formable if there is a transition path to a collusive equilibrium in the sense that, in this transition, each player maximizes its expected utility under the belief system of the weak PBE. Thus, formability is an additional property of an equilibrium that imposes a condition on the system of beliefs.

[13]For instance, the symmetric beliefs we consider in Section 3.5.1 violate Condition 1 but capture the idea that retailers may anticipate communication between manufacturers.

[14]See Section 3 in Rey and Vergé (2004).

[15]For any history $\mathcal{H}_t$ and Assumption 5, retailer expect a certain level of wholesale prices that results in a unique expected price level at the rival of p^e_{it}. We prove this in Proposition 1.

Definition 1. (Formable) A collusive weak PBE $\mathcal{C}$, consisting of a strategy profile and a system of beliefs, is formable if there exists a (transition) strategy profile, such that for the system of beliefs of $\mathcal{C}$ after an *arbitrary* history $\mathcal{H}_{t_0}$:

1. In the transition phase starting in period t_0 and lasting for x periods, with $x < \infty$, until t_0+x (exclusively), the strategy profile in the continuation game is sequentially rational for the system of beliefs of C;

2. The collusive phase starts in period $t_0 + x$. In that continuation game, the strategy profile of $\mathcal{C}$, adjusted with $t_0 + x$ as the first period, together with the system of beliefs of C, is a weak PBE.

Note that beliefs in the transition phase can be incorrect, such that the strategy profile and system of beliefs after the alternative history do not generally constitute a weak PBE.[16] The number of transition periods x can be long, such that our definition is not restrictive with regard to the speed of formation.

Next, we define how the concept of renegotiation-proofness relates to our model. Specifically, we rely on the notion of renegotiation-proofness by Farrell and Maskin (1989). We adapt the concept by restricting renegotiation to manufacturers. Renegotiation takes place implicitly in our model whenever *manufacturers* have a collective interest in revising their agreement. Formally, we hold retailers' beliefs and strategies constant when checking that there is no other equilibrium that Pareto-dominates it. This assumption is in line with other theories on renegotiation-proofness (see, for example, Bernheim et al., 1987 and Bernheim and Ray, 1989).[17]

Definition 2. A collusive weak PBE $\mathcal{C}$, consisting of a strategy profile and a system of beliefs, is *weakly* renegotiation-proof if it features opportunism-proofness and renegotiation-proof punishment:

1. (Opportunism-Proofness) A collusive equilibrium strategy for given beliefs (a weak PBE) is said to be opportunism-proof if the manufacturers, on the equilibrium path, do not benefit from jointly changing their contract offers.

2. (Renegotiation-Proof Punishment) A wPBE has the property of renegotiation-proof punishment if there is no punishment period in which the manufacturers would benefit if they jointly changed their contract offers.

[16]The definition, however, requires that the strategy profile of the continuation game starting with t_0 is perception perfect, i.e., sequentially rational given the beliefs.

[17]Our approach is line with the original idea presented by Farrell and Maskin (1989) who consider subgame perfect Nash equilibria (SPNE) in a collusive game and treat the buyers as non-strategic. Similarly, renegotiation-proofness is usually applied to the parties of a *relational contract* only, in our case the collusive agreement, e.g., Buehler and Gärtner (2013); Goldlücke and Kranz (2013). Alternatively, one can understand our approach as applying the exact concept of Farrell and Maskin (1989) to the reduced form SPNE taking retailers' strategies and beliefs from the wPBE as fixed.

The first condition excludes the scope for opportunistic manufacturer behavior during the collusive phase. Because demand does not change over time, manufacturers have no incentive to renegotiate the collusive price on the equilibrium path, but a coordinated one-shot deviation (against their retailers) may be profitable. Without opportunism-proofness, the implicit "collusive agreement" between the manufacturers is not renegotiation-proof. Whereas opportunism-proofness considers renegotiation-proofness of the collusive strategy on the collusive path, the criterion of renegotiation-proof punishment does so for the punishment phase after deviations. Intuitively, the condition states whether manufacturers want to negotiate the terms of the price war that follows after a unilateral deviation. Together, opportunism-proofness and renegotiation-proof punishment ensure *weak* renegotiation-proofness of the equilibrium analogously to the definition in Farrell and Maskin (1989). *Strong* renegotiation-proofness requires that not only all the continuation game equilibria do not invite joint deviation, but also that there is no other renegotiation-proof strategy profile that Pareto-dominates the candidate equilibrium of the whole game.

Discussion. In a standard model of horizontal collusion, the usual stability condition is necessary for collusion and also sufficient if the manufacturers can coordinate on the collusive outcome through a "meeting of the minds." If, however, colluding upstream firms sell to competing downstream firms, stability is not sufficient. Retailers might not accept the new contract offers with collusive prices, especially if the contract offers to downstream rivals are secret. Our formation condition captures the feasibility of the manufacturers to implement a collusive agreement vis-a-vis the downstream firms. For example, in the coffee cartel, collusion among the coffee roasters turned out to be unsuccessful because their retailers did not accept the collusive wholesale prices (Holler and Rickert, 2022). Collusion only became successful after introducing resale price maintenance.

Furthermore, we consider renegotiation of the collusive strategy and the option to renegotiate after a deviation. The collusive strategy is often viewed as all or nothing and cannot be adjusted. We show that manufacturers might want to adjust their collusive strategy (against their retailers) in a stable market environment, which we refer to as opportunism. Additionally, we check whether manufacturers can and want to renegotiate in the punishment phase after a deviation. As stated by Levenstein and Suslow (2006), "one of the most clearly established stylized facts is that cartels form, endure for a period, appear to break down, and then re-form again". They report that cheating and disagreement in cartels happens quite often but cartels frequently have multiple episodes of cooperation. Hence, it might be that cartel members renegotiate during price wars to return to the collusive outcome. As referenced in Levenstein and Suslow (2006), empiricists are able to differentiate between bargaining price wars and price wars that are part of a punishment strategy. We will show that collusion can be sustainable while not being renegotiation-proof and while being renegotiation-proof, explaining both kinds of price wars.

3.4 Solution

3.4.1 History-Independent Beliefs

In this section, we consider retailer beliefs that are independent of the history of actions in the game. Such beliefs arise naturally when the manufacturers' equilibrium strategies are stationary. For example, the competitive equilibrium in which both manufacturers set the competitive wholesale prices in each period is consistent with retailers having a constant and thus history-independent belief about wholesale prices. With these history-independent beliefs, we show that collusion is infeasible.

Definition 3. (History-Independent Beliefs) The price expectation p^e_{-it} of retailer i in period t about the price of retailer $-i$ is independent of the history of the actions in the game up to period $t-1$ and independent of the offer (w_{it}, F_{it}) made by its supplier in period t.

The proposed belief refinement above defines the retailers' out-of-equilibrium beliefs about the retail price of their competitor to be independent of the history. As discussed before, the fact that beliefs are passive within a period follows from Condition 1. The definition, however, does not impose a restriction on how the beliefs react to the past play of actions.

We first solve for the equilibrium of the game that results when both manufacturers maximize stage-game profits. In the last stage within each period, on the equilibrium path, each retailer has accepted the contract, but the rival's wholesale price remains secret. Each retailer i faces the own wholesale price w_{it} and holds a history-independent belief p^e_{-it} about the retail price of the rival. The retailers set the retail prices p_{it} simultaneously. Thus, the profit of retailer i is

$$\pi_{it}\left(p_{it}, p^e_{-it}\right) = \left(p_{it} - w_{it}\right) D_i\left(p_{it}, p^e_{-it}\right) - F_{it}.$$

Retailer i maximizes its profit with respect to p_{it}. The first-order condition is

$$D_i\left(p_{it}, p^e_{-it}\right) + \left(p_{it} - w_{it}\right) \frac{\partial D_i\left(p_{it}, p^e_{-it}\right)}{\partial p_{it}} = 0 \tag{3.1}$$

and defines retailer i's reaction function $p_i\left(w_{it}, p^e_{-it}\right)$. Anticipating the pricing outcome and resulting profits, each retailer decides whether to accept the wholesale tariff offered by its manufacturer in the second stage. Under Assumption 1, there exists a unique symmetric Nash equilibrium in the downstream market for offers (w_{it}, F_{it}) made by the manufacturers given retailers accept the offers.

In the first stage, the manufacturers offer their wholesale tariffs. Under manufacturer competition, each manufacturer i maximizes its profit, anticipating that its retailer sets a price of $p_i\left(w_{it}, p^e_{-it}\right)$:

$$\max_{w_{it},F_{it}} w_{it} \cdot D_i\left(p_i\left(w_{it}, p^e_{-it}\right), p_{-i}(w_{-it}, p^e_{it})\right) + F_{it},$$

subject to the retailer's participation constraint

$$\left(p_i\left(w_{it}, p^e_{-it}\right) - w_{it}\right) \cdot D_i\left(p_i\left(w_{it}, p^e_{-it}\right), p^e_{-it}\right) - F_{it} \geq 0. \tag{3.2}$$

The profit is determined by two parts. The variable part consists of the units sold times the unit wholesale price, and the fixed part consists of an up-front payment from the retailer to the manufacturer. The maximum fixed payment cannot be larger than the revenue that the retailer earns. This participation constraint binds in equilibrium.

Consider the one-shot game. In the one-shot equilibrium, the belief of the retailer about the retail price of the rival coincides with the expectation of the manufacturer about the retail prices. This implies $p_{-i}(w_{-it}, p^e_{it}) = p^e_{-it}$.

The optimality condition of the manufacturer problem implies that the only solution has $w_{it} = 0$. Denote the competitive wholesale price by w^P and the resulting retail price by $p^P = p_i(0, p^P)$. We have the following result:

Lemma 1. *With history-independent beliefs, in the equilibrium with competing manufacturers, the wholesale prices are equal to zero ($w^P = 0$), resulting in a competitive retail price level of p^P.*

The intuition behind the result is that, with selling the product at marginal costs of zero, each manufacturer ensures that retailers maximize the joint profits of a manufacturer-retailer pair, which can then be extracted via the fixed fee.

This history-independent benchmark defines the competitive level of wholesale prices which is at $w_i = 0$, the corresponding fixed fees, and the resulting retail prices. The fixed fee is equal to the retailer's profit, $F_i = \pi_i(p^P, p^P)$, corresponding to a non-colluding industry. In the next steps, we analyze the equilibrium of the repeated game if retailers hold history-independent beliefs. To show that firms are unable to collude on a price different from the competitive wholesale price of zero, we take the whole dynamic game into account. Manufacturers may collude using any dynamic strategy (for example, grim-trigger strategies). Each strategy, however, is a mapping from the previous history to an action chosen in period t.

Proposition 1. *Suppose retailers hold history-independent beliefs. Then, there exists a unique symmetric perfect Bayesian equilibrium in which the manufacturers set wholesale prices equal to zero and extract all profits of retailers via the fixed fee. Collusion cannot increase prices above the competitive level.*

This impossibility result can be interpreted as showing that the beliefs that support the competitive equilibrium cannot support a collusive equilibrium. Thus, for manufacturers in a competitive industry striving to start collusion, their joint will is not sufficient to increase prices above the competitive level. The strict impossibility result does not extend to linear tariffs but depends on efficient contracts (See Section 3.5.2).[18]

The intuition for the result is that the punishment strategy that underlies any collusive equilibrium becomes ineffective with these beliefs. Each manufacturer can ensure a profit larger than the collusive profit level that only depends on its own actions and the belief of its own retailer but is independent of the action of the other manufacturer. That payoffs are independent of the other manufacturer's action eliminates any credible punishment. Punishment, however, is necessary for collusion and thus there is an incentive to deviate from any supra-competitive price level.

Next we consider beliefs that are consistent with collusion. They are history-dependent and meet Condition 1, which implies that they are within-period passive.

3.4.2 Belief Dynamics

Consider beliefs that depend on the history of the game. We focus on the history of wholesale prices and disregard the history of retail prices, such that the retailers' beliefs cannot directly depend on the competing retailers' past actions. Otherwise, such beliefs could support retail collusion, which we want to abstract from.[19] Thus, the relevant history of the game in period t is $\mathcal{H}_t \equiv [(w_{A0}, w_{B0})\,;\ldots;(w_{At-1}, w_{Bt-1})]$, that is, the set of the pairs of wholesale prices that the manufacturers have set in all periods up to period $t-1$. More formally, for each retailer i, the belief is a function $p^e_{-it}(\mathcal{H}_t)$.

Grim-Trigger Beliefs

We construct history-dependent beliefs that have a grim-trigger property. For this, consider that manufacturers play grim-trigger strategies as described below.

[18] The classical opportunism problem can be mitigated under linear contract as shown by Gaudin (2019), where the monopolist achieves supra-competitive (marginal) wholesale prices. In contrast, with two-part tariffs, the monopolist sets the wholesale prices (w) equal to its marginal costs (McAfee and Schwartz, 1994).

[19] Suppose that each retailer believes that the price of the competitor is the monopoly price unless they have observed a different price in the past in which case they would believe that the price is the competitive price. Such beliefs could support a collusive action profile even with otherwise history-independent strategies.

We focus on industry-profit-maximizing collusion, such that $w^C = w^M$.[20] Any deviation by one manufacturer causes the other manufacturer to set the punitive price w^P forever. Because we are solving for perfect Bayesian equilibria, the punishment with w^P must be individually rational. This implies that $w^P = 0$ must hold in equilibrium, which follows from the logic presented in the proof of Proposition 1: Otherwise, each manufacturer would have an incentive to deviate from a $w^P \neq 0$ in a punishment period in which future beliefs and actions are fixed and unaffected by current actions. In punishment periods, each manufacturer must thus maximize short-term profits.

Collusive strategy. We consider collusion at the prices that maximize the integrated industry profit:

$$p^M = \max_{p_i} p_i \cdot D_i\left(p_i, p_{-i}\right) + p_{-i} \cdot D_{-i}\left(p_{-i}, p_i\right), \tag{3.3}$$

and define w^M implicitly through $p^M = p_i\left(w^M, w^M\right)$. Denote the maximal industry profit by

$$\Pi^M := p^M \cdot D_i\left(p^M, p^M\right) + p^M \cdot D_{-i}\left(p^M, p^M\right). \tag{3.4}$$

The optimal collusive wholesale price w^M and the respective belief about the retail price, p^M, are then determined by equation (3.3) as well. Thus, whenever the manufacturers collude at the industry-profit-maximizing wholesale prices w^M, each manufacturer earns a profit of

$$\Pi^C := \frac{\Pi^M}{2}.$$

The optimal collusive wholesale price maximizes the joint profit of the manufacturers given that the retailers' belief is identical to that price. The collusive profits are Pareto-optimal for the manufacturers if there is no joint deviation that is more profitable (see Definition 1). If weak renegotiation-proofness is fulfilled, then the equilibrium is also strongly renegotiation-proof because the profit on the equilibrium path would be Pareto-efficient for the colluding manufacturers.

Histories and beliefs. If manufacturers play grim-trigger strategies, their actions are only conditional on two kinds of histories: the collusive history $\mathcal{H}^{\mathcal{C}}$, where both manufacturers have only played w^C, and the deviation histories $\mathcal{H}^{\mathcal{D}}$ (any history other than $\mathcal{H}^{\mathcal{C}}$). Define the grim-trigger strategy as follows: Manufacturers set the collusive wholesale price w^Cin the first period. Then, in the t^{th} period, if both manufacturers have set the collusive price in each of the $t-1$ previous periods (history $\mathcal{H}^{\mathcal{C}}$), they set the collusive wholesale price w^C; otherwise, after histories $\mathcal{H}^{\mathcal{D}}$, manufacturers set a punishment price $w^P \neq w^C$

[20] Qualitatively, our results do not rely on the assumption that manufacturers collude on the monopoly price level. Collusion at the industry-profit maximizing level, however, facilitates renegotiation-proofness.

forever. Grim-trigger beliefs match these strategies by assigning two different beliefs to these histories. Grim-trigger beliefs $p_i^e(\mathcal{H})$ are thus history-dependent and differentiate between the two histories $\mathcal{H}^{\mathcal{C}}$ and $\mathcal{H}^{\mathcal{D}}$. Grim-trigger beliefs can be interpreted as the beliefs resulting from a subgame-perfect PBE with retailers fully anticipating manufacturers' strategies in every subgame.

Definition 4. (Grim-trigger beliefs)

- In the first period, retailers believe that the rival sets a retail price of p^C.
  - As long as both manufacturers play w^C, that is, the collusive history $\mathcal{H}^{\mathcal{C}}$ prevails, each retailer believes that the other retailer sets p^C in the current period.
  - Once one manufacturer has deviated, the history is $\mathcal{H}^{\mathcal{D}}$, and both retailers believe that the other retailer sets the competitive price p^P. This corresponds to a situation in which both retailers have common knowledge that the wholesale prices are $w^P = 0$ in the current period.

Given the passive nature of beliefs in a period (Condition 1), the beliefs are not correct in deviation periods. Neither a retailer that is offered the equilibrium contract updates its belief, nor updates the retailer that receives a deviating offer.

Because a deviation does not occur on the equilibrium path, the beliefs are nevertheless correct on the equilibrium path. Off the equilibrium path, the retailers' beliefs anticipate that manufacturers play grim-trigger strategies in that they also punish a deviation by one manufacturer in period t in all future periods.

Equilibrium. To determine an equilibrium of the dynamic game, we must consider deviations from the collusive strategy. In equilibrium, each manufacturer realizes per-period profits of $\Pi^C = \Pi^M/2$. In a deviation period, both retailers believe that the wholesale price is w^C and anticipate that the other retailer sets p^C. This results in a belief $p_{it}^e := p_i\left(w_{it}, p^C\right)$. Suppose that manufacturer i maximizes its deviation profit in period t in view of history $\mathcal{H}^{\mathcal{C}}$. When there is a deviation, that is, $w_{it} \neq w^C$ holds, grim-trigger beliefs imply that the level of w_{it} has no impact on future beliefs. The deviation profit is given by

$$\begin{aligned} & w_{it} \cdot D_i\left(p_i\left(w_{it}, p^C\right), p^C\right) + \left[p_i\left(w_{it}, p^C\right) - w_{it}\right] \cdot D_i\left(p_i\left(w_{it}, p^C\right), p^C\right) \\ & = p_i\left(w_{it}, p^C\right) \cdot D_i\left(p_i\left(w_{it}, p^C\right), p^C\right). \end{aligned} \tag{3.5}$$

Maximizing with respect to w_{it} yields the first-order condition

$$\begin{aligned}\frac{\partial p_i\left(w_{it},p^C\right)}{\partial w_{it}}D_i\left(p_i\left(w_{it},p^C\right),p^C\right)+\frac{\partial D_i\left(p_i\left(w_{it},p^C\right),p^C\right)}{\partial p_i}\frac{\partial p_i\left(w_{it},p^C\right)}{\partial w_{it}}p_i(w_{it},p^C)=0\\ \iff \frac{\partial p_{it}\left(w_{it},p^C\right)}{\partial w_{it}}\frac{\partial D_{it}\left(p_{it}\left(w_{it},p^C\right),p^C\right)}{\partial p_{it}}w_{it}=0.\end{aligned} \tag{3.6}$$

The last step follows from the first-order condition of the retailers in equation (3.1). Inserting the expression in the first line above yields the second line. Because the first two factors in equation (3.6) are non-zero, the manufacturer optimally deviates to $w^D = 0$. This results in

$$\Pi^D := p_i(0,p^C)\cdot D_i\left(p_i(0,p^C),p^C\right).$$

After any deviation by a manufacturer, the beliefs revert to $p^e_{-i} = p^P$ forever, that is, the belief in the punishment period. This results in profits of

$$\Pi^P := p^P \cdot D_i(p^P,p^P), \tag{3.7}$$

where $p^P = p_{it}(0,p^P)$ is the competitive price.

Collusion is sustainable when no manufacturer wants to deviate from the grim-trigger strategy. Using the one-shot deviation principle, the relevant incentive constraint for stability is

$$\frac{\Pi^C}{1-\delta} \geq \Pi^D + \frac{\delta\Pi^P}{1-\delta}. \tag{3.8}$$

The left-hand side contains the present value on the equilibrium path and the right-hand side the present value of a deviation. We can rewrite the incentive constraint for manufacturer i as follows:

$$\frac{p^C\cdot D_i(p^C,p^C)}{(1-\delta)} \geq p_i(0,p^C)\cdot D_i\left(p_i(0,p^C),p^C\right)+\frac{\delta}{1-\delta}\left(p^P\cdot D_i(p^P,p^P)\right). \tag{3.9}$$

This inequality is equivalent to the incentive constraint for standard horizontal collusion (when manufacturers and retailers are pairwise integrated), where the following order holds: $\Pi^D > \Pi^C = \Pi^M/2 > \Pi^P$.

Let us check whether the equilibrium is opportunism-proof (see Definition 1). Jointly, the manufacturers may have an incentive to reduce their wholesale prices for given beliefs. Suppose $\mathcal{H}_C$ is the history of the game, such that each retailer believes that its competitor sets p^C and anticipates to set $p_i\left(w_{it},p^C\right)$. First, we show that jointly deviating manufacturers set a wholesale price of $w_{it} = w_{-it} < 0$ to maximize spot profits. To see this, let us inspect the profit per manufacturer in the case of a joint deviation:

$$\begin{aligned}\Pi^{JD} :=& \frac{1}{2} \max_{w_{it}, w_{-it}} w_{it} \cdot D_i \left(p_i(w_{it}, p^C), p_{-i}(w_{-it}, p^C)\right) + \left[p_i(w_{it}, p^C) - w_{it}\right] \cdot D_i \left(p_i(w_{it}, p^C), p^C\right) \\ &+ w_{-it} \cdot D_{-i} \left(p_{-i}(w_{-it}, p^C), p_i(w_{it}, p^C)\right) + \left[p_{-i}(w_{-it}, p^C) - w_{-it}\right] \cdot D_{-i} \left(p_{-i}(w_{-it}, p^C), p^C\right).\end{aligned}$$

We assume that this profit is quasi-concave, such that we can use first-order conditions.[21] We rewrite the first-order conditions in equation (3.1) by applying symmetry. This is possible because we assume that the manufacturer profits are well behaved in the sense that the optimal joint action of the manufacturers is symmetric. This yields

$$\left[\underbrace{D_i \left(p_i(w, p^C), p_{-i}(w, p^C)\right) - D_i \left(p_i(w, p^C), p^C\right)}_{<0}\right] + w \underbrace{\frac{\partial p_i}{\partial w_{it}}}_{>0} \left[\underbrace{\frac{\partial D_i}{\partial p_i} + \frac{\partial D_{-i}}{\partial p_i}}_{<0}\right] = 0, \forall i. \tag{3.10}$$

Equation (3.10) only holds for $w < 0$. Hence, manufacturers optimally deviate to $w < 0$ jointly. A manufacturer makes a higher profit in the case of a joint deviation than when deviating unilaterally: $\Pi^{JD} > \Pi^D$. To explain the last inequality, note that the manufacturers could replicate the profit of Π^D for each of them by setting $w = 0$. The manufacturers, however, optimally set a price w below zero because this yields strictly larger profits. Setting a price below zero has a negative externality on the rival retailer that they do not internalize; it enables them to profitably exploit the incorrect beliefs. Following a deviation, the manufacturers make profits Π^P in future periods due to the grim-trigger beliefs. This results in the following opportunism-proofness condition

$$\frac{\Pi^C}{1-\delta} \geq \Pi^{JD} + \frac{\delta \Pi^P}{1-\delta}. \tag{3.11}$$

Comparing equation (3.11) with equation (3.8), the only difference is the deviation profit on the right-hand side. Because $\Pi^{JD} > \Pi^D$, the condition for opportunism-proofness is harder to satisfy than the stability condition above.

Proposition 2. *With grim-trigger beliefs, there exists an equilibrium in which manufacturers are able to sustain collusion on the industry-profit-maximizing wholesale price using grim trigger-strategies if the discount factor is high enough to satisfy equation* (3.9)*. This condition is equivalent to the incentive constraint when manufacturers and retailers are pairwise integrated, that is, under horizontal collusion. Under the stronger condition* (3.11)*, collusion is also opportunism-proof. The punishment is renegotiation-proof, but collusion is not formable.*

[21] Note that this holds for linear demand. In general, quasi-concavity depends on higher-order derivatives of demand at different loci of demand.

First, in contrast to Proposition 1, manufacturers are able to sustain collusion when retailers hold grim-trigger beliefs. Intuitively, grim-trigger beliefs allow for punishment, which is ineffective with history-independent beliefs. Moreover, the equilibrium above satisfies subgame perfection, that is, retailers correctly anticipate the manufacturers' actions in information sets that are reached with zero probability. Retailers have consistent beliefs on the equilibrium path in the collusive phase and in punishment phase subgames, as grim-trigger beliefs mirror the manufacturer's grim-trigger strategies. Second, as derived above, the opportunism problem gives rise to the incentive condition (3.11) that is harder to satisfy (that is, only for more patient firms) than condition (3.8) for stable collusion. Hence, opportunism can make cartels less sustainable. Similar to the literature on the opportunism problem in single-shot games, this result depends on the beliefs. Due to their passive beliefs, the retailers do not react immediately to opportunism, which allows the manufacturers to "trick" the retailers because the price that retailers expect will turn out to be incorrect if manufacturers jointly deviate. This joint deviation differs from a unilateral deviation of one manufacturer-retailer pair from a candidate equilibrium in which instead the belief of the deviating retailer is correct. In the latter case, it is optimal for the deviating manufacturer to set the wholesale price equal to the true costs. In the former case of a joint deviation, each retailer wrongly believes that the other retailer will buy at a high wholesale price and will thus sell at a high price, such that demand is high. It is, hence, profitable for the manufacturer to set a high fixed fee in return for a marginal wholesale price below costs because the retailer believes that it can sell a large quantity. A marginal wholesale price below costs, however, only becomes profitable for a manufacturer when the retailer has a wrong belief. It could thus signal the retailer that its belief is wrong because, if it were correct, the manufacturer's offer would be dominated by an offer with a wholesale price of $w = 0$. The stricter condition for opportunism-proofness is thus an implication of the passive nature of the beliefs. We later consider symmetric beliefs, where retailers' beliefs instantly react to any change in wholesale prices.

If marginal wholesale prices below marginal costs are impossible (for instance, due to competition law), the incentive constraint for opportunism and for stability are identical, implying that any stable cartel is also opportunism-proof. In any case, the opportunism problem is mitigated under grim-trigger beliefs compared to the one-shot game, or history-independent beliefs, because retailers with grim-trigger beliefs react to the deviation from the expected cartel wholesale prices by adapting their beliefs from the next period on to the level of competitive prices. This effectively punishes manufacturer opportunism, such that joint deviations are not profitable when manufacturers are sufficiently patient.

Interestingly, however, we find that under grim-trigger beliefs, collusion is *not* formable. Because these beliefs are unforgiving, they are beneficial in supporting the collusive equilibrium, but do not allow for new cycles of collusion, even after a long time.

The characterized equilibrium is also renegotiation-proof in the sense of (Farrell and Maskin, 1989). Because the opportunism-proofness condition (3.11) ensures that manufacturers have no incentives to deviate jointly from the collusive agreement, and credible punishment is fulfilled as well, the equilibrium is *weakly* renegotiation-proof. Because the equilibrium features collusion on the industry-profit-maximizing prices that are Pareto-optimal for the manufacturers, the equilibrium is also *strongly* renegotiation-proof.

Corollary 1. *The equilibrium with grim-trigger beliefs described in Proposition 2 is strongly renegotiation-proof if the opportunism-proofness condition* (3.11) *is satisfied.*

Note that this corollary extends naturally to general trigger-strategies and more general beliefs as employed in the next section. Recall that the concept of renegotiation-proofness is applied to the coalition of manufacturers; but the result may be more general if retailers are allowed to be part of the negotiation, as long as beliefs are not affected. Recall that we focus on collusion in which retailers do not take part. Fixing beliefs, however, manufacturers leave zero profits to retailers, making them effectively indifferent in renegotiation, such that a renegotiation-proof equilibrium would also be renegotiation-proof when including retailers.

For reference, consider a standard model of horizontal collusion absent a vertical dimension. In our model, this arises when each manufacturer-retailer pair is vertically integrated. Suppose that firms collude with grim-trigger strategies at the industry-profit-maximizing level and the usual stability condition is met. Such a collusion is formable because switching to grim-trigger strategies is an equilibrium of the continuation game independent of the history of actions in the game. Collusion is then also opportunism-proof because joint profits are maximized. In this industry structure, punishment is not renegotiation-proof because the firms would prefer to renegotiate to return to collusive prices. Hence, there are no renegotiation-proof equilibria in the case of standard horizontal collusion. By contrast, we will show that secret contracting can feature renegotiation-proof punishment and that opportunism-proofness can impose additional conditions on collusion.

Trigger Beliefs

In this section, we define trigger beliefs and derive the conditions for the sustainability and for other properties of collusion. First, consider that the manufacturers play trigger-strategies similar to the grim-trigger strategy above, but a deviation is forgiven after κ periods.

We define trigger-strategies as in Green and Porter (1984): Let κ denote a time length measured in periods.

Define period t to be collusive if

(a) $t = 0$, or

(b) $t - 1$ was collusive and $w_{it-1} = w^C$ for all i, or

(c) $t - \kappa$ was collusive and $w_{it-\kappa} \neq w^C$ for some i.

Define t to be reversionary otherwise. Manufacturer i sets

$$w_{it} = \begin{cases} w^C & \text{if } t \text{ is collusive,} \\ w^P & \text{if } t \text{ is reversionary.} \end{cases}$$

Assume that the collusive price equals the monopoly price ($w^C = w^M$), as defined in Subsection 3.4.2. Any deviation by one manufacturer causes the other manufacturer to carry out a punitive action of $w^P = 0$. Again, $w^P = 0$ must hold because a deviation during the reversionary periods is not punished as the future actions and beliefs are fixed. Hence, the punishment action must be the same as the short-term optimal action, that is $w^P = 0$ as we demonstrated before. This maximizes the manufacturer profits because it aligns the incentives of retailer and manufacturer. Similar to the previous section, trigger beliefs emerge when considering a wPBE in which manufacturers use trigger strategies instead of grim-trigger strategies.

Definition 5. (Trigger beliefs)

Choose a collusive price level p^C and a punishment price level p^P. Formally, the retailers' beliefs correspond to the manufacturer strategies for collusive and reversionary periods as defined above:

$$p^e_{-it} = \begin{cases} p^C & \text{if } t \text{ is collusive,} \\ p^P & \text{if } t \text{ is reversionary.} \end{cases}$$

When both manufacturers play trigger-strategies with actions w^C and w^P, the retailers' corresponding trigger-beliefs are correct on the collusive equilibrium path. They are also correct in the punishment phase.

Qualitatively, our insights for grim-trigger strategies carry over to more general trigger strategies. Trigger strategies with limited punishment imply stricter conditions for stability and opportunism-proofness compared to grim-trigger strategies. Nevertheless, they may be attractive for very relevant reasons that we do not model, including cost and demand shocks as well as other uncertainty that could result in unwarranted punishment, which – in case of grim-trigger strategies – would be very costly. We summarize our findings in the following proposition:

Proposition 3. *With trigger beliefs, there exists an equilibrium in which manufacturers are able to sustain collusion on the industry-profit-maximizing wholesale price (with trigger strategies) if the discount factor is high enough to satisfy condition* (3.21)*. This condition is equivalent to the incentive constraint when manufacturers and retailers are pairwise integrated. Furthermore, the condition requires more patience than under grim-trigger beliefs and strategies. Only under the stricter Condition* (3.22)*, collusion is also opportunism-proof. Punishment is always renegotiation-proof, and collusion is never formable.*

The derived equilibrium with trigger beliefs is Pareto-efficient as well as subgame perfect and features renegotiation-proof punishment. If, in addition, the manufacturers are patient enough for the condition of opportunism-proofness to hold, the equilibrium is strongly renegotiation-proof.

Corollary 2. *The equilibrium with trigger beliefs described in Proposition 3 is strongly renegotiation-proof if the opportunism-proofness condition* (3.22) *is satisfied.*

Adaptive Beliefs

After analyzing beliefs that anticipate collusion with (grim-)trigger strategies, we provide a simple example of a belief system where retailers are initially agnostic whether manufactures collude or compete. These beliefs try to capture that retailers learn heuristically about future wholesale price offers by adapting to past observed behavior. Imposing Condition (1) implies that retailers do not revise their belief within a period, but they 'learn' afterwards from manufacturers' observed behavior.We introduce adaptive beliefs to analyze equilibria with (grim-)trigger strategies. We characterize the beliefs in a way that they only depend on the actions of the manufacturers in the last T periods and not on the full history of the game.

Definition 6. (Adaptive Beliefs)

Beliefs are passive within each period t, that is, w_{it} does not affect the belief p^e_{-it}. Beliefs are dynamic: p^e_{-it} can depend on the history of past wholesale prices. There are three relevant histories:

1. In period t, the manufacturers offer contracts that are identical to the ex ante beliefs in period t. Then, both retailers retain the same belief in period $t+1$.

2. In period t, both manufacturers play the same $w \in W^*$ that differs from the ex ante beliefs of the retailers in period t. The same holds for all previous periods up to $t-(T-1)$, with $T \in \{1, 2, ...\}$ being a parameter measuring the adaptation length in periods. In $t+1$, both retailers hold the new (passive) belief p^*. The set W^* contains wholesale prices that the retailers accept as possible equilibria (for example, the collusive wholesale price w^C).

3. In period t, at least one of the manufacturers does not play a price consistent with a retailer's ex ante belief, and it is not the case that both manufacturers play a price $w \in W^*$. In $t+1$, both retailers hold the new belief p^P, where p^P is the wholesale price of the perfect Bayesian Nash equilibrium of the stage game (Nash reversal).

The definition of adaptive beliefs differs from the definition of (grim-) trigger beliefs in that we allow retailers to take on any new belief $w \in W^*$ after observing w for sufficiently many periods (see point 2). The remaining requirements (point 1 and 3) are also present for (grim-) trigger beliefs.

In the following, we use the example of grim-trigger strategies to analyze the necessary incentive constraints implied by opportunism-proofness, renegotiation-proof punishment, and formability of collusion that arise from considering adaptive beliefs. We also solve for trigger strategies and present the results in the proposition, relegating the exposition to the proof.

Stability. The stability condition for collusion, once established, is

$$\frac{\Pi^C}{1-\delta} \geq \Pi^D + \frac{\delta \Pi^P}{1-\delta}. \tag{3.12}$$

In this case, retailers with adaptive beliefs already have the belief p^C and revert to the belief p^P after a deviation. Hence, with grim-trigger strategies, the condition is the same as the incentive condition (3.8). As a consequence, collusion is stable if manufacturers are sufficiently patient (δ is large enough). Again, the condition is identical to the stability condition under horizontal collusion.

Formability. Formability is fulfilled if there exists a transition path to the collusive wPBE, where – on the path – the manufacturers' actions are mutually best responses. Suppose that, starting from any history in period s, the manufacturers both start to play the collusive wholesale price w^C with the usual grim-trigger strategies that punish any deviation. If the manufacturers want to form collusion, the worst history in terms of our beliefs is that there was competition in $s-1$ (with profits Π^P). This implies that, for the next T periods, the retailers' beliefs are fixed at p^P, such that the manufacturers can extract lower transfers. Hence, in the transition periods, the manufacturer profits are lower than under competition, $\Pi^F < \Pi^P$. The reason is that the beliefs are identical in both cases, and whereas the manufacturers play their unique best response to the belief in periods of competition, which results in Π^P, they play a worse action with respect to stage-game profits as response to the same belief in transition periods, resulting in a profit of Π^F.

With grim-trigger strategies, manufacturers have an incentive to jointly start to collude in a competitive period, such that they eventually arrive at the collusive equilibrium path if

$$\frac{\Pi^P}{1-\delta} \leq \frac{1-\delta^T}{1-\delta}\Pi^F + \frac{\delta^T}{1-\delta}\Pi^C. \tag{3.13}$$

The left-hand side contains the present value of perpetual competition. The first term on the right-hand-side is the discounted profit of the T periods in which retailers' belief is p^P, while manufacturers actually set w^C; the second term is the discounted profit of perpetual collusion starting after T formation periods.

Because the manufacturer profits are lower during the formation period than under competition, each manufacturer may have an incentive to deviate to a lower wholesale price during formation. Deviating during the formation phase yields a period profit of $\Pi^{F,D}$ but triggers a punitive action forever. Consider the incentives to stick to w^C in the formation phase. No manufacturer wants to deviate unilaterally in the formation phase, which implies that actions in the transition are mutually best responses, if

$$\frac{1-\delta^T}{1-\delta}\Pi^F + \frac{\delta^T}{1-\delta}\Pi^C \geq \Pi^{F,D} + \frac{\delta\Pi^P}{1-\delta}. \tag{3.14}$$

Comparing the inequalities (3.13) and (3.14) shows that the latter is stricter if and only if $\Pi^{F,D} \geq \Pi^P$. This is always the case, as we demonstrate in the proof of Proposition 4. Hence, the deviation condition (3.14) is not only necessary but also sufficient for formability. Note that formability decreases in T, that is, it holds for a smaller set of discount factors because the left-hand side decreases in T. This holds because an increase in T, which only affects the left-hand side of condition (3.14), shifts the weight Π^C to the smaller term Π^F.

Renegotiation-proofness. Collusion is opportunism-proof if the manufacturers have no incentive to deviate jointly from the collusive price. Suppose the manufacturers jointly behave opportunistically in the present period. They can earn an opportunism profit of Π^{JD} by lowering w_i and increasing F_i for each retailer. As a result, the retailers believe in competition in the next period. Confronted with competitive beliefs, the reformation phase starts so that the manufacturers need to play w^C for T periods to convince retailers of collusive prices again. This yields the following condition of opportunism-proofness:

$$\frac{\Pi^C}{1-\delta} \geq \Pi^{JD} + \delta\left(\frac{\left(1-\delta^T\right)}{1-\delta}\Pi^F + \frac{\delta^T}{1-\delta}\Pi^C\right). \tag{3.15}$$

Collusion is formable and opportunism-proof if the manufacturers are sufficiently patient, that is, if condition (3.14) for formability and condition (3.15) for opportunism-proofness hold. Increasing the adaptation length T of the beliefs makes collusion less formable but more opportunism-proof, that is, relaxes condition (3.15), but tightens condition (3.14). Opportunsim-proofness is harder to satisfy than stability whenever collusion is formable.

To see this, compare the right-hand side of condition (3.15) with the stability condition (3.12) and note that $\Pi^{JD} > \Pi^{D}$.

Punishment is not renegotiation-proof whenever collusion is formable. To see this, note that whenever the formability condition holds, and the manufacturers are supposed to punish, they are better off forming collusion again. This implies that formation dominates competitive pricing in the punishment phase and manufacturers prefer to enter the cooperative phase again. Note that grim-trigger strategies are a special case of trigger strategies with $\kappa \to \infty$. For trigger strategies, we generalize the result in the following proposition focusing on collusion at industry-profit maximizing wholesale prices:

Proposition 4. *With adaptive beliefs and trigger strategies, the stability condition is the same as for vertically integrated collusion with the equivalent trigger strategies. Stability increases in the number of punishment periods κ. Collusion is formable and opportunism-proof if δ is sufficiently large. Increasing the adaptation periods T of the retailer beliefs makes collusion less formable but more opportunism-proof. Punishment is not renegotiation-proof for any $\kappa > 0$ whenever collusion is formable, but collusion may be sustained even without punishment, that is, for $\kappa = 0$.*

The take-away is that there exists an equilibrium with adaptive beliefs in which collusion is formable, sustainable, and opportunism-proof. Formability requires that beliefs can adapt to collusion after a 'history' of competition or punishment. The adaptivity can have the cost that punishment is not renegotiation-proof. This is different from trigger-beliefs where retailers do not accept collusive contracts after a deviation. Instead, with adaptive beliefs they do so, such that, whenever an adaptive belief features formability, punishment is not renegotiation-proof. An interesting observation is that, for adaptive beliefs, there is a trade-off between opportunism-proofness and formability: The longer beliefs take to adapt the harder it is to start collusion, whereas opportunistic behavior that counts on restarting collusion becomes less of a problem. This moreover implies that an equilibrium with adaptive beliefs is not subgame perfect in constrast to an equilibrium with (grim-) trigger beliefs.

As a polar case, we find that the strategy to always collude, the degenerate trigger strategy with a punishment length of zero, can support a collusive equilibrium. If manufacturers play "always collusion", the incentive constraint for stability is

$$\frac{\Pi^C}{1-\delta} \geq \Pi^D + \delta\left(\frac{1-\delta^T}{1-\delta}\Pi^F + \frac{\delta^T}{1-\delta}\Pi^C\right). \tag{3.16}$$

Note that the stability is supported by the beliefs because deviation requires a new formation of collusion of length T. While this strategy is the least stable strategy, it is the only strategy that features formability and, in addition, features renegotiation-proof punishment.

Corollary 3. *For $\kappa > 0$, there exists a strongly renegotiation-proof equilibrium that is not formable if T and δ are sufficiently large. For $\kappa > 0$ and if the formability conditions* (3.13) *and* (3.14) *hold, there is no renegotiation-proof equilibrium with trigger strategies. For $\kappa = 0$, a strongly renegotiation-proof equilibrium exists if the opportunism-proofness condition* (3.15) *holds (which implies stability).*

3.5 Extensions

3.5.1 Symmetric Beliefs

After considering passive beliefs, we turn to the analysis of the case that retailers have symmetric beliefs defined as follows:

Definition 7. (Symmetric Beliefs) The price expectation p^e_{-it} of retailer i in period t about the price of retailer $-i$ is $p_i\left(w_{it}, p(w_{it})\right)$.

In other words, when a retailer receives an unexpected offer deviating from the candidate equilibrium, the retailer revises its belief and believes that its rival has received the same offer by its manufacturer. Symmetric beliefs are history-independent because they only rely on the information contained in the current wholesale price offer. We assume that manufacturers play (grim-)trigger as in the previous sections. The case of symmetric beliefs is also analyzed in Liu and Thomes (2020) who consider the linear demand case and, hence, offer closed-form solutions for the critical discount factor. Symmetric beliefs may be particularly plausible if retailers expect that manufacturers coordinate their actions and expect that they do this in a symmetric way.

We focus again on industry-profit-maximizing collusion that naturally arises when manufacturers jointly maximize their profits given symmetric beliefs. Denote the price expectation of retailer i with symmetric beliefs by p^e_{it} as above. This allows the manufacturer to essentially choose the symmetric price level, such that the joint maximization problem of the manufacturers can be rewritten as

$$\Pi^C := \frac{1}{2}\max_p p \cdot D_i\left(p, p\right) + p \cdot D_{-i}\left(p, p\right) = \Pi^M.$$

Hence, the joint-profit-maximizing wholesale price of the manufacturers is equal to the industry-profit-maximizing price. Moreover, this implies that any joint deviation by the manufacturers will always be the industry-profit-maximizing price if retailers hold symmetric beliefs. In contrast to (grim-)trigger beliefs, where the optimal joint deviation of the manufacturers is to charge a wholesale price below zero, such a deviation is not optimal because only the collusive price maximizes the manufacturers' profits. Because the manufacturers collude at the Pareto-efficient level, collusion is opportunism-proof. Additionally, punishment is not renegotiation-proof because the manufacturers would prefer to revert back to collusion in every punishment period.

Collusion is formable because symmetric beliefs instantly adapt to the new wholesale price in every period. Manufacturers only need to agree on the collusive price and set it in any period. In a period in which the manufacturers set wholesale prices of w^C, the retailers' expectations are immediately equal to $p_i\left(w^C, p(w^C)\right) = p_i(w^C, p^C) = p^C$, which corresponds to the expectation of collusion. Forming collusion immediately leads to stable collusion as long as the stability condition is fulfilled.

From Pagnozzi and Piccolo (2011), we know that symmetric beliefs affect competition between vertically separated manufacturers. Competition is less fierce due to a so-called belief effect, which increases the competitive wholesale price above marginal costs. If punishment, however, relies on the competitive wholesale prices and profits, the stability of collusion is affected. Manufacturers must be more patient to satisfy the condition of stable collusion.

Symmetric beliefs violate the "no-signaling-what-you-don't-know" Condition 1.[22] We thus look for a perfect Bayesian equilibrium in the context of symmetric beliefs:

Proposition 5. *With symmetric beliefs, there exists a collusive equilibrium with (grim-) trigger strategies if condition 3.27 holds. Collusion is also formable if the stability condition holds. Collusion is always opportunism-proof, but punishment is not renegotiation-proof. If condition 3.27 holds, both the competitive equilibrium and the collusive equilibrium are not renegotiation-proof.*

With symmetric beliefs there is no opportunism problem because lowering the wholesale prices negatively affects the belief such that retailers do not accept collusive contracts. Because the symmetric belief follows the manufacturers' actions, there is no long-term response of retailers and thus retailers do not discipline the manufacturers as, for example, the retailers' response does when they are holding trigger beliefs. This implies that formation, in the sense of needing to convince retailers, is not an issue with symmetric beliefs. As a downside for collusion, symmetric beliefs do not support the renegotiation-proofness of punishment because they make the punishment phase prone to renegotiation incentives. The reason is that the retailers always accept the collusive contracts, believing that wholesale prices are at the collusive level in the whole industry. During a punishment phase both manufacturers are able to renegotiate a joint increase of their wholesale prices again, which in turn will be accepted by each retailer. Therefore, manufacturers are able to return to the Pareto-dominant collusive equilibrium. This intuition also applies to the competitive equilibrium given that the manufacturers discount factor is sufficiently large to potentially support collusion.

[22]Pagnozzi and Piccolo (2011) discuss this observation and also show that, with common cost shocks, symmetric beliefs can be consistent with a PBE.

3.5.2 Linear Wholesale Prices

For this extension we restrict the manufacturers' contracts to be linear. Each manufacturer offers a contract that only includes a wholesale price w_i.

Let us reconsider the case of history-independent beliefs. For two-part tariffs we have shown that there is no collusive equilibrium. In contrast, for linear tariffs we construct grim-trigger equilibria and show that for sufficient patience an equilibrium with collusive prices above the competitive price level always exists. Let us first define the competitive equilibrium. The one-shot wholesale prices with correct beliefs $p^N \equiv p^*(w^N, w^N)$ are defined by

$$w^N = \arg\max_w \; wD(p^*(w, p^N), p^*(w^N, p^N)).$$

This results in the competitive profits $\Pi^N = \max_w wD(p^*(w, p^N), p^N)$. Playing w^N in each period and retailers believing p^N is a wPBE of the repeated game.

Consider that manufacturers collude symmetrically at wholesale price of w^C using grim-trigger strategies. We denote the wholesale prices in punishment periods by w^P. A deviation manufacturer chooses a wholesale price denoted by w^D.

If there exists an equilibrium with a collusive price level of w^C, the beliefs have to be correct on the equilibrium path and equal $p^e = p^*(w^C, p^e)$. We construct such an equilibrium with time-constant beliefs to show the existence of a collusive equilibrium. If manufacturer A deviates in a period, the belief of the retailer of the non-deviating manufacturer B are still incorrectly at p^e. Similarly, in punishment periods, both retailers may hold incorrect beliefs, expecting the price to be at the collusive level, whereas the actual price level depends on w^P that is generally not equal to w^C. Different to the case of two-part tariffs, wrong retailers do not render collusion impossible with linear tariffs as the retailers mechanically accept different wholesale price levels, so that the effect of the wholesale prices on manufacturer profits are maintained.

Let us define the profits necessary to evaluate the stability of collusion:

$$\Pi^C \equiv w^C D(p^e, p^e),$$

$$\Pi^D = \max_w w \cdot D(p^*(w, p^e), p^e).$$

The punishment prices being mutual best responses yields

$$\begin{aligned} w^P &= \arg\max_w w \cdot D(p^*(w, p^e), p^*(w^P, p^e)), \\ \Pi^P &= w^P \cdot D(p^*(w^P, p^e), p^*(w^P, p^e)), \end{aligned}$$

where $p^*(w^P, p^e) < p^e$ holds for any $w^P < w^C$.

Assumption 6. *We assume that at the locus at $w^C = w^D = w^P = w^N$, the implicitly defined wholesale prices w^N ,w^D, and w^P are unique and $0 < \frac{\partial w^D(w^C)}{\partial w^C} < 1$ and $0 < \frac{\partial w^P(w^C)}{\partial w^C} < 1$ hold.*

This assumption holds globally for linear demand. The assumption implies that the wholesale prices have the usual ordering in a neighborhood around the competitive price level: If $w^C > w^N$, that is, if there is effective collusion, then $w^C > w^D > w^P > w^N$ holds (see proof of Proposition 6).

An equilibrium with grim-trigger strategies exists if there is a critical discount factor below one, such that the stability condition holds. This can be rewritten as

$$\Pi^D - \Pi^C < \Pi^C - \Pi^P. \tag{3.17}$$

Proposition 6. *Suppose that retailers hold history-independent beliefs and Assumption 6 holds. With linear wholesale prices, there is a collusive equilibrium with $w^C > w^N$ if manufacturers are sufficiently patient. This collusive equilibrium is also formable.*

While collusive equilibria with linear tariffs and prices just above the competitive price level exist if manufacturers are sufficiently patient, linear tariffs reduce the ability of manufactures to extract retail rents. Thus, overall, linear tariffs do not necessarily result in higher profits for manufacturers than two-part tariffs.

3.6 Policy Discussion

Our results provide new insights for various business practices from the perspective of competition policy.

Resale price maintenance (RPM). There is an ongoing policy debate about RPM that is typically considered anti-competitive in European competition policy and more benign in the US, at least since the Supreme Court's Leegin decision.[23] Our theory suggests that RPM can be pro-collusive by helping colluding manufacturers to convince retailers that competing products will have high prices. For this it is necessary that RPM is not just implemented as a secret contract clause but rather communicated in a way that competing retailers become aware of it. This may happen through the trade press, press releases, or other announcements – at least if RPM is legal. If it is illegal, recommended retail prices, which may for instance be printed on the products' price tags, or so-called minimum advertised price (MAP) restrictions can serve a similar purpose if the manufacturers incentivize the retailers to not deviate from them.[24]

[23]See Leegin Creative Leather Products, Inc. v. PSKS, Inc., 551 U.S. 877 (2007).

[24]See Asker and Bar-Isaac (2020) for a review of MAP restrictions. They also find that MAP facilitates collusion among manufacturers, where our model adds an additional channel, namely to overcome the opportunism problem towards retailers.

Vertical integration. Vertical integration is often viewed as beneficial by solving various coordination issues within supply chains, such as that of double marginalization. A major concern, however, is market foreclosure (Rey and Tirole, 2007). Our theory suggests that vertical integration can facilitate collusion because retailers no longer need to be convinced of higher competing prices, which, as we show, may fail under various retailer beliefs. A potential downside of vertical integration is that, in certain cases, punishments for a deviation from collusion may not be renegotiation proof, although it can be for certain retailer beliefs under vertical separation (see Corollary 1).

Communication. Communication within supply chains is generally viewed to be beneficial to overcome coordination problems, whereas communication between competitors about sales prices and similar strategic variables is typically considered to be suspicious and may constitute a violation of the cartel prohibition.[25] Our theory suggests that communication among competing retailers about simultaneous new wholesale tariff offers of different manufacturers may help to turn competitive retailer beliefs into collusive beliefs, which facilitates manufacturer collusion. Hence, communication between retailers might not only facilitate collusion among them, but can also facilitate collusion in the upstream market.

Collusion through downsizing of packs. There are instances of coordinated downsizing of pack sizes by manufacturers, as observed in the chocolate case in Germany (Ritter). In their detailed analysis of a washing powder cartel in Europe, Laitenberger and Smuda (2015) report that the three firms involved engaged in various anticompetitive practices. As one such practice, the firms agreed on indirect price increases by keeping prices unchanged when, among others, the product volume or the number of wash loads per package was reduced. The collective reduction of the pack sizes while maintaining the old price of the larger pack may be used to reduce strategic uncertainty for retailers. The colluding manufacturers do not have to implement a price increase, which, depending on the retailer beliefs, may fail. Instead, it might be sufficient to show their retailers smaller packs of the competing brands to convince them that selling the smaller pack at the old price is competitive. If manufacturers present the retailers in a price negotiation with actually downsized packs of the own and a competing product, the claim may be more credible than the claim that other retailers face higher wholesale prices for competing products, especially if the adjustment of pack size is costly.

Buyback policies. Buyback policies can reduce the risk of retailers ending up with unsold units if their prices are not competitive. If retailers need to be convinced that the

[25] See the "Guidelines on the applicability of Article 101 of the Treaty on the Functioning of the European Union to horizontal co-operation agreements", European Commission, Commission Communication 2023/C 259/01.

higher (collusive) wholesale prices are competitive to accept the manufacturer's contract, a buyback policy that insures the retailer in case its offer is not competitive can help.

3.7 Conclusion

Our approach of explicitly characterizing belief systems and discussing the equilibria consistent with such beliefs differs from the usual logic of Nash equilibrium, where players know other players' strategies fully. In particular, we consider that retailers do not fully anticipate the collusive strategy of manufacturers. While retailers' beliefs are correct on the equilibrium path, that is, retailers predict the correct price level in any collusive equilibrium, the existence of such an equilibrium depends also on retailers playing their role in the punishment phase which is necessary to support the equilibrium. We analyze three problems for the existence and plausibility of manufacturer collusion. First, the collusive equilibrium may fail to exist due to a lack of effective punishment. Second, the collusive equilibrium may lack credibility if the implied agreement is not renegotiation-proof, either on the collusive equilibrium path or during a punishment phase. In this context, we find that an opportunism problem, like the one a monopolist with secret contracts faces, is recreated by colluding firms.[26] Third, a collusive equilibrium may be implausible if formation, such as the transition from a competitive market to a collusive market, is not possible given the retailers' reactions.

When retailer beliefs do not anticipate manufacturer collusion, we demonstrate that collusion is infeasible with two-part tariffs. Such beliefs may arise in industries that have long-standing competitive conduct. Belief differences between industries may explain why some industries stay competitive, whereas other industries give birth to collusion repeatedly. Because these beliefs give rise to equilibria of the infinitely repeated game and, consequently, are correct on the equilibrium path, they are self-fulfilling and may never be challenged.

When retailer beliefs react to observed past actions, trigger-based manufacturer collusion becomes feasible and the punishment may even be renegotiation-proof, although this would not be the case under vertical integration of the industry. We show that opportunism can still be the most important challenge for the colluding manufacturers, more so than the usual unilateral deviation incentives.

Trigger beliefs and particularly grim-trigger beliefs can feature renegotiation-proof punishment and opportunism-proofness because they make it difficult for manufacturers to obtain high wholesale prices again after a breakdown. However, they do not allow for the formation of collusion. Formation requires that, to transition from a competitive to a

[26]This highlights that communication among manufacturers may be a source of misinterpretation for retailers. This contrasts the understanding that communication usually reduces strategic uncertainty, as discussed by Blume (1994) and Blume and Heidhues (2008).

collusive outcome, beliefs must adapt to the change in contract offers. The intuition here is that trigger beliefs cannot handle the transition from a non-collusive state to successful collusion because they would be too "pessimistic" about the future whenever they observe a non-collusive price in the history of the game.

We introduce new beliefs that are adaptive to the manufacturers' pricing over time. These retailer beliefs facilitate the formation of manufacturer collusion. As the beliefs do not mirror the collusive strategy, they do not support renegotiation-proof punishment; they may still support a credible collusive strategy because retailers "punish" deviation by believing in competitive conduct after a deviation. Adaptive beliefs can also satisfy the conditions for stability and opportunism-proofness.

Because manufacturer cartels are ubiquitous, our results can help competition authorities to screen markets. Our model shows that, whenever supply contracts are not public or easily renegotiable, the ability to form and sustain collusion critically depends on retailers' beliefs about the supply conditions of other retailers. This may make it easier to sustain collusion in markets in which the retailers are used to manufacturer collusion. Our findings suggest that it should be in the interest of colluding manufacturers to manage and influence their retailers' beliefs about the conditions in the wholesale market. One conjecture is thus that the opportunism problem may be one of the causes behind the widespread use of resale price maintenance and hub-and-spoke arrangements when manufacturers collude. A more direct control of retail prices by manufacturers in form of resale-price maintenance may circumvent the problem of skeptical retailer beliefs. Similarly, coordinated downsizing of packages by manufacturers, as observed in the chocolate case in Germany, may be used to reduce strategic uncertainty for retailers.[27] Consequently, coordinated behavior and communication of manufacturers vis à vis their retailers may deserve more antitrust scrutiny because such coordination can be essential for making manufacturer cartels work.

Competition authorities should also be aware that announcements by industry associations to raise prices can help member firms to overcome the opportunism problem in the context of collusion. When such an announcement comes from an authoritative body, downstream firms can expect that price increases will affect the whole industry and that they will not lose out vis à vis their competitors when accepting higher input prices.

Another aspect that should be of interest for competition authorities is that (sudden) changes of (industry-wide) input prices – for example, due to inflation or an uncertain economic environment – can be used by manufacturers to increase their prices collusively to a larger extent than such cost changes would imply. These changes may lead to a

[27] See Section 3.1.

situation in which downstream firms' expectations are less skeptical with regard to price increases, such that the transition to supra-competitive prices and profits is facilitated.[28]

[28]Recent developments in the retail food sector point in this direction (see, for example, Der Spiegel, "Is retail ripping off consumers?", 03/23/2023; last access 06/25/2023).

A Appendix: Omitted Proofs

Proof of Lemma 1. Plugging $p_{-i}(w_{-it}, p^e_{it}) = p^e_{-it}$ in the manufacturer problem 3.2. Simplifying yields

$$\Pi_{it}\left(w_{it}, p^e_{-it}\right) = p_i\left(w_{it}, p^e_{-it}\right) \cdot D_i\left(p_{it}\left(w_{it}, p^e_{-it}\right), p^e_{-it}\right).$$

Maximizing with respect to w_{it} gives the first-order condition

$$\begin{aligned} & \frac{\partial p_i\left(w_{it}, p^e_{-it}\right)}{\partial w_{it}} D_i\left(p_i\left(w_{it}, p^e_{-it}\right), p^e_{-it}\right) \qquad (3.18) \\ + & \frac{\partial D_i\left(p_i\left(w_{it}, p^e_{-it}\right), p^e_{-it}\right)}{\partial p_i} \frac{\partial p_i\left(w_{it}, p^e_{-it}\right)}{\partial w_{it}} p_i\left(w_{it}, p^e_{-it}\right) = 0 \\ \iff & \frac{\partial p_i\left(w_{it}, p^e_{-it}\right)}{\partial w_{it}} \frac{\partial D_i\left(p_i\left(w_{it}, p^e_{-it}\right), p^e_{-it}\right)}{\partial p_i} w_{it} = 0. \qquad (3.19) \end{aligned}$$

Because the first term on the left-hand side in the previous line is assumed to be strictly positive and the second term is assumed to be strictly negative (Assumption 5), the only solution to the first-order condition is $w_{it} = 0$. $\square$

Proof of Proposition 1. Step (i): Let us construct this equilibrium by showing that it is uniquely optimal for each manufacturer to set $(w_i = 0, F_{it} = \pi^*_i(w_i = 0, p^e_{-it}))$ in each period independent of the strategy of the other manufacturer.

The participation constraint of the retailer is binding in equilibrium because otherwise the manufacturer could increase profits by raising the fixed fee without affecting p^e_{-it} by the assumption of passive beliefs. Hence, $F_{it} = \pi^*_i(w_i = 0, p^e_{-it}))$ holds.

On the equilibrium path, $p^e_{-it} = p^*_{-it}$, that is, retailers' beliefs are correct and thus identical to manufacturer's conjectures such that the in-period manufacturer profits can be simplified as in the stage game to:

$$\Pi_{it}\left(w_{it}, p^e_{-it}\right) = p_i\left(w_{it}, p^e_{-it}\right) \cdot D_i\left(p_{it}\left(w_{it}, p^e_{-it}\right), p^e_{-it}\right).$$

Note that this in-period profit is insulated from the actual actions of the other manufacturer, and w_{it} only affects the manufacturer profits through the price setting of the retailer. Because the manufacturers' profit in each period, on any equilibrium path, only depends on w_{it} and the belief p^e_{-it}, the discounted equilibrium profits of a manufacturer do not depend on the strategy of the other manufacturer either. Fixing p^e_{-it}, equation (3.19) implies that $w_i = 0$ is optimal (independent of w_{-it} in each period). Hence, there is an equilibrium path with each manufacturer setting $w_i = 0$ in every subgame and a matching time-constant belief by retailers, where the time-constant retail price follows because the retail price equilibrium is unique for $w_i = w_{-i} = 0$.

Step (ii) Next, we exclude other equilibrium paths in which $w_{it} \neq 0$ by contradiction: Suppose that there is an equilibrium path with $w_{it} \neq 0$ in some periods. It follows from equation 3.19 that each manufacturer can increase current period profits through setting $w_{it} = 0$ and $F_{it} = \pi(0, p^e_{-it})$, resulting in $\Pi_{it}\left(0, p^e_{-it}\right)$. In particular, in every period, manufacturer i can ensure at least a profit of $\Pi(0, p^e) = \max_p \pi(0, p^e) = \max_p pD(p, p^e)$, which is independent of the actual wholesale price of the other manufacturer. This profit is strictly larger than the profit in any candidate equilibrium with $w' \neq 0$. Formally, this is $\Pi(w', p^e) < \Pi(0, p^e) \Leftrightarrow w'D(p^e) + (p^e - w')\,D(p^e) = p^e D(p^e) < \max_p pD(p, p^e)$, where the strict inequality follows from the uniqueness of the maximizer and the strict monotonicity of p_i in w_i .

Since each manufacturer can ensure this deviation profit of $\Pi(0, p^e)$ in any period independently of the other manufacturers actions, a profitable deviation from any candidate equilibrium with $w' \neq 0$ always exists and this deviation is immune to punishment by the other manufacturer.

As we employ beliefs about p^e_{it} instead of beliefs about w^e_{it}, we establish that this is equivalent in our setting if retailers have common knowledge about the belief system and expect sequentially rational choices of their rivals. More precisely, for any history $\mathcal{H}_t$ and belief w^e_{it}, there exists a unique implied belief. Expecting a sequentially rational choice implies that the belief of retailer A, p^e_{Bt}, is the result of $p^e_{Bt} = \arg\max_{p_{Bt}} \pi_B(w^e_{Bt}, w^e_{At} = w^{2e}_{At})$, where w^{2e}_{At} is the higher-order belief of retailer A about retailer B's belief. The identity $w^e_{At} = w^{2e}_{At}$ follows from common knowledge about the belief system and Condition 1. Due to condition 1 beliefs are within-period passive, such that the higher-order belief w^{2e}_{At} is independent of retailer A's private information w_{At}. Common knowledge about the belief system implies then that retailers hold correct conjectures about higher order beliefs. Therefore, given w^e_{-it}, retailer i correctly predicts w^{2e}_{it} and thus the unique $p^e_{-it} = \arg\max_{p_{-it}} \pi_{-i}(w^e_{-it}, w^{2e}_{it})$. This follows from Assumption 5, which implies that there is a unique solution to the maximization problem. □

Proof of Proposition 2. We established in the main text that manufacturers are able to *sustain* collusion if equation (3.9) is fulfilled. To show that the incentive constraint is the same if manufacturers and retailers are pairwise integrated, it is sufficient to prove that the profits are the same: $\Pi^j_V = \Pi^j_I$ for $j = C, D, P$ (V: vertical, I: integrated). Because we define Π^C_V to be half of the integrated industry-maximizing profit (3.4), the profit from collusion in the integrated case Π^C_I is the same.

The deviation profit Π^D_V is given by the first line of equation (3.5) and simplifies further in the second line because the manufacturers' true actions and retailers' beliefs are aligned. The second line, however, is equal to the maximization problem of an integrated manufacturer that maximizes the profit with respect to the retail price. Lastly, the punishment profits are aligned as well, following the same argument as before. The punishment profits are given by equation (3.7). Manufacturers set $w^P = 0$, which results in the same retail prices and profits as in a vertically integrated industry.

To see that punishment is renegotiation-proof, note that with grim-trigger beliefs, any deviation leads to the belief p^P forever such that the current actions of the manufacturers have no effect on the belief of retailers. To establish this, we check for a joint action by manufacturers that would yield a Pareto improvement for manufactures. Note that it is the best response is, in each period, for each manufacturer, to set $w = 0$ individually. Next, we consider the profit maximization for manufacturers when they would optimize jointly during a punishment period:

$$\frac{1}{2}\max_{w_{it},w_{-i,t}} w_{it} \cdot D_i\left(p_i(w_{it},p^P),p_{-i}(w_{-it},p^P)\right) + \left[p_i(w_{it},p^P) - w_{it}\right] \cdot D_i\left(p_i(w_{it},p^P),p^P\right)$$
$$+w_{-it} \cdot D_{-i}\left(p_{-i}(w_{-it},p^P),p_i(w_{it},p^P)\right) + \left[p_{-i}(w_{-it},p^P) - w_{-it}\right] \cdot D_{-i}\left(p_{-i}(w_{-it},p^P),p^P\right).$$

We again assume that this profit is quasi-concave such that we can use first-order conditions. Using the retailers' first-order conditions and applying symmetry, using that manufacturer profits are well-behaved by assumption such that the optimum is symmetric, the first-order conditions can be rewritten as

$$\left[\underbrace{D_i\left(p_i(w,p^P),p_{-i}(w,p^P)\right) - D_i\left(p_i(w,p^P),p^P\right)}_{\geq 0}\right] + w\underbrace{\frac{\partial p_i}{\partial w_{it}}}_{>0}\left[\underbrace{\frac{\partial D_i}{\partial p_i} + \frac{\partial D_{-i}}{\partial p_i}}_{<0}\right] = 0\ \forall i. \tag{3.20}$$

This holds for $w = 0$. Hence, playing $w = 0$ is the best manufacturers can do such that there is no continuation game that can be reached by manufacturers and that yields larger profits.

To see that collusion is *not formable*, consider any history with $w_i \neq w^C$ in $t = 0$. With this history, grim-trigger beliefs imply that the belief is p^P forever, which violates the second condition of the definition of formability that a collusive PBE can be obtained in some future period for any history. □

Proof of Proposition 3. Consider the conditions that are needed to sustain collusion. After any deviation by a manufacturer, the beliefs revert to $p^e_{-i} = p^P$ for κ periods, which results in profits of $\Pi^P = p^P \cdot D_i(p^P, p^P)$. After κ periods, however, the retailers believe in collusion at a price of p^C again. Collusion is sustainable when no manufacturer wants to deviate from the trigger strategy given by the incentive condition:

$$\frac{\Pi^C}{1-\delta} \geq \Pi^D + \delta\left(\frac{1-\delta^\kappa}{1-\delta}\Pi^P + \frac{\delta^\kappa}{1-\delta}\Pi^C\right). \tag{3.21}$$

This condition is more difficult to fulfill than the incentive condition (3.9) for grim-trigger strategies and beliefs. The punishment in condition (3.21) is less harsh and ends after κ periods, such that the expression on the right-hand side is larger than in the condition with grim-trigger strategies. Recall that manufacturers make the same profits as pairwise integrated manufacturer-retailer pairs (Proposition 2). Note that the individual profits, Π^C, Π^D, and Π^P, are still identical to an integrated firm's profits. This implies that the stability condition (3.21) is the same for vertically separated and vertically integrated manufacturer-retailer pairs whenever they play trigger strategies of length κ.

To check whether the equilibrium is opportunism-proof, we must consider a revised version of condition (3.11) that applies to trigger beliefs. Because trigger beliefs are forgiving after κ periods, a joint deviation of both manufacturers is not "punished by the beliefs" forever.[29] Hence, the conditions can be written as

$$\frac{\Pi^C}{1-\delta} \geq \Pi^{JD} + \delta\left(\frac{1-\delta^\kappa}{1-\delta}\Pi^P + \frac{\delta^\kappa}{1-\delta}\Pi^C\right). \tag{3.22}$$

Again, the opportunism-proofness condition (3.22) resembles the stability condition (3.21), except that $\Pi^{JD} > \Pi^D$, which makes the condition harder to meet. If condition (3.22) holds, collusion is robust against joint deviations by the manufacturers, that is, opportunistic behavior.

The punishment is renegotiation-proof because, in a punishment phase, both retailers believe $w^P = 0$ to which the best response is $w = 0$ as well. Beliefs are constant in the wholesale price played in the punishment phase. Thus, the argument in Subsection 3.4.2 applies, such that focusing on short-term best responses is valid.

To see that collusion is not formable with trigger beliefs, consider an argument by contradiction. Formability requires that the PBE under consideration can arise after an arbitrary history of the game. Let us consider a specific history. Suppose that, for the first κ periods, the manufacturers do not collude. With this history, all the following periods starting with period $\kappa + 1$ are labeled as non-collusive. It follows that it is impossible for retailers to hold collusive beliefs at any future point of the game. Recall that periods are only labeled collusive after the first period if in period t, the previous period $t-1$ was collusive or period $t-\kappa$ was collusive. □

Proof of Proposition 4. Similarly to the incentive constraints above the proposition for grim-trigger strategies, we can also describe the incentive constraints when manufacturers play more general trigger strategies. The difference between these strategies is the length of the punishment phase κ. Because neither the formation condition (3.13) nor the

[29] By contrast, the grim-trigger beliefs switch forever to the competitive price level in response to a deviation. This effectively punishes the manufacturers that face pessimistic retailers from then on.

opportunism-proofness condition (3.15) rely on punishment between the colluding firms, the conditions remain the same.

In the remainder of the proof, we characterize the stability condition and the stability condition of the formation phase for trigger strategies. If firms play trigger strategies with punishment length κ, the incentive constraint for stability is

$$\frac{\Pi^C}{1-\delta} \geq \Pi^D + \delta\left(\frac{1-\delta^\kappa}{1-\delta}\right)\Pi^P + \delta^{\kappa+1}\left(\frac{1-\delta^T}{1-\delta}\right)\Pi^F + \delta^{\kappa+1}\frac{\delta^T}{1-\delta}\Pi^C.$$

As in all the above cases, the punishment must be $w^P = 0$ because there is no unilateral action in a punishment phase that has any effect on future beliefs. Because the future beliefs and actions of the rival manufacturer are fixed in any punishment period, each manufacturer must play the one-shot best response of $w^P = 0$.

The deviation of one manufacturer triggers a punishment of length $\kappa \in [0, \infty)$. After the punishment phase, collusion is resumed, but the adaptive retailer beliefs require a formation phase of length T. A deviation in the formation phase is not profitable if

$$\frac{1-\delta^T}{1-\delta}\Pi^F + \frac{\delta^T}{1-\delta}\Pi^C \geq \Pi^{F,D} + \delta\left(\frac{1-\delta^\kappa}{1-\delta}\right)\Pi^P + \delta^{\kappa+1}\left(\frac{1-\delta^T}{1-\delta}\right)\Pi^F + \delta^{\kappa+1}\frac{\delta^T}{1-\delta}\Pi^C, \quad (3.23)$$

which simplifies to

$$\delta^T\left(\Pi^C - \Pi^F\right) \geq \left(\frac{(1-\delta)}{1-\delta^{\kappa+1}}\right)\Pi^{F,D} - \Pi^F + \delta\left(\frac{1-\delta^\kappa}{1-\delta^{\kappa+1}}\right)\Pi^P. \quad (3.24)$$

Next, we show that condition (3.13) is always stricter than condition (3.23). We also demonstrate that this is the relevant condition for formation. To see this, we consider the polar cases that correspond to a strategy "always collude" for $\kappa = 0$ and the grim-trigger strategy as $\kappa \to \infty$. We already know from the analysis before the proposition that the postulated relation of the conditions holds for $\kappa \to \infty$.

For $\kappa = 0$, the incentive constraint for formation can be written as

$$\delta^T\left(\Pi^C - \Pi^F\right) \geq \Pi^P - \Pi^F.$$

A deviation from the collusive price in the formation phase results in a profit of $\Pi^{F,D}$. However, for $\kappa = 0$, no punishment is triggered. The stability condition for formation becomes

$$\frac{1-\delta^T}{1-\delta}\Pi^F + \frac{\delta^T}{1-\delta}\Pi^C \geq \Pi^{F,D} + \delta\left(\frac{1-\delta^T}{1-\delta}\frac{\Pi^F}{1-\delta} + \frac{\delta^T}{1-\delta}\Pi^C\right),$$

which simplifies to

$$\delta^T\left(\Pi^C - \Pi^F\right) \geq \Pi^{F,D} - \Pi^F. \quad (3.25)$$

Observe that a deviation in the formation phase is not profitable under a stricter condition than the incentive condition for formation if $\Pi^{F,D} \geq \Pi^{P}$. Additionally, a necessary condition is $\Pi^{C} \geq \Pi^{F,D}$ because the discount factor δ is in the range $(0,1)$. Because $\delta < 1$, the left-hand side of condition (3.25) decreases in T, such that formability becomes harder to satisfy the larger T.

Thus, for both cases of $\kappa = 0$ and $\kappa \to \infty$, we demonstrated that sticking to the formation phase requires higher values of δ (the condition is more strict) if $\Pi^{F,D} \geq \Pi^{P}$. Let us now analyze intermediary values of κ.

Inequality (3.24) differs from the formation incentives. Recall that the formation condition (3.13) is

$$\delta^{T}\left(\Pi^{C} - \Pi^{F}\right) \geq \Pi^{P} - \Pi^{F}$$

and thus independent of κ. For a given value of δ, the weight of the term $\Pi^{F,D}$ is strictly monotonically decreasing in κ, whereas the weight of Π^{P} is strictly monotonically increasing in κ. Hence, the condition $\Pi^{F,D} \geq \Pi^{P}$ is sufficient for any $\kappa \in [0,\infty)$ to guarantee that condition (3.24) is tighter than condition (3.13). This implies that the stability condition for formation is the relevant condition for general trigger strategies.

Next, we show that

$$\Pi^{F,D} \geq \Pi^{P}$$

holds because the profit of i increases in the wholesale price of the competitor for given beliefs due to the fact that this increases the price of the competing retailer and increases demand for i. The profit under competition is

$$\Pi^{P}\left(w_i = 0, p^{e}_{-i} = p^{P}\right) = \underbrace{p_i\left(0, p^{P}\right)}_{p_i^{P}} \cdot D_i\left(p^{P}, p^{P}\right).$$

A firm that deviates from formation plays $w_i = 0$ because this maximizes the unilateral spot profit. The deviation profit $\Pi^{F,D}$ is similar to the punishment profit and only consists of the revenue $p_i(\cdot) \cdot D_i(\cdot)$. For $\Pi^{F,D}$, however, both the manufacturer and the retailer expect a higher wholesale price of w^{C} of the other manufacturer, yielding a higher price at the competing retailer. Because profits increase in the rival's price, $\Pi^{F,D} > \Pi^{P}$ holds.

Finally, observe that punishment is not renegotiation-proof whenever collusion is formable. To see this consider that the formability condition implies that is better to jointly deviate to a new formation of collusion in any period in which punishment profits are expected. This also implies that only in the case $\kappa = 0$, formability and renegotiation-proof punishment are both met because punishment is not part of the strategy. □

Proof of Corollary 3. The equilibrium path is always Pareto-efficient by assumption, such that any weakly renegotiation-proof equilibrium is also strongly renegotiation-proof.

To find the critical value of T, such that collusion is not formable, but opportunism-proof for $\kappa > 0$, we must compare the relevant conditions. If collusion is not formable, punishment will be renegotiation-proof. The equilibrium is thus weakly renegotiation-proof for certain values of T. As shown in the proof of Proposition 4, the relevant condition for formation is the condition whereby a deviation from formation is not profitable:

$$\delta^T \left(\Pi^C - \Pi^F\right) \geq \left(\frac{1-\delta}{1-\delta^{\kappa+1}}\right) \Pi^{F,D} - \Pi^F + \delta \left(\frac{1-\delta^\kappa}{1-\delta^{\kappa+1}}\right) \Pi^P > 0.$$

Because $\delta < 1$, this condition is violated for sufficiently large T, such that δ^T and thus the left-hand side of the formation condition becomes arbitrarily small.

The condition for opportunism-proofness 3.15 can be written as

$$\delta^T \left(\Pi^C - \Pi^F\right) \leq \frac{1}{\delta} \left[\Pi^C - (1-\delta)\, \Pi^{JD} - \delta \Pi^F\right].$$

Because for $\delta < 1$, the left-hand side of the opportunism-proofness condition also becomes arbitrarily small as T increases. Hence, there exists a T, such that the opportunism condition holds, whenever the right-hand side is non-negative (which holds for sufficiently large δ), while the formation condition is violated. $\square$

Proof of Proposition 5. We will show all results for general trigger strategies as defined in Subsection 3.4.2, which includes grim-trigger strategies as a subcase for $\kappa \to \infty$. Let us first consider the stability condition. As argued in the text before the proposition, manufacturers collude on the industry-maximizing level and earn profits of $\Pi^C = \Pi^M/2$. If the manufacturers make a one-sided deviation from the collusive agreement, they must consider that changing their wholesale price also influences the belief of their own retailer. A deviating manufacturer maximizes the following problem:

$$\begin{aligned} w^D = \arg\max\nolimits_{w_i} w_{it} \cdot D_i \left(p_i(w_{it}, p(w_{it})), p_{-i}(w^C, p(w^C))\right) \\ + \left(p_i(w_{it}, p(w_{it})) - w_{it}\right) \cdot D_i \left(p_i(w_{it}, p(w_{it})), p(w_{it})\right)\right). \end{aligned}$$

This results in a profit Π^D for the deviating manufacturer. Note that $w^D > 0$. Punishment is assumed to be carried out on the competitive wholesale price level. The competitive benchmark corresponds to the case analyzed in Pagnozzi and Piccolo (2011). Under symmetric beliefs, manufactures solve

$$\begin{aligned} \max_{w_i} w_{it} \cdot D_i \left(p_i \left(w_{it}, p(w_{it})\right), p_{-i} \left(w_{-it}, p(w_{-it})\right)\right) \\ + \left(p_i(w_{it}, p(w_{it})) - w_{it}\right) \cdot D_i \left(p_i \left(w_{it}, p(w_{it})\right), p(w_{it})\right). \end{aligned}$$

This results in Π^P. By using the envelope theorem, the first-order condition simplifies to

$$\underbrace{w_i \underbrace{\frac{\partial D_i(\cdot)}{\partial p_i}}_{<0}}_{<0} + \underbrace{(p(w_i, p(w_i)) - w_i) \cdot \underbrace{\frac{\partial D_i(\cdot)}{\partial p_{-i}}}_{>0}}_{\text{belief effect:}>0} = 0. \tag{3.26}$$

Applying symmetry to equation (3.26) defines the equilibrium wholesale prices w^P. Note that the second term of (3.26) is positive at $w_i = 0$, which implies that $w^P > 0$. Under competition, with symmetric beliefs, prices are above the price level under competition with passive beliefs. The stability condition is given by

$$\frac{\Pi^C}{1-\delta} \geq \Pi^D + \delta\left(\frac{1-\delta^\kappa}{1-\delta}\Pi^P + \frac{\delta^\kappa}{1-\delta}\Pi^C\right). \tag{3.27}$$

To see why collusion is formable with symmetric beliefs, recall the definition of formability. Symmetric beliefs allow for forming collusion if there exists a strategy profile, such that for this belief, best responses are played in period $t \geq s$, and there exists a weak PBE that results in payoff V^C. Given any history before period s – the period in which collusion is about to be formed –, manufacturers play mutually best responses in the following periods when setting wholesale prices that result in the collusive price defined by equation (3.3). The joint profit maximization of both manufacturers and setting w^C is a weak PBE if the stability condition for collusion – equation (3.27) – is fulfilled. In this equilibrium, both manufacturers earn a payoff of $V^C = \Pi^C/(1-\delta)$. Hence, collusion is formable according to our definition.

As in Subsection 3.4.2, it should be considered whether there exist incentives to deviate from formation. Due to symmetric beliefs, collusion is in place directly after it is formed, such that the condition for deviating from formation is identical to deviating from collusion in equation (3.27).

To check whether collusion is opportunism-proof, we consider the joint maximization problem of the manufacturers. As shown in the text, this leads to the Pareto-efficient wholesale price level. That is, manufacturers always prefer to set w^C when jointly maximizing. Thus, there is no scope for opportunistic behavior because the beliefs directly react to any change in wholesale prices. Following the same argument, punishment is not renegotiation-proof because manufacturers would prefer to renegotiate and jointly revert to setting w^C.

The competitive equilibrium is not renegotiation proof whenever a collusive equilibrium with higher manufacturer profits exists. Hence, if condition 3.27 holds, there exists a Pareto-dominant equilibrium for manufacturers that can be reached by a joint deviation of the manufacturer. With symmetric beliefs retailers are immediately willing to accept the collusive contract expecting wholesale prices to increase in the whole industry. □

Proof of Proposition 6. If $0 < \frac{\partial w^D(w^C)}{\partial w^C}|_{w^C=w^D} < 1$ holds, then increasing w^C implies $w^C > w^D$. Then, $w^C > w^N$ (by construction) implies that $w^D > w^N$. Using $w^C > w^D$, we show that

$w^D > w^P$ holds: Note that $w^D > w^P \Leftrightarrow p^e > p^*(w^P, p^e) \Leftrightarrow w^P < w^C$. The last inequality follows again from the argument that starting at $w^C = w^N$, increasing w^C implies by $\frac{\partial w^P(w^C)}{\partial w^C} < 1$ that $w^C > w^N$ implies $w^C > w^P$. Hence, $w^C > w^P$ implies that $w^D > w^P$must hold.

Rewrite condition (3.17) as

$$2\Pi^C - \Pi^P - \Pi^D > 0.$$

This can be rewritten to

$$\int_{w^p}^{w^c} wD(p^*(w,p^e), p^*(w,p^e))dw + \int_{w^D}^{w^c} wD(p^*(w,p^e), p^e)dw$$

$$= \int_{w^p}^{w^D} wD(p^*(w,p^e), p^*(w,p^e))dw + \int_{w^D}^{w^c} w\left[D(p^*(w,p^e), p^e) + D(p^*(w,p^e), p^e)\right]dw. \tag{3.28}$$

Note $2\Pi^C - \Pi^P - \Pi^D = 0$ holds at $w^C = w^N$. Thus if, at this locus, the derivative of (3.28) with respect to w^C is positive, then the stability condition holds for some $w^C > w^N$and some $\delta < 1$. Thus, there is a collusive equilibrium if, at $w^C = w^N$,

$$\partial\left(\int_{w^p}^{w^D} wD(p^*(w,p^e), p^*(w,p^e))dw + \int_{w^D}^{w^c} w\left[D(p^*(w,p^e), p^e) + D(p^*(w,p^e), p^e)\right]dw\right)/\partial w^C > 0.$$

Note that, at $w^C = w^N$, it holds that $w^D = w^P = w^N = w^C$. By Leibniz integral rule we have

$$\begin{aligned}
&\partial\left(\int_{w^p}^{w^D} wD(p^*(w,p^e), p^*(w,p^e))dw + \int_{w^D}^{w^c} w\left[D(p^*(w,p^e), p^e) + D(p^*(w,p^e), p^e)\right]dw\right)/\partial w^C \\
&= w^D D(p^*(w^D,p^e), p^*(w^D,p^e))w^{D\prime} - w^P D(p^*(w^P,p^e), p^*(w^P,p^e))w^{P\prime} \\
&+ \int_{w^p}^{w^D} \left[wD(p^*(w,p^e), p^*(w,p^e))\right]' dw + w^C\left[D(p^*(w^C,p^e), p^e) + D(p^*(w^C,p^e), p^e)\right] \\
&-w^D\left[D(p^*(w^D,p^e), p^e) + D(p^*(w^D,p^e), p^e)\right]w^{D\prime} + \int_{w^D}^{w^c}\left[w\left[D(p^*(w,p^e), p^e) + D(p^*(w,p^e), p^e)\right]\right]' dw
\end{aligned}$$

Evaluation the derivative at $w^D = w^P = w^N = w^C = w^*$ yields

$$w^D D(p^*(w^D, p^e), p^*(w^D, p^e))w^{D\prime} - w^P D(p^*(w^P, p^e), p^*(w^P, p^e))w^{P\prime}$$

$$+ \underbrace{\int_{w^p}^{w^D} [wD(p^*(w, p^e), p^*(w, p^e))]' \, dw}_{=0} + w^C \left[D(p^*(w^C, p^e), p^e) + D(p^*(w^C, p^e), p^e)\right]$$

$$-w^D \left[D(p^*(w^D, p^e), p^e) + D(p^*(w^D, p^e), p^e)\right] w^{D\prime} + \underbrace{\int_{w^D}^{w^c} [w\,[D(p^*(w, p^e), p^e) + D(p^*(w, p^e), p^e)]]' \, dw}_{=0} .$$

Define $D^* = D(p^*(w^*, p^e), p^*(w^*, p^e)) = D(p^*(w^*, p^e), p^e)$ with $p^*(w^*, p^e) = p^e$. The above expression can then be rewritten as

$$w^* D(p^*(w^*, p^e), p^*(w^*, p^e))(w^{D\prime} - w^{P\prime}) +$$

$$w^* \left[D(p^*(w^*, p^e), p^e) + D(p^*(w^*, p^e), p^e)\right](1 - w^{D\prime})$$

$$= w^* D^* (2 - w^{D\prime} - w^{P\prime}).$$

As $w^* D^* > 0$, this is larger than zero if and only if

$$w^{D\prime} + w^{P\prime} < 2,$$

where $w^{D\prime} = \frac{\partial w^D(w^C)}{\partial w^C}|_{w^D = w^C}$. This holds if $w^{D\prime} < 1$ and $w^{P\prime} < 1$, which holds by Assumption 6.

Formability requires that the wPBE can be reached from an arbitrary history through a sequentially rational transition. For the collusive wPBE, the system of beliefs in this case is simply that, starting in period t_0, the beliefs are constant in w^C for both retailers. Thus, if the stability condition for collusion holds, collusion is a wPBE of the continuation game starting in period t_0 (there is no transition required). $\square$

Bibliography

ACM, "Market Study into Mobile App Stores," Report, The Netherlands Authority for Consumers and Markets 2019.

Almar, Frederik, Benjamin Friedrich, Ana Reynoso, Bastian Schulz, and Rune Vejlin, "Marital Sorting and Inequality: How Educational Categorization Matters," CESifo Working Paper 2023.

Anderson, Simon P. and Stephen Coate, "Market Provision of Broadcasting: A Welfare Analysis," *The Review of Economic Studies*, 2005, *72* (4), 947–972.

Andersson, Ola and Erik Wengström, "Do Antitrust Laws Facilitate Collusion? Experimental Evidence on Costly Communication in Duopolies," *Scandinavian Journal of Economics*, 2007, *109* (2), 321–339.

Antler, Yair and Benjamin Bachi, "Searching Forever After," *American Economic Journal: Microeconomics*, 2022, *14* (3), 558–590.

_ **, Daniel Bird, and Daniel Fershtman**, "Search, Dating, and Segregation in Marriage," Discussion Paper 7-2023, The Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 2023.

_ **,** _ **, and** _ , "Search, Matching, and Online Platforms," Working Paper, The Eitan Berglas School of Economi, Tel Aviv University, Tel Aviv 2024.

Aoyagi, Masaki, Guillaume Fréchette, and Sevgi Yuksel, "Beliefs in Repeated Games: An Experiment," *American Economic Review*, 2024, *114* (12), 3944–3975.

Aridor, Guy, "Market Definition in the Attention Economy: An Experimental Approach," *American Economic Journal: Microeconomics*, Forthcoming.

Armstrong, Mark, "Competition in Two-Sided Markets," *The RAND Journal of Economics*, 2006, *37* (3), 668–691.

_ **and John Vickers**, "Consumer Protection and Contingent Charges," *Journal of Economic Literature*, 2012, *50* (2), 477–493.

Asker, John and Heski Bar-Isaac, "Vertical Information Restraints: Pro-and Anti-competitive Impacts of Minimum-Advertised-Price Restrictions," *The Journal of Law and Economics*, 2020, *63* (1), 111–148.

Athey, Susan, Kyle Bagwell, and Chris Sanchirico, "Collusion and Price Rigidity," *Review of Economic Studies*, 2004, *71* (2), 317–349.

Bagwell, Kyle, "The Economic Analysis of Advertising," *Handbook of Industrial Organization*, 2007, *3*, 1701–1844.

_ **and Garey Ramey**, "Advertising and Limit Pricing," *The RAND Journal of Economics*, 1988, pp. 59–71.

Banks, Jeffrey S. and Joel Sobel, "Equilibrium Selection in Signaling Games," *Econometrica: Journal of the Econometric Society*, 1987, pp. 647–661.

Becker, Gary S., "A Theory of Marriage: Part I," *Journal of Political Economy*, 1973, *81* (4), 813–846.

_ , "A Theory of Social Interactions," *Journal of Political Economy*, 1974, *82* (6), 1063–1093.

Bernheim, B. Douglas, "Rationalizable Strategic Behavior," *Econometrica: Journal of the Econometric Society*, 1984, pp. 1007–1028.

_ **and Debraj Ray**, "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, 1989, *1* (4), 295–326.

_ **, Bezalel Peleg, and Michael D. Whinston**, "Coalition-Proof Nash Equilibria I. Concepts," *Journal of Economic Theory*, 1987, *42* (1), 1–12.

Bertsimas, Dimitris and John N. Tsitsiklis, *Introduction to Linear Optimization*, Vol. 6, Athena Scientific Belmont, MA, 1997.

Bloch, Francis and Harl Ryder, "Two-Sided Search, Marriages, and Matchmakers," *International Economic Review*, 2000, *41* (1), 93–116.

Blume, Andreas, "Intraplay Communication in Repeated Games," *Games and Economic Behavior*, 1994, *6* (2), 181–211.

_ **and Paul Heidhues**, "Modeling Tacit Collusion in Auctions," *Journal of Institutional and Theoretical Economics*, 2008, *164* (1), 163–184.

Bos, Iwan and Joseph E. Harrington, "Endogenous Cartel Formation with Heterogenous Firms," *The RAND Journal of Economics*, 2010, *41* (1), 92–117.

_ , **Marco A. Marini, and Riccardo D. Saulle**, "Cartel Formation with Quality Differentiation," *Mathematical Social Sciences*, 2020, *106* (C), 36–50.

_ , **Ronald Peeters, and Erik Pot**, "Do Antitrust Agencies Facilitate Meetings in Smoke-Filled Rooms?," *Applied Economics Letters*, 2013, *20* (6), 611–614.

_ , **Wilko Letterie, and Dries Vermeulen**, "Antitrust as Facilitating Factor for Collusion," *B.E. Journal of Economic Analysis & Policy*, 2015, *15* (2), 797–814.

Bruch, Elizabeth E. and Mark E.J. Newman, "Aspirational Pursuit of Mates in Online Dating Markets," *Science Advances*, 2018, *4* (8), eaap9815.

_ **and** _ , "Structure of Online Dating Markets in US Cities," *Sociological Science*, 2019, *6*, 219–234.

Buehler, Stefan and Dennis L. Gärtner, "Making Sense of Nonbinding Retail-Price Recommendations," *American Economic Review*, 2013, *103* (1), 335–59.

Burdett, Kenneth and Melvyn G. Coles, "Long-Term Partnership Formation: Marriage and Employment," *The Economic Journal*, 1999, *109* (456), 307–334.

Burks, Stephen V., Jeffrey P. Carpenter, Lorenz Goette, and Aldo Rustichini, "Overconfidence and Social Signalling," *Review of Economic Studies*, 2013, *80* (3), 949–983.

Caillaud, Bernard and Bruno Jullien, "Chicken & Egg: Competition among Intermediation Service Providers," *The RAND Journal of Economics*, 2003, *34* (2), 309–328.

Chade, Hector, Jan Eeckhout, and Lones Smith, "Sorting Through Search and Matching Models in Economics," *Journal of Economic Literature*, 2017, *55* (2), 493–544.

Charness, Gary, Aldo Rustichini, and Jeroen Van de Ven, "Self-Confidence and Strategic Behavior," *Experimental Economics*, 2018, *21*, 72–98.

Chen, Daniel, "The Market for Attention," Working Paper, Available at SSRN 2022.

Chen, Ying, "Perturbed Communication Games with Honest Senders and Naive Receivers," *Journal of Economic Theory*, 2011, *146* (2), 401–424.

Cho, In-Koo and David M. Kreps, "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 1987, *102* (2), 179–221.

_ **and Joel Sobel**, "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory*, 1990, *50* (2), 381–413.

Cong, Lin William, Xi Li, Ke Tang, and Yang Yang, "Crypto Wash Trading," *Management Science*, 2023, *69* (11), 6427–6454.

Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer, *Competition Policy for the Digital Era*, Publications Office, 2019. Published by the European Commission: Directorate-General for Competition.

Damiano, Ettore and Hao Li, "Price Discrimination and Efficient Matching," *Economic Theory*, 2007, *30* (2), 243–263.

_ **and Li Hao**, "Competing Matchmaking," *Journal of the European Economic Association*, 2008, *6* (4), 789–818.

Dantzig, George B., *Linear Programming and Extensions*, Princeton: Princeton University Press, 1963.

Dargnies, Marie-Pierre, Rustamdjan Hakimov, and Dorothea Kübler, "Self-Confidence and Unraveling in Matching Markets," *Management Science*, 2019, *65* (12), 5603–5618.

d'Aspremont, Claude, Alexis Jacquemin, Jean Gabszewicz, and John Weymark, "On the Stability of Collusive Price Leadership," *Canadian Journal of Economics*, 1983, *16* (1), 17–25.

DellaVigna, Stefano and Ulrike Malmendier, "Paying Not to Go to the Gym," *American Economic Review*, 2006, *96* (3), 694–719.

Deng, Yiting, Anja Lambrecht, and Yongdong Liu, "Spillover Effects and Freemium Strategy in the Mobile App Market," *Management Science*, 2023, *69* (9), 5018–5041.

Dertwinkel-Kalt, Markus, Vincent Eulenberg, and Christian Wey, "Defining What the Relevant Market is: A New Method for Consumer Research and Antitrust," *Available at SSRN*, 2024.

Do, Jihwan and Jeanine Miklós-Thal, "Opportunism in Vertical Contracting: A Dynamic Perspective," *available at SSRN*, 2021.

Dolado, Juan J., Marcel Jansen, and Juan F. Jimeno, "On-the-Job Search in a Matching Model with Heterogeneous Jobs and Workers," *The Economic Journal*, 2009, *119* (534), 200–228.

Dubra, Jean-Pierre Benoît Juan, "Does the Better-Than-Average Effect Show That People Are Overconfident?: Two Experiments," *Journal of the European Economic Association*, 2015, *13* (2), 293–329.

Echenique, Federico, Nicole Immorlica, and Vijay V. Vazirani, *Online and Matching-Based Market Design*, Cambridge University Press, 2023.

Eeckhout, Jan, "Bilateral search and Vertical Heterogeneity," *International Economic Review*, 1999, *40* (4), 869–887.

Egebark, Johan, Mathias Ekström, Erik Plug, and Mirjam Van Praag, "Brains or Beauty? Causal Evidence on the Returns to Education and Attractiveness in the Online Dating Market," *Journal of Public Economics*, 2021, *196*, 104372.

Eika, Lasse, Magne Mogstad, and Basit Zafar, "Educational Assortative Mating and Household Income Inequality," *Journal of Political Economy*, 2019, *127* (6), 2795–2835.

Eliaz, Kfir and Ran Spiegler, "On the Strategic Use of Attention Grabbers," *Theoretical Economics*, 2011, *6* (1), 127–155.

_ **and** _ , "A Simple Model of Search Engine Pricing," *The Economic Journal*, 2011, *121* (556), F329–F339.

_ **and** _ , "Search Design and Broad Matching," *American Economic Review*, 2016, *106* (3), 563–586.

Evans, David S. and Richard Schmalensee, "The New Economics of Multi-Sided Platforms: A Guide to the Vocabulary," *Available at SSRN 2793021*, 2016.

Farrell, Joseph and Eric Maskin, "Renegotiation in Repeated Games," *Games and Economic Behavior*, 1989, *1* (4), 327–360.

Fershtman, Daniel and Alessandro Pavan, "Matching Auctions," *The RAND Journal of Economics*, 2022, *53* (1), 32–62.

Fishman, Michael J. and Kathleen M. Hagerty, "Mandatory versus Voluntary Disclosure in Markets with Informed and Uninformed customers," *Journal of Law, Economics, and Organization*, 2003, *19* (1), 45–63.

Fletcher, Amelia, Gregory S. Crawford, Jacques Crémer, David Dinielli, Paul Heidhues, Michael Luca, Tobias Salz, Monika Schnitzer, Fiona M. Scott Morton, Katja Seim, and Michael Sinkinson, "Consumer Protection for Online Markets and Large Digital Platforms," Policy Discussion Paper No. 1, Yale Tobin Center for Economic Policy 2021.

Fudenberg, Drew and Jean Tirole, *Game Theory*, 1 ed., Vol. 1, MIT Press, 1991.

— **and** — , "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 1991, *53* (2), 236–260.

Gal-Or, Esther, "Market Segmentation on Dating Platforms," *International Journal of Industrial Organization*, 2020, *68*, 102558.

Gale, David and Lloyd S. Shapley, "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 1962, *69* (1), 9–15.

Gaudin, Germain, "Vertical Relations, Opportunism, and Welfare," *The RAND Journal of Economics*, 2019, *50* (2), 342–358.

Gilo, David and Yaron Yehezkel, "Vertical Collusion," *The RAND Journal of Economics*, 2020, *51* (1), 133–157.

Glazer, Jacob, Helios Herrera, and Motty Perry, "Fake Reviews," *The Economic Journal*, 2021, *131* (636), 1772–1787.

Goldlücke, Susanne and Sebastian Kranz, "Renegotiation-Proof Relational Contracts," *Games and Economic Behavior*, 2013, *80*, 157–178.

Gomes, Renato and Alessandro Pavan, "Many-to-Many Matching and Price Discrimination," *Theoretical Economics*, 2016, *11* (3), 1005–1052.

— **and** — , "Price Customization and Targeting in Matching Markets," *The RAND Journal of Economics*, 2024, *55* (2), 230–265.

González-Díaz, Julio and Miguel A. Meléndez-Jiménez, "On the Notion of Perfect Bayesian Equilibrium," *Top*, 2014, *22*, 128–143.

Green, Edward and Robert Porter, "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 1984, *52* (1), 87–100.

Greitemeyer, Tobias, "Unattractive People are Unaware of Their (Un) Attractiveness," *Scandinavian Journal of Psychology*, 2020, *61* (4), 471–483.

Grossman, Sanford J, "The Informational Role of Warranties and Private Disclosure about Product Quality," *The Journal of Law and Economics*, 1981, *24* (3), 461–483.

Gürtzgen, Nicole, Benjamin Lochner, Laura Pohlan, and Gerard J. van den Berg, "Does Online Search Improve the Match Quality of New Hires?," *Labour Economics*, 2021, *70*, 101981.

Hagiu, Andrei and Bruno Jullien, "Why Do Intermediaries Divert Search?," *The RAND Journal of Economics*, 2011, *42* (2), 337–362.

_ **and** _ , "Search Diversion and Platform Competition," *International Journal of Industrial Organization*, 2014, *33*, 48–60.

Halaburda, Hanna, Mikołaj Jan Piskorski, and Pınar Yıldırım, "Competing by Restricting Choice: The Case of Matching Platforms," *Management Science*, 2018, *64* (8), 3574–3594.

_ **, Mikołaj Jan Piskorski, and Pınar Yıldırım**, "Competing by Restricting Choice: The Case of Matching Platforms," *Management Science*, 2018, *64* (8), 3574–3594.

Hardy, Godfrey Harold, John Edensor Littlewood, and George Pólya, *Inequalities*, Cambridge University Press, 1952.

Harrington, Joseph E., "A Theory of Collusion with Partial Mutual Understanding," *Research in Economics*, 2017, *71* (1), 140–158.

_ , "How Do Hub-and-Spoke Cartels Operate? Lessons from Nine Case Studies," Working Paper, Available at SSRN 2018.

_ **and Myong-Hun Chang**, "Modeling the Birth and Death of Cartels with an Application to Evaluating Competition Policy," *Journal of the European Economic Association*, 2009, *7* (6), 1400–1435.

_ **and Wei Zhao**, "Signaling and Tacit Collusion in an Infinitely Repeated Prisoners' Dilemma," *Mathematical Social Sciences*, 2012, *64* (3), 277–289.

Hart, Oliver, Jean Tirole, Dennis W. Carlton, and Oliver E. Williamson, "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity. Microeconomics*, 1990, *1990*, 205–286.

He, Sherry, Brett Hollenbeck, and Davide Proserpio, "The Market for Fake Reviews," *Marketing Science*, 2022, *41* (5), 896–921.

Heidhues, Paul and Botond Kőszegi, "Behavioral Industrial Organization," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, chapter 6, pp. 517–612.

Hitsch, Günter J., Ali Hortaçsu, and Dan Ariely, "Matching and Sorting in Online Dating," *American Economic Review*, 2010, *100* (1), 130–163.

Hoffman, Mitchell and Stephen V. Burks, "Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Firm Profits," *Quantitative Economics*, 2020, *11* (1), 315–348.

Holler, Emanuel and Dennis Rickert, "How Resale Price Maintenance and Loss Leading Affect Upstream Cartel Stability: Anatomy of a Coffee Cartel," *International Journal of Industrial Organization*, 2022, *85*, 102871.

Hunold, Matthias and Johannes Muthers, "Manufacturer Cartels and Resale Price Maintenance," Economics Working Papers 2006, Department of Economics, Johannes Kepler University Linz, Austria 2020.

Johnen, Johannes and Robert Somogyi, "Deceptive Features on Platforms," *The Economic Journal*, 2024, *134* (662), 2470–2493.

Johnson, Terence R., "Matching Through Position Auctions," *Journal of Economic Theory*, 2013, *148* (4), 1700–1713.

Jullien, Bruno and Alessandro Pavan, "Information Management and Pricing in Platform Markets," *The Review of Economic Studies*, 2019, *86* (4), 1666–1703.

_ and Patrick Rey, "Resale Price Maintenance and Collusion," *The RAND Journal of Economics*, 2007, *38* (4), 983–1001.

Kang, Zi Yang and Ellen V. Muir, "Contracting and Vertical Control by a Dominant Platform," in "EC" 2022, pp. 694–695.

Kartik, Navin, "Strategic Communication with Lying Costs," *The Review of Economic Studies*, 2009, *76* (4), 1359–1395.

_ , Marco Ottaviani, and Francesco Squintani, "Credulity, Lies, and Costly Talk," *Journal of Economic theory*, 2007, *134* (1), 93–116.

Ke, T. Tony and Yuting Zhu, "Cheap Talk on Freelance Platforms," *Management Science*, 2021, *67* (9), 5901–5920.

Kihlstrom, Richard E and Michael H Riordan, "Advertising as a Signal," *Journal of Political Economy*, 1984, *92* (3), 427–450.

Kircher, Philipp, "Schumpeter Lecture 2022: Job Search in the 21St Century," *Journal of the European Economic Association*, 2022, *20* (6), 2317–2352.

Knapp, Boris, "Fake Reviews and Naive Consumers," *Available at SSRN*, 2022.

Koopmans, Tjalling C. and Martin Beckmann, "Assignment Problems and the Location of Economic Activities," *Econometrica: Journal of the Econometric Society*, 1957, pp. 53–76.

Kreps, David M. and Joel Sobel, "Signalling," *Handbook of Game Theory with Economic Applications*, 1994, *2*, 849–867.

Kroft, Kory and Devin G. Pope, "Does Online Search Crowd out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist," *Journal of Labor Economics*, 2014, *32* (2), 259–303.

Kummer, Michael and Patrick Schulte, "When Private Information Settles the Bill: Money and Privacy in Google's Market for Smartphone Applications," *Management Science*, 2019, *65* (8), 3470–3494.

Laitenberger, Ulrich and Florian Smuda, "Estimating Consumer Damages in Cartel Cases," *Journal of Competition Law & Economics*, 2015, *11* (4), 955–973.

Lauermann, Stephan and Georg Nöldeke, "Stable Marriages and Search Frictions," *Journal of Economic Theory*, 2014, *151*, 163–195.

Lee, Soohyung, "Effect of Online Dating on Assortative Mating: Evidence from South Korea," *Journal of Applied Econometrics*, 2016, *31* (6), 1120–1139.

Levenstein, Margaret C. and Valerie Y. Suslow, "What Determines Cartel Success?," *Journal of Economic Literature*, 2006, *44* (1), 43–95.

Liu, Fang and Tim Paul Thomes, "Manufacturer Collusion-Vertical Integration vs. Delegation with Private Contracts," mimeo 2020.

Mailath, George J., "Incentive Compatibility in Signaling Games with a Continuum of Types," *Econometrica: Journal of the Econometric Society*, 1987, pp. 1349–1365.

_ **and Ernst-Ludwig von Thadden**, "Incentive Compatibility and Differentiability: New Results and Classic Applications," *Journal of Economic Theory*, 2013, *148* (5), 1841–1861.

Malmendier, Ulrike and Geoffrey Tate, "CEO Overconfidence and Corporate Investment," *The Journal of Finance*, 2005, *60* (6), 2661–2700.

_ **and** _ , "Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction," *Journal of Financial Economics*, 2008, *89* (1), 20–43.

Martin, Stephen, Hans-Theo Normann, and Christopher M. Snyder, "Vertical Foreclosure in Experimental Markets," *The RAND Journal of Economics*, 2001, *32* (3), 466–496.

Mas-Colell, Andreu, Michael Dennis Whinston, and Jerry R. Green, *Microeconomic Theory*, Vol. 1, Oxford University Press, New York, 1995.

Mathios, Alan D, "The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market," *The Journal of Law and Economics*, 2000, *43* (2), 651–678.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier, "Promotional Reviews: An Empirical Investigation of Online Review Manipulation," *American Economic Review*, 2014, *104* (8), 2421–2455.

McAfee, Randolph and Marius Schwartz, "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity," *American Economic Review*, 1994, *84* (1), 210–230.

McCutcheon, Barbara, "Do Meetings in Smoke-Filled Rooms Facilitate Collusion?," *Journal of Political Economy*, 1997, *105* (2), 330–350.

McGowan, Müge Adalet and Dan Andrews, "Labour Market Mismatch and Labour Productivity: Evidence from PIAAC Data," OECD Economics Department Working Papers 1209, OECD 2015.

Miklós-Thal, Jeanine and Greg Shaffer, "Naked Exclusion with Private Offers," *American Economic Journal: Microeconomics*, 2016, *8* (4), 174–194.

Milgrom, Paul and John Roberts, "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 1986, *94* (4), 796–821.

Milgrom, Paul R, "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, pp. 380–391.

Mortensen, Dale T. and Christopher A. Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," *The Review of Economic Studies*, 1994, *61* (3), 397–415.

_ and _ , "New Developments in Models of Search in the Labor Market," in "Handbook of Labor Economics," Vol. 3, Elsevier, 1999, chapter 39, pp. 2567–2627.

Morton, Fiona M. Scott, Pascal Bouvier, Ariel Ezrachi, Bruno Jullien, Roberta Katz, Gene Kimmelman, A. Douglas Melamed, and Jamie Morgenstern, "Committee for the Study of Digital Platforms: Market Structure and Antitrust Subcommittee Report," *Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business*, 2019, *36.*

Mueller, Andreas I., Johannes Spinnewijn, and Giorgio Topa, "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias," *American Economic Review*, 2021, *111* (1), 324–363.

Murphy, Sean C., William von Hippel, Shelli L. Dubbs, Michael J. Angilletta Jr., Robbie S. Wilson, Robert Trivers, and Fiona Kate Barlow, "The Role of Overconfidence in Romantic Desirability and Competition," *Personality and Social Psychology Bulletin*, 2015, *41* (8), 1036–1052.

Nocke, Volker and Lucy White, "Do Vertical Mergers Facilitate Upstream Collusion?," *American Economic Review*, 2007, *97* (4), 1321–1339.

_ **and Patrick Rey**, "Consumer Search, Steering, and Choice Overload," *Journal of Political Economy*, 2024, *132* (5), 1684–1739.

Normann, Hans-Theo, "Vertical Integration, Raising Rivals' Costs and Upstream Collusion," *European Economic Review*, 2009, *53* (4), 461–480.

O'Brien, Daniel P. and Greg Shaffer, "Vertical Control with Bilateral Contracts," *The RAND Journal of Economics*, 1992, *23* (3), 299–308.

Ottaviani, Marco and Francesco Squintani, "Naive Audience and Communication Bias," *International Journal of Game Theory*, 2006, *35* (1), 129–150.

Pagnozzi, Marco and Salvatore Piccolo, "Vertical Separation with Private Contracts," *Economic Journal*, 2011, *122* (559), 173–207.

Pearce, David G, "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica: Journal of the Econometric Society*, 1984, pp. 1029–1050.

Pestel, Nico, "Marital Sorting, Inequality and the Role of Female Labour Supply: Evidence from East and West Germany," *Economica*, 2017, *84* (333), 104–127.

Piccolo, Salvatore and Jeanine Miklós-Thal, "Colluding through Suppliers," *The RAND Journal of Economics*, 2012, *43* (3), 492–513.

_ **and Markus Reisinger**, "Exclusive Territories and Manufacturers' Collusion," *Management Science*, 2011, *57* (7), 1250–1266.

Pissarides, Christopher A, "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages," *The American Economic Review*, 1985, *75* (4), 676–690.

Prat, Andrea and Tommaso Valletti, "Attention Oligopoly," *American Economic Journal: Microeconomics*, 2022, *14* (3), 530–557.

Ramey, Garey, "D1 Signaling Equilibria with Multiple Signals and a Continuum of Types," *Journal of Economic Theory*, 1996, *69* (2), 508–531.

Reillier, Laure Claire and Benoit Reillier, *Platform Strategy: How to Unlock the Power of Communities and Networks to Grow Your Business*, Routledge, 2017.

Rey, Patrick and Jean Tirole, "A Primer on Foreclosure," in Mark Armstrong and Robert Porter, eds., *Handbook of Industrial Organization*, 1 ed., Vol. 3, Elsevier, 2007, chapter 33, pp. 2145–2220.

__ **and Thibaud Vergé**, "Bilateral Control with Vertical Contracts," *The RAND Journal of Economics*, 2004, *35* (4), 728–746.

__ **and** __ , "Resale Price Maintenance and Interlocking Relationships," *The Journal of Industrial Economics*, 2010, *58* (4), 928–961.

Rhodes, Andrew and Chris M Wilson, "False Advertising," *The RAND Journal of Economics*, 2018, *49* (2), 348–369.

Robles, Theodore F., Richard B. Slatcher, Joseph M. Trombello, and Meghan M. McGinn, "Marital Quality and Health: A Meta-Analytic Review," *Psychological Bulletin*, 2014, *140* (1), 140.

Rochet, Jean-Charles and Jean Tirole, "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 2003, *1* (4), 990–1029.

__ **and** __ , "Two-Sided Markets: A Progress Report," *The RAND Journal of Economics*, 2006, *37* (3), 645–667.

Rosenfeld, Michael J., Reuben J. Thomas, and Sonia Hausen, "Disintermediating your Friends: How Online Dating in the United States Displaces other Ways of Meeting," *Proceedings of the National Academy of Sciences*, 2019, *116* (36), 17753–17758.

Roth, Alvin E., "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 1982, *7* (4), 617–628.

__ **and Marilda Sotomayor**, "Two-Sided Matching," *Handbook of Game Theory with Economic Applications*, 1992, *1*, 485–541.

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante, "Mismatch Unemployment," *American Economic Review*, 2014, *104* (11), 3529–3564.

Schinkel, Pieter Maarten, Jan Tuinstra, and Jakob Rüggeberg, "Illinois Walls: How Barring Indirect Purchaser Suits Facilitates Collusion," *The RAND Journal of Economics*, 2008, *39* (3), 683–698.

Schirrmacher, Nina-Birte, Jan Ondrus, and Thomas Kude, "Launch Strategies of Digital Platforms: Platforms with Switching and Non-Switching Users," Technical Report 2017.

Segal, Ilya, "Contracting with Externalities," *Quarterly Journal of Economics*, 1999, *114* (2), 337–388.

Selten, Reinhard, "A Simple Model of Imperfect Competition, where 4 are Few and 6 are Many," *International Journal of Game Theory*, 1973, *2* (1), 141–201.

Shapley, Lloyd S. and Martin Shubik, "The Assignment Game I: The Core," *International Journal of Game Theory*, 1971, *1* (1), 111–130.

Sharabi, Liesel L. and Elizabeth Dorrance-Hall, "The Online Dating Effect: Where a Couple Meets Predicts the Quality of their Marriage," *Computers in Human Behavior*, 2024, *150*, 107973.

Shimer, Robert and Lones Smith, "Assortative Matching and Search," *Econometrica*, 2000, *68* (2), 343–369.

Smith, Lones, "The Marriage Model with Search Frictions," *Journal of Political Economy*, 2006, *114* (6), 1124–1144.

Spinnewijn, Johannes, "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," *Journal of the European Economic Association*, 2015, *13* (1), 130–167.

Srinivasan, Karthik, "Paying Attention," Working Paper, Booth School of Business, University of Chicago 2023.

Wets, Roger J-B., "On the Continuity of the Value of a Linear Program and of Related Polyhedral-Valued Multifunctions," *Mathematical Programming Essays in Honor of George B. Dantzig Part I*, 1985, pp. 14–29.

Yasui, Yuta, "Controlling Fake Reviews," Working Paper, Available at SSRN 2020.

Zhang, Daiqiang, "Testing Passive versus Symmetric Beliefs in Contracting with Externalities," *International Economic Review*, 2021, *62* (2), 723–767.

Heinrich Heine Universität Düsseldorf

Eidesstattliche Versicherung

Ich, Frau Jana Gieselmann, versichere an Eides statt, dass die vorliegende Dissertation von mir selbstständig und ohne unzulässige fremde Hilfe unter Beachtung der „Grundsätze zur Sicherung guter wissenschaftlicher Praxis an der Heinrich-Heine-Universität Düsseldorf“ erstellt worden ist.

Düsseldorf, der 21. Mai 2025

J. Gieselmann

Unterschrift