J. Econ. Theory 220 (2024) 105883

Contents lists available at ScienceDirect

Journal of Economic Theory

journal homepage: www.elsevier.com/locate/jet

Sorting versus screening in decentralized markets with adverse selection ☆

Sarah Auster [a,*], Piero Gottardi [b,c]

[a] Department of Economics, University of Bonn, Germany
[b] Department of Economics, Essex University, United Kingdom of Great Britain and Northern Ireland
[c] University of Venice, Italy

ARTICLE INFO

Keywords:
Competitive search
Adverse selection
Market segmentation

ABSTRACT

We study the role of traders' meeting capacities in decentralized markets with adverse selection. Uninformed principals choose trading mechanisms to find an agent with whom to trade. Agents are privately informed about their quality and aim to match with one of the principals. We consider a rich set of meeting technologies and characterize the properties of the equilibrium allocations for each of them. In equilibrium, different agent types can be separated either via sorting—they self-select into different submarkets—or screening within the trading mechanism, or a combination of the two. We show that, as the meeting capacity increases, the equilibrium features more screening and less sorting. Interestingly, this increases price variability and reduces the average quality of trade as well as the total level of trade in the economy. The trading losses are, however, more than compensated by savings in entry costs so that welfare increases.

1. Introduction

The recent years have seen significant innovations in the way potential trading partners can contact each other in decentralized markets. A key role has been played by the internet and the development of a variety of platforms to facilitate meetings between market participants. In the case of services, a distinctive and novel feature of several recently emerged platforms is that they allow customers to specify their needs, after which providers are invited to submit bids. For instance, on Thumbtack and HomeAdvisor one can indicate the kind of work that she or he needs to be done at home—carpentry, plumbing, electricity, lawn maintenance—providing various details on the task. The platform then shares those details with multiple professionals in the area who will compete for the provision of the service by providing a quote and other information. Other websites operate similarly for different services, such as moving,[1] or translation and editing services (Upwork). The examples above refer to procurement problems, albeit simple, where the customer seeks the provision of a specific, customized service. In most of these situations, a crucial element will be the quality of the provided service, about which the customer may have limited information. While these markets are clearly decentralized, platforms put customers in touch with a number of providers, who can then compete over the terms at which the service is provided.

☆ We would like to thank Ronald Wolthoff and various seminar participants for very helpful comments. Sarah Auster gratefully acknowledges funding from the German Research Foundation (DFG) through Germany's Excellence Strategy—EXC 2126/1—390838866 and CRC TR 224 (Project B02). Piero Gottardi gratefully acknowledges financial support from PRIN 2022 grant 2022PXE3B7.

* Corresponding author.
E-mail addresses: sauster@uni-bonn.de (S. Auster), piero.gottardi@essex.ac.uk (P. Gottardi).

[1] See, e.g., https://www.moving.com/movers/moving-company-quotes.asp.

https://doi.org/10.1016/j.jet.2024.105883
Received 25 July 2022; Received in revised form 10 July 2024; Accepted 21 July 2024

Available online 25 July 2024

The objective of this paper is to present a theoretical framework that allows to analyze the impact of these developments on the level and quality of the transactions occurring in markets with characteristics as the ones described above. Our framework has three key features. First, the market is decentralized, meaning that meetings between market participants are subject to search frictions. Second, agents may meet multiple potential counter-parties and can use this feature to enhance competition among them via the design of the trading mechanism. Third, the market exhibits adverse selection: some agents are privately informed about their quality and quality affects the valuation of both trading parties. Hence, values are interdependent, as in Akerlof (1970). For concreteness, we will cast the model in terms of service procurement, but we emphasize that the framework is more general. In particular, as we discuss in more detail below, our findings apply to other markets featuring search frictions, many-to-one meetings, and adverse selection, such as labor markets and decentralized financial markets.

We consider a directed search model of a market where uninformed customers seek to hire a provider whose quality can be high or low. We allow customers to post general direct mechanisms that specify trading probabilities and transfers for participating providers, contingent not only on their own reported type but also on the number and the reported types of other providers meeting the same customer. This feature can be interpreted as customers choosing one among several available platforms, characterized by different trading mechanisms. Given the posted mechanisms, each provider picks one of them in order to find a match. The number of providers meeting a specific customer is random and its distribution depends on the market tightness—i.e. the relative number of customers and providers choosing a given mechanism—as well as on the meeting technology. The latter reflects the extent to which a customer is able to be contacted by several providers at the same time.

To clarify ideas, under a meeting technology with no capacity constraints, a customer can meet with all the providers who attempt to contact her, for instance by letting all of them participate in a suitably designed auction. In contrast, under a meeting technology with very tight capacity constraints, a customer can only meet a single provider at a given point in time (as in the case of bilateral meetings). In this paper, we consider a rich set of meeting technologies, ranging between these two extremes. The only restriction is that a customer's probability of meeting a particular provider decreases in the number of other providers attempting to meet the same customer. Hence, we assume that there is some crowding out in the meeting process, though the degree to which providers affect each other's chances of meeting a given customer can be arbitrarily small or large.

The properties of the meeting technology matter for how in equilibrium customers discriminate privately informed providers of different quality. In general, such discrimination may occur via *sorting*, that is, via market segmentation with different mechanisms posted in equilibrium to attract different types of agents, or via *screening*, that is, within the mechanism (Eeckhout and Kircher, 2010). In the presence of search frictions, sorting is costly, as market fragmentation increases the likelihood that potential trading partners remain unmatched. Screening within the trading mechanism, on the other hand, is hampered by the presence of crowding out in the meeting technology since a mechanism can discriminate only those agents that are present in the same meeting. With adverse selection, trade is not only constrained by search frictions but also by incentive compatibility. This, as we will show, has important, distinctive implications for the properties of equilibrium allocations and for the effects of varying the degree of crowding out in the meeting technology.

We first establish some general properties of the equilibrium allocation. We show that for all considered meeting technologies, equilibria satisfy the following three features. First, whenever a customer meets at least one provider, the customer will trade. Hence, there is no rationing within the trading mechanisms that are posted in equilibrium, in contrast to what one might see in screening menus or in auctions with reserve prices. Second, the incentive constraints of low-quality providers always bind in the posted mechanisms. Third, these mechanisms exhibit the property that whenever the customer meets at least one low-quality provider, she will trade low quality. Hence, equilibrium trading mechanisms perfectly discriminate: among the providers a customer meets, the low-quality ones are granted priority. Notably, this property does not depend on the relative number of low-quality providers in the economy and holds even when the relative gains from trade with high-quality providers are arbitrarily large.

In the second part of the paper, we consider a parametrized family of meeting technologies, proposed by Cai et al. (2022), where the degree to which providers crowd each other out in the meeting process is captured by a single parameter σ. This family allows for all possible levels of crowding out, with $\sigma = 0$ corresponding to the case of bilateral meetings (complete crowding out) and $\sigma = 1$ corresponding to the case where customers can meet all the providers that select them (no crowding out). We characterize the properties of equilibrium allocations as a function of σ and of the fraction of low-quality providers in the population.

We show that there are three relevant parameter regions of equilibria. When σ is below a threshold $\sigma^S > 0$, perfect market segmentation occurs in equilibrium, with low- and high-quality providers trading in separate submarkets. The equilibrium thus exhibits pure sorting, as in the case of bilateral meetings (see Guerrieri et al. (2010)). Since, in this type of equilibrium, the pool of providers any given customer attracts is homogeneous, the trading mechanisms take the form of posted prices. As σ increases and the customers' capacity to meet providers expands, screening different types of providers within the trading mechanism becomes more cost-effective. For $\sigma > \sigma^S$, screening is sufficiently profitable so that the equilibrium always features a market that attracts both types of providers, who are then screened via the trading mechanism. The mechanisms in this market now take the form of an auction. In such auctions, low-quality providers always outbid high-quality providers due to their lower costs, so the latter get to trade only when no low-type providers participate in the auction. If the fraction of low-type providers in the population is sufficiently large, there is a second submarket active in equilibrium, attracting only low-type providers. Because of the homogeneity of the pool, the trading mechanism in the second market again takes the form of a posted price. The equilibrium thus has a mix of sorting and screening. Conversely, if the fraction of low-type providers is small, all providers search in a single submarket and trade via the same type of auction. In this case, the equilibrium exhibits screening but no sorting.

We characterize the implications of changes in σ for the level and pattern of trade, entry, and welfare in equilibrium. To this end, we focus on values of σ where the equilibrium features both sorting and screening. As σ increases, an increasing number of low-type providers migrate from the submarket with price posting to the submarket where both types are present and the trading mechanism is an auction. Since low-quality providers beat high-quality ones in the auctions of the latter market, this migration reduces the trading probability of high-type providers as well as the average quality of trade. We show that welfare increases despite these effects: high-quality providers receive higher payments, which more than compensate the decrease in trading probability. The payoff of low-type providers, instead, remains constant,[2] as does the payoff of customers, who always break even in equilibrium. An increase in σ, therefore, constitutes a Pareto improvement. To understand why welfare improves in spite of the decline in average quality and quantity of trades, the customer's entry decision has to be taken into account. We argue that fewer customers enter the market when σ increases and that the reduction in entry costs contributes positively to welfare. We conclude the analysis with a discussion of the general welfare properties of equilibrium allocations, showing that they are always constrained inefficient (in ex-ante terms) due to an excessive degree of separation in trading probabilities of the two types.

As already mentioned, the implications of our analysis go beyond the market for service procurement. In the context of labor markets, our results may shed light on the evidence presented by Hall and Krueger (2012) and by Batra et al. (2023), showing that a large fraction of job postings contain only limited information about wages, with the aim of capturing broader sets of applicants. As (Batra et al., 2023, pp. 3-4) observe, 'the explosion of internet job search may have paradoxically reduced the amount of wage information available." Our analysis shows that, in situations where workers' skills are not easy to observe but are important to employers, the selection of workers is more likely to occur via screening in the application process than via sorting through wage postings when meetings become easier. Also OTC markets have experienced important transformations in the last years, with a rise of platforms where trades occur via request for quote protocols (RFQ). On these platforms, customers can ask for quotes from multiple dealers, making the trading mechanism similar to an auction. This situation arises as a property of the equilibrium in our setting when meeting capacities expand. Our findings thus help to explain recent developments in both of these markets.

Summing up, the paper's contribution is twofold. First of all, we provide new economic insights on how improvements in the meeting technology and in the ability to process multiple requests to trade affect market segmentation, trade patterns, and welfare in economies with adverse selection. Our analysis shows that improvements in the meeting technology lead informed agents to be screened more often *within the trading mechanism* rather than *across submarkets.* While this generates welfare gains, as one might expect, the mechanism through which these gains are realized is subtle: the gradual consolidation of the market arising from improvements in the meeting technology is beneficial *not* because it increases the number of high-surplus trades but because it allows to save on entry costs. Furthermore, total trading volume decreases, and high-quality types end up trading less, though at higher prices. Hence price variability increases. As we argue below, these effects are driven by the interdependence of values and thus significantly differ from the effects of improvements in the meeting technology on equilibrium outcomes found in earlier work with independent private values.

Second, on the methodological side, we provide a tractable approach to derive the search equilibrium allocation in settings where equilibria may not be constrained efficient and, therefore, cannot be found by decentralizing the planner's solution. The approach is sufficiently flexible to accommodate arbitrary meeting technologies and arbitrary sets of feasible trading mechanisms. In such a framework, we develop a novel method to establish the existence of a competitive search equilibrium. The proof relies on a fixed-point argument, where the key step is to reinterpret providers' market utilities as competitive prices and simplify the customers' design problem by reducing it to a standard demand problem. The proof of existence is then akin to that commonly used for equilibria in Walrasian markets. Due to its generality, we believe that this method is applicable well beyond our setting.

Related literature Eeckhout and Kircher (2010) were the first to highlight the fact that the properties of the meeting technology have important consequences for equilibrium outcomes and sorting patterns. They consider a market with independent private values and compare the two extreme cases on which most of the literature has focused its attention[3]: bilateral meetings (complete crowding out) and urn-ball meetings, where the principal can meet all the agents that selected her (no crowding out). They show that in the first case, the equilibrium features complete market segmentation with principals posting different prices, while in the second case, all principals post the same mechanism, given by a second price auction, discriminating agents within the mechanism. In a similar independent private value environment, Cai et al. (2022) allow for a rich set of meeting technologies and demonstrate that the equilibrium may feature a mix of sorting and screening. We show that the consideration of adverse selection leads to important changes in the sorting patterns, which are driven by the interaction of competition forces with incentive constraints. We should highlight that, while in their setting improvements in the meeting technology lead to increases both in the volume and in the average surplus per trade, the opposite happens in our setup. We discuss these differences in more detail in Section 4.

A number of papers have investigated the properties of competitive search equilibria in markets with adverse selection. Focusing on the case of bilateral meetings, Gale (1992), Inderst and Müller (2002), Guerrieri et al. (2010) showed that an equilibrium always exists and features complete market segmentation.[4] The case of urn-ball meetings and general mechanisms was instead analyzed by Auster and Gottardi (2019).[5] They prove the existence of an equilibrium in which all principals post the same screening mechanism (no sorting). In the current paper, we build on the formalization developed in this previous work and extend it to a rich class of

[2] Low-quality providers can be viewed as being in excess supply here since not all of them can be attracted in the submarket where high-quality providers trade, while still satisfying incentive compatibility.

[3] E.g., Moen (1997), Shimer (2005), Albrecht et al. (2014).

[4] As investigated by Chang (2018) and Williams (2021), this property may not extend to the case of multidimensional private information.

[5] Also Auster et al. (2022) consider urn-ball meetings and adverse selection. The key innovation is to let agents contact multiple principals, whose choice over trading mechanisms is, however, restricted to posted prices.

meeting technologies. While this class nests bilateral and urn-ball meetings as special cases, our main focus lies on more realistic technologies with non-extreme meeting capacities. The analysis of this richer setting poses several challenges. In particular, the existence of an equilibrium can no longer be shown using a constructive argument but requires a new, more general approach. The equilibrium characterization allows us to study the implications of small and large changes in customers' ability to meet counterparts for equilibrium outcomes, such as sorting patterns, market entry, and welfare.

Finally, we should also mention that the effects of allowing one-to-many meetings have also been investigated in random search models by considering the case where an agent may meet more than one principal (in contrast, we assume that a principal may meet multiple agents). Equilibrium implications in those models are rather different from the ones found in this paper, as they concern primarily the principals' bargaining power. Building on the earlier work by Burdett and Judd (1983), Lester et al. (2019) have carried out a systematic and very interesting analysis of these effects for adverse selection economies.

2. Environment

There is a measure one of providers and a large measure of homogeneous customers. Each provider offers a service of uncertain quality. Quality is identically and independently distributed across providers. It can be either high or low and is private information of the provider. The fraction of providers that offer a low-quality service is μ. Customers can freely enter the market at a cost K. The valuation/cost of customers and providers for the low-quality service is denoted, respectively, by $\underline{v}$ and $\underline{c}$, and that for the high-quality service by $\overline{v}$ and $\overline{c}$. We assume customers care for quality, and high quality is more costly: $\overline{v} > \underline{v}$, $\overline{c} > \underline{c}$. Furthermore, there are always positive gains from trade, with those for the high-quality service being weakly greater than those for the low-quality service:

$$\overline{v} - \overline{c} \geq \underline{v} - \underline{c} > K.$$

Meeting technologies The trading process operates as follows. Customers simultaneously choose mechanisms that specify how trade takes place with the providers they meet. Providers observe the chosen mechanisms and select one of the mechanisms they prefer, as well as one of the customers offering it. We refer to the collection of customers and providers choosing the same mechanism as constituting a submarket. The probability for a customer to meet a certain number of providers in a given submarket depends on the queue length λ in the submarket, defined as the ratio of the measure of providers to the measure of customers present in that submarket.[6] Let the probability that a customer meets $n = 0, 1, 2, \ldots$ providers as a function of the queue length λ be denoted by $P_n(\lambda)$.

We assume that the meeting probability does not depend on the providers' types; that is, conditional on a customer meeting n providers, the types of these providers are n independent draws from the population of providers present in the submarket. It thus follows that the probability for a customer to meet at least one provider of a certain type when the queue of that type is λ' and the overall queue length in the submarket is $\lambda \geq \lambda'$ is:

$$\phi(\lambda', \lambda) := 1 - \sum_{n=0}^{+\infty} P_n(\lambda) \left(1 - \frac{\lambda'}{\lambda}\right)^n.$$

The probability $\phi(\lambda', \lambda)$ is the complement to the probability of meeting no provider or meeting only providers of a different type. As shown in Cai et al. (2022), under the assumption that meeting probabilities are type-independent, the function ϕ fully characterizes the meeting technology.

The two cases that have received most attention in the literature are bilateral meetings and urn-ball meetings. The urn-ball matching technology captures a situation where, in any submarket, each provider always meets one randomly selected customer, and each customer meets all the providers who selected her (the customer has no capacity constraint in her ability to meet providers). Hence, a provider is certain to meet the customer he selects but does not know how many other providers show up. Moreover, a customer's probability of meeting a provider of a given type is fully determined by the queue length of that type of provider in the submarket, while it does not depend on the presence of other types of providers. The class of meeting technologies that have this property is called 'invariant' (Lester et al., 2015). In contrast, under the bilateral meeting technology, each customer can meet at most one provider. This means that—opposite to urn-ball matching—a provider is not certain to meet the selected customer and the presence of other providers negatively affects his meeting chances.

In this paper, we are interested in more general situations where customers can meet multiple providers—as under urn-ball matching—but where providers visiting a submarket impose externalities on the meeting possibilities of other providers in that market—as under bilateral matching. The idea is that customers face constraints in their ability to meet providers, but their capacity constraint needs not to be equal to one.

Mechanisms and payoffs For the specification of mechanisms, payoffs, and the definition of equilibrium, we follow closely Auster and Gottardi (2019). A direct mechanism m is a map from the number of low and high messages sent by the providers meeting a given customer to a set of trading probabilities and transfers conditional on having sent a low or high message. Given the queue lengths of low and high types in the submarket, $\underline{\lambda}$ and $\overline{\lambda}$, the probability of trade under mechanism m when sending a low message is denoted by $\underline{x}_m(\underline{\lambda}, \overline{\lambda})$. Given $(\underline{\lambda}, \overline{\lambda})$, this probability is obtained from the probability distribution over the number and types of other providers

[6] In principle, we can have a continuum of active submarkets, each with a measure zero of customers and providers. In that case, we can use the Radon-Nikodym derivatives to define queue lengths.

meeting the same customer, as specified by the meeting technology, and the probability of trade, as specified by the mechanism, assuming that all other providers report truthfully. The trading probability associated to sending the high message, $\overline{x}_m(\underline{\lambda},\overline{\lambda})$, and the expected transfers, $\underline{t}_m(\underline{\lambda},\overline{\lambda})$ and $\overline{t}_m(\underline{\lambda},\overline{\lambda})$, are defined analogously. We derive these expressions formally in Appendix A.1.

The expected payoff for low- and high-type providers when choosing mechanism m and revealing their type truthfully, net of the utility when they do not trade, is then given by:

$$\underline{u}(m|\underline{\lambda},\overline{\lambda}) = \underline{t}_m(\underline{\lambda},\overline{\lambda}) - \underline{x}_m(\underline{\lambda},\overline{\lambda})\underline{c},$$
$$\overline{u}(m|\underline{\lambda},\overline{\lambda}) = \overline{t}_m(\underline{\lambda},\overline{\lambda}) - \overline{x}_m(\underline{\lambda},\overline{\lambda})\overline{c}.$$

The incentive compatibility constraints, ensuring the optimality of truthful reporting, are as follows:

$$\underline{t}_m(\underline{\lambda},\overline{\lambda}) - \underline{x}_m(\underline{\lambda},\overline{\lambda})\underline{c} \geq \overline{t}_m(\underline{\lambda},\overline{\lambda}) - \overline{x}_m(\underline{\lambda},\overline{\lambda})\underline{c}, \tag{1}$$
$$\overline{t}_m(\underline{\lambda},\overline{\lambda}) - \overline{x}_m(\underline{\lambda},\overline{\lambda})\overline{c} \geq \underline{t}_m(\underline{\lambda},\overline{\lambda}) - \underline{x}_m(\underline{\lambda},\overline{\lambda})\overline{c}. \tag{2}$$

The expected payoff for a customer posting mechanism m, when providers report truthfully and the expected number of high- and low-type providers, respectively, is $\overline{\lambda}$ and $\underline{\lambda}$, is then:

$$\pi(m|\underline{\lambda},\overline{\lambda}) = \overline{\lambda}(\overline{x}_m(\underline{\lambda},\overline{\lambda})\overline{v} - \overline{t}_m(\underline{\lambda},\overline{\lambda})) + \underline{\lambda}(\underline{x}_m(\underline{\lambda},\overline{\lambda})\underline{v} - \underline{t}_m(\underline{\lambda},\overline{\lambda})).$$

Equilibrium An allocation in this setting is defined by a probability measure β on the set of feasible mechanisms M with support M^β, where $\beta(M')$ describes the measure of customers that post mechanisms in $M' \subseteq M$, and two maps $\underline{\lambda},\overline{\lambda} : M^\beta \to \mathbb{R}^+$ specifying, respectively, the queue lengths of low- and high-type providers selecting mechanism m. We say that an allocation is *feasible* if:

$$\frac{1}{\beta(M^\beta)} \int_{M^\beta} \frac{\underline{\lambda}(m)}{\underline{\lambda}(m)+\overline{\lambda}(m)} d\beta(m) = \mu. \tag{3}$$

Likewise, we say an allocation is *incentive compatible* if, for all $m \in M^\beta$, (1) and (2) hold.

An equilibrium is then given by a feasible and incentive compatible allocation such that the values of β and $\underline{\lambda},\overline{\lambda}$ are consistent with the optimal choices of customers and providers.[7] The beliefs of customers over the queue lengths of low- and high-type providers that mechanisms not offered in equilibrium would attract are described by extending the maps $\underline{\lambda},\overline{\lambda}$ to the domain $M\backslash M^\beta$. As is standard in the literature (e.g. Guerrieri et al. (2010), Auster and Gottardi (2019)), we require these beliefs to satisfy a consistency condition with providers' optimal choices out of equilibrium that is analogous to the one holding for mechanisms posted in equilibrium. More specifically, a customer believes that a deviating mechanism attracts some low-type providers only if low-type providers are indifferent between this mechanism and the one they choose in equilibrium, and similarly for high-type providers. The definition of a directed search equilibrium is stated informally below, while a formal description can be found in Appendix A.2.

Definition 1. A directed search equilibrium is a feasible and incentive compatible allocation, given by a measure β with support M^β and two maps $\underline{\lambda},\overline{\lambda} : M \to \mathbb{R}^+ \cup +\infty$, such that the following conditions hold:

- Customers choose incentive compatible mechanisms in M to maximize expected payoffs, given their beliefs. For all $m \in M^\beta$, these payoffs are equal to K.
- Providers choose mechanisms in M^β to maximize their expected payoffs.
- Beliefs for all mechanisms in $M\backslash M^\beta$ are consistent in the sense described above.

3. Analysis

A distinguishing feature of the environment with adverse selection is that the equilibrium cannot be found by decentralizing the constrained planner's solution. To characterize directed search equilibria we must therefore study the optimal choices of customers when they compete among them. To this end, it is convenient to rewrite the customers' optimization problem in terms of an auxiliary problem. The latter consists in choosing the expected trading probabilities, expected transfers, and queue lengths for each type so as to maximize profits, taking as given the utility gain attained in the market by low- and high-type providers relative to their endowment points, denoted by $\underline{U}$ and $\overline{U}$ (in short, their market utilities).[8] Auster and Gottardi (2019) demonstrated the validity of this approach for the case of urn-ball meetings. We now extend the result to the general class of meeting technologies considered here.

3.1. A customer's auxiliary optimization problem

Recalling the definition of the function ϕ, we can express the probability that a given customer meets at least one low-type provider in a submarket with a low-type queue length $\underline{\lambda}$ and a high-type queue length $\overline{\lambda}$ by $\phi(\underline{\lambda},\underline{\lambda}+\overline{\lambda})$, the probability of meeting at least one

[7] The restriction to incentive compatible allocations is w.l.o.g.
[8] In equilibrium, $\underline{U}$ and $\overline{U}$ coincide with the expected payoff obtained by low and high type providers.

high-type provider by $\phi(\overline{\lambda}, \underline{\lambda} + \overline{\lambda})$, and the probability of meeting at least one provider by $\phi(\underline{\lambda} + \underline{\lambda}, \underline{\lambda} + \overline{\lambda}) = 1 - P_0(\underline{\lambda} + \overline{\lambda})$. We then have[9]:

Proposition 1. *Consider an equilibrium where the set of posted mechanisms is M^β. For every $m \in M^\beta$ with queue lengths $\underline{\lambda} = \underline{\lambda}(m), \overline{\lambda} = \overline{\lambda}(m)$, the vector*

$$\left(\underline{x}_m(\underline{\lambda}, \overline{\lambda}), \overline{x}_m(\underline{\lambda}, \overline{\lambda}), \underline{t}_m(\underline{\lambda}, \overline{\lambda}), \overline{t}_m(\underline{\lambda}, \overline{\lambda}), \underline{\lambda}, \overline{\lambda}\right)$$

solves the problem

$$\max_{\underline{x}, \overline{x}, \underline{t}, \overline{t}, \underline{\lambda}, \overline{\lambda}} \overline{\lambda}(\overline{x}\,\overline{v} - \overline{t}) + \underline{\lambda}(\underline{x}\,\underline{v} - \underline{t}) \tag{P^{aux}}$$

subject to:

$$\underline{t} - \underline{x}\,\underline{c} \leq \underline{U} \quad \textit{holding with equality if } \underline{\lambda} > 0, \tag{4}$$

$$\overline{t} - \overline{x}\,\overline{c} \leq \overline{U} \quad \textit{holding with equality if } \overline{\lambda} > 0, \tag{5}$$

$$\underline{t} - \underline{x}\,\underline{c} \geq \overline{t} - \overline{x}\,\underline{c}, \tag{6}$$

$$\overline{t} - \overline{x}\,\overline{c} \geq \underline{t} - \underline{x}\,\overline{c}, \tag{7}$$

$$\underline{\lambda}\,\underline{x} \leq \phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}), \tag{8}$$

$$\overline{\lambda}\,\overline{x} \leq \phi(\overline{\lambda}, \underline{\lambda} + \overline{\lambda}), \tag{9}$$

$$\overline{\lambda}\,\overline{x} + \underline{\lambda}\,\underline{x} \leq \phi(\underline{\lambda} + \overline{\lambda}, \underline{\lambda} + \overline{\lambda}). \tag{10}$$

To understand the statement of the auxiliary problem in the proposition, note that an equilibrium mechanism attracting low- or high-type providers must yield these providers a payoff equal to their market utility. Hence, the associated trading probabilities and transfers must satisfy constraints (4)-(5). Likewise, they must satisfy the providers' incentive constraints, which correspond to conditions (6)-(7). The remaining three conditions say that the mechanism is feasible, in the sense that total trading probabilities do not exceed one,[10] and that meetings take place according to technology ϕ. In particular, inequality (8) requires that a customer's probability of trading with a low-type provider, $\underline{\lambda}\,\underline{x}$, is weakly smaller than a customer's probability of meeting at least one low-type provider, $\phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda})$. Inequalities (9) and (10) are the analogous conditions for the high-type provider and overall trade.

While it is easy to see that any feasible and incentive compatible equilibrium mechanism must satisfy conditions (4)-(10), the main role of Proposition 1 is to show that equilibrium mechanisms must solve problem P^{aux}. This means that in order to find an equilibrium we can directly study the solutions of P^{aux}. Whenever the values of $\underline{U}$ and $\overline{U}$ are such that the queue lengths at the solutions of the customer's auxiliary problem satisfy the property that the average fraction of low- and high-type providers equals the fraction in the population, these solutions give us the expected trading probabilities and transfers in an equilibrium.

3.2. General properties of search equilibria

We impose next some basic conditions on the meeting technology.

Assumption 1. The meeting technology function ϕ is twice differentiable in both arguments and satisfies[11]

$$\frac{\partial \phi}{\partial \overline{\lambda}}(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}) < 0, \qquad \frac{\partial^2 \phi}{\partial \overline{\lambda}^2}(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}) > 0.$$

Assumption 1 says that providers impose some negative externalities on each other in the meeting process. The first inequality in Assumption 1 says that having additional high-type providers in a submarket reduces the chances that a customer meets some low-type providers. The second inequality says that the magnitude of this effect is decreasing in the queue length of the high types. This assumption is violated under the urn-ball meetings technology, where an increase in the number of providers attempting to meet a given customer has no effect on the probability of meeting that customer of the providers who already selected him. Under Assumption 1 instead, customers face some constraints in their ability to meet providers, so there is a chance that providers crowd each other out in the meeting process.

Given these conditions, we can derive our first main result:

[9] All proofs can be found in the Appendix.

[10] See (14) in Appendix A.1 for a formal statement of this condition.

[11] As explained above, the function ϕ is symmetric with respect to the type appearing in the first entry, so the condition equally applies to $\phi(\overline{\lambda}, \underline{\lambda} + \overline{\lambda})$.

Proposition 2. *Let Assumption 1 be satisfied. A directed search equilibrium exists. In any equilibrium, the posted mechanisms satisfy the following properties:*

1. *All meetings lead to trade.*
2. *In any meeting where low-type providers are present, the customer will trade with one of them.*
3. *Low-type incentive constraints are binding.*

We start by discussing each of the three stated properties that equilibrium mechanisms must satisfy. Using these properties, we then explain how the existence of a directed search equilibrium is established.

No rationing The first property says that in equilibrium customers only offer mechanisms under which every meeting ends in trade. For instance, if a customer offers an auction in equilibrium, the reserve price of the auction is always acceptable to all providers that show up. This property is in contrast with what happens when customers and providers meet according to the urn-ball technology. In that case, we showed in our previous work (Auster and Gottardi, 2019) that when the fraction of low-quality agents is sufficiently high, multiple payoff equivalent equilibria co-exist, only one of which has the property that every meeting leads to trade, whereas all the others exhibit rationing of high types within the mechanism. Proposition 2 demonstrates that when there are some negative externalities in the meeting process, no equilibrium exists that features rationing within the mechanism. Put differently, if customers are constrained in their ability to meet providers—even slightly—they do not find it optimal to attract providers when they may end up not trading with any of them. This implies that within the set of equilibria described in Auster and Gottardi (2019), all but one equilibrium are non-robust to introducing a small level of a negative externality in the meeting process. To establish the 'no rationing' property in our setting, we show that under Assumption 1, at any solution of P^{aux}, a customer's probability to trade equals her probability of meeting a provider; that is, the overall feasibility constraint (10) holds as equality.

Priority for low-type providers The second feature that all equilibrium mechanisms satisfy is that high-type providers get to trade only in meetings where no low-type provider is present. Formally, we show that, at any solution of P^{aux}, the low-type feasibility constraint (8) holds as equality: a customer's probability of buying a low-quality service is equal to her probability of meeting a low-type provider. Note this is true even though gains from trade are larger for high types. An important step for proving the above result is the observation that the incentive compatibility of equilibrium mechanisms requires that $\underline{U} > \overline{U}$. That is, the expected utility gain obtained in equilibrium by trading in the market is strictly higher for low-type providers than for high-type providers (see Appendix A.4). Given this property, any mechanism that does not give priority to low-type providers is dominated by one that does. In particular, for any non-priority mechanism, consider an alternative mechanism that attracts fewer low types and more high types but gives priority to low types so that the overall probability of trading with a low- and high-type provider remains unchanged for the customer. At the alternative mechanism, providers get the same utility as with the mechanisms posted in the market, but the customer's profits are higher. Since high-quality providers have a lower utility gain in equilibrium than low types, they are less costly to attract. Deviating to the alternative mechanism is reminiscent of cream skimming. It should be noted that this result does not rely on the crowding out condition on the meeting technology imposed in Assumption 1, it also holds for the urn-ball technology.

Low-type incentive constraint binds Proposition 2 further establishes that in equilibrium the low-type incentive constraint (6) is satisfied as equality. If this property is violated, for any mechanism attracting both types, a customer could increase her payoff by replacing some low-type providers with the same number of high-type providers. Such replacement keeps the customer's probability of trade unchanged but weakly increases the gains from trade that are generated (recall: $\overline{v} - \overline{c} \geq \underline{v} - \underline{c}$) and reduces the rents paid to the providers since $\underline{U} > \overline{U}$. It also increases a high-type provider's probability of trade but, as long as the low-type incentive constraint is slack, this is not an issue, provided the replacement is sufficiently small. We are then left with the possibility of incentive constraints being slack at a separating equilibrium, where low- and high-type providers search in separate submarkets. This would require customers to be indifferent between both markets. Since, however, high-quality providers generate higher gains from trade and have a lower market utility, indifference can only be sustained when the incentive compatibility constraints for low-type providers bind and limit trade in the high-quality market.

Existence Given the three properties we established, a customer's auxiliary problem can be written in the following simplified way:

$$\max_{\underline{\lambda},\overline{\lambda}} \quad \phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})(\underline{v}-\underline{c}) + [\phi(\underline{\lambda}+\overline{\lambda}, \underline{\lambda}+\overline{\lambda}) - \phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})](\overline{v}-\overline{c}) - \underline{\lambda}\,\underline{U} - \overline{\lambda}\,\overline{U} \tag{11}$$

subject to

$$[\phi(\underline{\lambda}+\overline{\lambda}, \underline{\lambda}+\overline{\lambda}) - \phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})](\overline{c}-\underline{c}) = \overline{\lambda}(\underline{U}-\overline{U}). \tag{12}$$

The expression in (11) is the expected payoff of a customer for a mechanism that leads to trade in every meeting with priority for low types, as a function of the queue lengths of low- and high-type providers. With probability $\phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})$, the customer meets some low-type provider, and the gains from trade from the low-quality service are realized, while with probability $\phi(\underline{\lambda}+\overline{\lambda}, \underline{\lambda}+\overline{\lambda}) - \phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})$, the customer meets no low-type provider but some high-type provider, so the gains from trade of the high-quality service are realized. The rent the customer has to pay to the providers she attracts is $\underline{\lambda}\,\underline{U} + \overline{\lambda}\,\overline{U}$. For all $\overline{\lambda} > 0$, the value of the low-type queue length $\underline{\lambda}$ is then determined by the low-type providers' incentive constraint (12). For $\overline{\lambda} = 0$, equation (12) is always satisfied and thus imposes no restrictions on $\underline{\lambda}$.

To prove an equilibrium exists, we need to show that we can find values of $\underline{U}, \overline{U}$ such that the maximum of (11) subject to (12) is equal to K and a measure β can be assigned to the set of solutions so that the feasibility condition (3) is satisfied. We establish the existence of such values using Kakutani's fixed point theorem. The argument is similar to that of proofs for the existence of Walrasian equilibria, with queue lengths $\underline{\lambda}, \overline{\lambda}$ taking the role of quantities consumed and market utilities $\underline{U}, \overline{U}$ taking the role of prices. The argument provides a novel and general method that could be used also in other settings to prove the existence of directed search equilibria when efficiency is not guaranteed and hence the equilibrium cannot be found as a solution of a constrained planner problem.[12]

4. Meeting ability and market segmentation

In the previous section, we established the existence of a search equilibrium and several properties of equilibrium mechanisms, which hold for all meeting technologies in the rich set we are considering, characterized by Assumption 1. In this section, we turn to the question of how the meeting technology ϕ affects other important properties of the search equilibrium, such as market segmentation, trading volumes and market entry. To address this question, we restrict attention to a class of meeting technologies where customers have a stochastic meeting capacity (see Cai et al. (2022)) and the degree of crowding out in the meeting process is captured by a single parameter $\sigma \in [0, 1]$:

$$\phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}) = \frac{\underline{\lambda}}{1 + \sigma\underline{\lambda} + (1 - \sigma)(\underline{\lambda} + \overline{\lambda})}. \tag{13}$$

It is instructive to consider two special cases. For $\sigma = 0$, the customer's probability of meeting at least one low-type provider can be conveniently written as follows:

$$\phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}) = \frac{\underline{\lambda} + \overline{\lambda}}{1 + \underline{\lambda} + \overline{\lambda}} \frac{\underline{\lambda}}{\underline{\lambda} + \overline{\lambda}} = \phi(\underline{\lambda} + \overline{\lambda}, \underline{\lambda} + \overline{\lambda}) \frac{\underline{\lambda}}{\underline{\lambda} + \overline{\lambda}},$$

that is, as the probability of meeting at least one provider multiplied with the probability that a randomly selected provider is of low type. Meetings are thus bilateral. On the other hand, for $\sigma = 1$, the probability of meeting at least one low-type provider, $\phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}) = \frac{\underline{\lambda}}{1+\underline{\lambda}}$, only depends on the low-type queue length but not on the queue length of the high type. In that case, the meeting technology is invariant, which means that high- and low-type providers do not crowd each other out (as it happens under urn-ball matching). Hence Assumption 1 boils down to the restriction $\sigma < 1$. Note that, given the queue lengths $\underline{\lambda}$ and $\overline{\lambda}$, the probability of meeting *some* provider is unaffected by the parameter σ: $\phi(\underline{\lambda} + \overline{\lambda}, \underline{\lambda} + \overline{\lambda}) = (\underline{\lambda} + \overline{\lambda})/(1 + \underline{\lambda} + \overline{\lambda})$. Hence, varying σ changes the expected number of providers in a meeting without changing the total number of meetings.

The class of meetings technologies given by (13) with $\sigma \in [0, 1)$ is still quite rich and allows for different levels of crowding out in meetings, with crowding out decreasing in σ. Restricting attention to this class, the following lemma shows that any search equilibrium features at most two active submarkets. Moreover, only one of these markets attracts high-type providers.

Lemma 3. *Assume ϕ satisfies* (13) *with $\sigma < 1$. A search equilibrium is of one of the following three types:*

1. *Pure sorting: two submarkets are active, one attracts low-type providers, the other attracts high-type providers.*
2. *Partial sorting: two submarkets are active, one attracts low-type providers, the other attracts both types of providers.*
3. *Pure screening: only one submarket is active and it attracts both types of providers.*

To prove the lemma, we consider the simplified customer's auxiliary problem, stated in (11). Recall this problem consists of maximizing the customer's expected payoff over the queue lengths $\underline{\lambda}$ and $\overline{\lambda}$ subject to the low-type incentive constraint (12). As we argued above, the latter constraint is satisfied as equality at any solution of the auxiliary problem with $\overline{\lambda} > 0$. Using this property and the functional form of ϕ in (13), we can solve the binding constraint for the queue length $\underline{\lambda}$ as a function of $\overline{\lambda} > 0$ in order to write the customer's expected payoff as a function of a single variable $\overline{\lambda}$. We then show in the proof that this payoff is strictly concave in $\overline{\lambda}$, implying that the customer's auxiliary problem can have at most one solution with $\overline{\lambda} > 0$. The other potential solution lies at $\overline{\lambda} = 0$.

Which of the three types of equilibria obtains depends on the extent to which customers are constrained in their ability to meet providers, i.e., on the value of the parameter σ. As shown in Proposition 2, whenever in equilibrium customers choose a mechanism that attracts both types of providers, the mechanism will treat the two types differently, giving priority to low-type providers. The effectiveness of such screening clearly depends on the properties of the meeting technology captured by σ. When σ is close to 0, that is, capacity is tight, the priority rule does not allow for a substantial differentiation of the trading probability of the two types (and hence the price a customer pays). The alternative way to differentiate providers is to gear the terms of trade to one of the two types

[12] In contrast, both Guerrieri et al. (2010) and Auster and Gottardi (2019) use a constructive argument to establish existence of an equilibrium, respectively, with pairwise and urn-ball meetings.

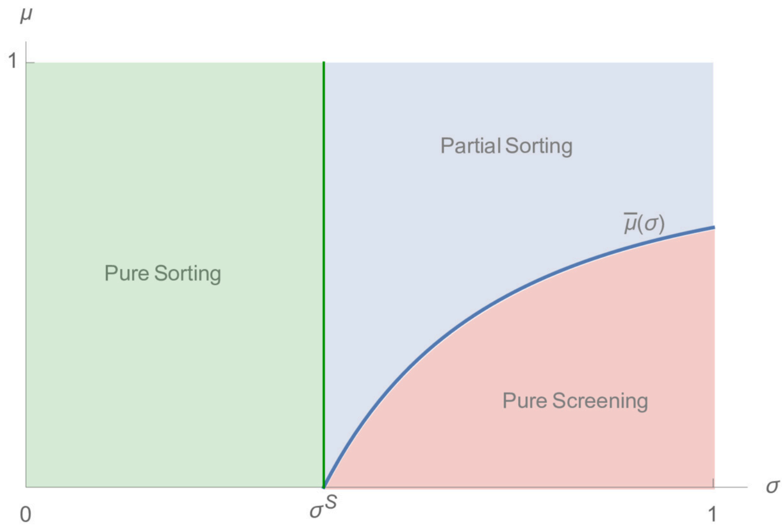

Fig. 1. Equilibrium market segmentation for the following parameter specification: $\overline{v} = 1.2$, $\underline{v} = 1$, $\overline{c} = 0.1$, $\underline{c} = 0$, $K = 0.1$.

so that the providers sort themselves across different markets. In this case, the mechanisms offered are simple posted prices, i.e., customers commit to trade at the announced price with one randomly selected provider among the ones they meet. This is effectively the same as having bilateral meetings, where one of the providers attempting to meet a given customer is randomly selected for the meeting. With different markets active in equilibrium, there is no limit to the difference that can be achieved in the trading probabilities of the two types, as the queue lengths in the two active markets can be suitably varied. Market segmentation is however costly, as it reduces the total number of meetings.[13] Building on these considerations, we show the following.

Proposition 4. *Assume ϕ satisfies (13) with $\sigma < 1$. There exists a parameter $\sigma^S \in (0, 1)$ and a function $\overline{\mu} : (\sigma^S, 1] \to [0, 1]$ satisfying $0 < \overline{\mu}(1) < 1$ such that:*

1. *if $\sigma \leq \sigma^S$, there is a pure sorting equilibrium;*
2. *if $\sigma > \sigma^S$ and $\mu > \overline{\mu}(\sigma)$, there is a partial sorting equilibrium;*
3. *if $\sigma > \sigma^S$ and $\mu \leq \overline{\mu}(\sigma)$, there is a pure screening equilibrium.*

Fig. 1 illustrates, for all admissible values of the parameters σ and μ, which of the three different types of equilibria described in Proposition 4 is guaranteed to exist.[14] When σ is below the threshold σ^S, there exists an equilibrium where mechanisms take the form of posted prices. Providers sort themselves into separate submarkets according to their type, with high-quality providers trading at a higher price in a submarket with a sufficiently longer queue, as required by incentive compatibility. Hence, customers are more likely to meet a provider in the high-quality market than in the low-quality market. For customers to be indifferent between the two markets, they must then make higher profits with low types than with high types conditional on trading.

An increase in σ augments the number of providers that are present in a given meeting, but as long as $\sigma \leq \sigma^S$, it has no effect on the equilibrium allocation: posted prices and queue lengths are constant in σ, as are the market utilities of low- and high-type providers. What changes with σ in this parameter region is the profitability of a customer's deviation to a screening mechanism. In particular, as σ increases above σ^S, the following deviation becomes profitable: customers attracting high-quality providers post a mechanism that also attracts some low-quality providers and grants them priority in meetings. The advantage of this mechanism is that it generates a positive probability of contracting with a low-quality provider at a relatively low price while keeping the option to hire a high-quality provider with high probability in case no low-quality provider shows up. While priority allows the low types to jump the queue in meetings where the other providers are of high type, for low values of σ, the incidence of meetings with multiple providers is relatively low. Hence, in such situations, the priority rule has little bite in enhancing the probability of trade for low

[13] It can be easily verified that when ϕ is as in (13), for any $\gamma_1, \gamma_2, \lambda_1, \lambda_2 > 0$, we have $\gamma_1\phi(\lambda_1, \lambda_1) + \gamma_2\phi(\lambda_2, \lambda_2) \leq (\gamma_1 + \gamma_2)\phi((\gamma_1\lambda_1 + \gamma_2\lambda_2)/(\gamma_1 + \gamma_2))$, holding as equality if and only if $\lambda_1 = \lambda_2$.

[14] Note that we cannot rule out the possibility that multiple equilibria exist. The candidate partial sorting equilibrium and the candidate pure screening equilibrium are determined by a solution to the first-order conditions of the auxiliary problem and the customers' free entry condition and such a solution may not be unique. We do know, however, that pure sorting equilibria can only exist to the left of σ^S, where σ^S is uniquely defined. Likewise, we can define for any candidate partial sorting equilibrium a map $\overline{\mu}(\sigma)$. When $\mu < \overline{\mu}(\sigma)$ for all such maps, no partial sorting equilibrium exists, in which case we know, by the existence result in Proposition 2, that a pure screening equilibrium exists. We should also point out that when $\sigma = 1$, in addition to the equilibria with partial sorting and pure screening—which are the limits of the equilibria we found as σ approaches 1—there is a continuum of payoff equivalent equilibria with rationing, as shown in Auster and Gottardi (2019). (Recall Assumption 1 is violated in that case, so Proposition 4 does not apply.)

types. As σ increases, the priority rule becomes more effective, which reduces the compensation customers have to pay to low-type providers for the longer queue in the high-type market, making the deviation more profitable.

For values of σ above σ^S, the equilibrium always features some degree of pooling. In particular, as stated in Proposition 4, all customers offering mechanisms that attract high-type providers also attract some low-type providers. Two possibilities then arise. First, only a fraction of low-type providers choose the market selected by high types (where the mechanism takes the form of an auction), while the remaining low types continue to trade via posted prices in a separate market. The equilibrium is thus of the partial sorting type, where the different types of providers are separated partly via their choice of mechanism and partly within the chosen mechanism. The second possibility is that all customers and providers choose the same mechanism, so separation only occurs within this mechanism. As shown in Proposition 4 and illustrated in Fig. 1, which of the two kinds of equilibria obtains depends both on the values of σ and μ.

To understand the role played by the meeting technology, note that for values of σ slightly above σ^S the deviation from the pure-sorting equilibrium described earlier, consisting of attracting some low-quality providers in the high-quality market, is profitable but only slightly. Hence, for σ slightly above σ^S, we can expect that the fraction of low-type providers present in the market where high-type providers trade is small. As σ increases, the priority rule becomes more effective, which implies that customers can afford to attract more low-type providers to the market chosen by the high-type providers and still guarantee a sufficiently high trading probability for the low types. Hence, as σ increases, more and more low-type providers migrate from the low-quality market to the market where both types of services are traded.

To determine whether the equilibrium features partial sorting or pure screening, we need to examine the role played by the fraction μ of low types in the population. Consider a candidate partial sorting equilibrium: in this equilibrium, a positive mass of low types still chooses the low-quality market. Terms of trade in that market do not vary with σ, so the low types' market utility $\underline{U}$ is constant. For any given $\sigma > \sigma^S$, let us indicate with $\overline{\mu}(\sigma)$ the mass of low-quality providers choosing the mixed market, or, equivalently, the demand for low types by customers who, in this candidate equilibrium, choose to operate in the mixed market when the cost of attracting low types is $\underline{U}$. When $\mu > \overline{\mu}(\sigma)$, the candidate equilibrium allocation with partial sorting is admissible, as the mass of low types in the population exceeds the demand for low types coming from customers active in the mixed market. In the threshold case where $\mu = \overline{\mu}(\sigma)$, such candidate equilibrium allocation is still admissible, but there are no providers left in the low-quality market. Hence, the equilibrium features pure screening. Instead, when the mass of low types in the population is sufficiently low, so that $\mu < \overline{\mu}(\sigma)$, the demand from the mixed market for low types exceeds the supply, so the candidate equilibrium allocation with partial sorting is not admissible. In this case, the equilibrium is of the pure screening type, but with a higher level of the low-type market utility $\underline{U}$.

Finally, according to Proposition 4, each of the three types of equilibria exists for some values of the parameters μ, σ, as in the situation depicted in Fig. 1. That is, there are three non-empty subsets of the region of admissible values of $\sigma, \mu \in [0,1) \times (0,1)$, where, respectively, a pure sorting, partial sorting and pure screening equilibrium exists. This partitioning of the parameter space clearly demonstrates the effects of the customers' meeting capacity and the severity of adverse selection on the main qualitative properties of equilibria. The providers' equilibrium trading probabilities and transfers always differ for the two types, but how this difference in treatment is achieved, via sorting and/or screening, depends on the values of μ and σ.

4.1. *Entry and welfare*

In what follows, we want to analyze in more detail how the properties of the equilibrium allocation vary with the characteristics of the meeting technology and, in particular, the degree of crowding out captured by σ. To this end, we will focus our attention on the parameter region where the equilibrium features both sorting and screening. Within this region, changes in σ modify the equilibrium allocation in various important ways: they affect the level of sorting across the two markets, the entry-level of customers (so far implicit in the analysis), and the terms of trade in the market where both types trade.[15] In contrast, the terms of trade in the low-quality market do not vary with σ.

Our next proposition shows that, in the region where the equilibrium features partial sorting, an increase in σ leads to a reduction of the high type's trading probability. Despite the lower trading probability, the market utility of high-quality providers increases. Hence, high types are overcompensated for the loss in trading probability by the price in the event of trade. In contrast, the welfare of low-type providers does not vary with σ.

Proposition 5. *Assume ϕ satisfies (13). Let $\sigma \in (\sigma^S, 1)$ and $\mu > \overline{\mu}(\sigma)$. A marginal increase in σ leads to:*

- *a decrease of the trading probability $\overline{x}$ and an increase of the utility gain $\overline{U}$ attained in equilibrium by high-type providers;*
- *no change of the utility gain $\underline{U}$ attained by low-type providers.*

To gain some intuition for this result, recall that higher values of σ increase the profitability of attracting low-type providers to the high-type market, due to the increased effectiveness of the priority rule. Operating in the market where high types trade thus becomes more attractive for customers. As a result, the competition for high-type providers gets fiercer, and this drives up their market utility $\overline{U}$ (to the point that customers still break even in both markets). There is instead no increased competition for low types, as long as

[15] In the pure screening region, only the last two effects are present, while in the pure sorting case, the allocation is invariant to σ.

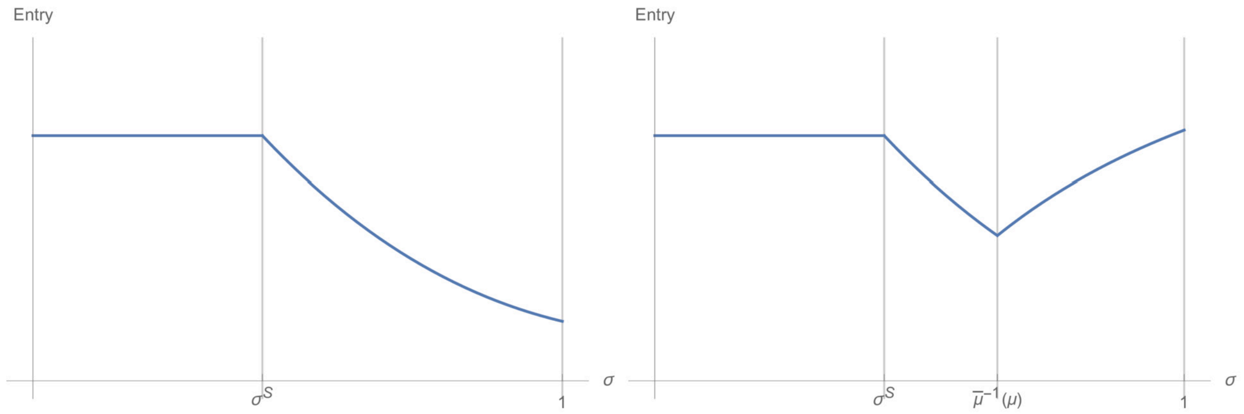

Fig. 2. Entry levels for $\overline{v} = 1.2$, $\underline{v} = 1$, $\overline{c} = 0.1$ and $\underline{c} = 0$, $k = 0.1$, with $\mu = 0.7$ (left panel) and $\mu = 0.4$ (right panel).

$\mu > \overline{\mu}(\sigma)$, since the demand for them coming from customers operating in the mixed market is less than the supply of low types in the economy. The market utility $\underline{U}$ thus remains unchanged. This implies that in the parameter region where we have partial sorting, an increase in the ability to meet potential counterparts, as described by a rise in σ, induces a Pareto improvement: high-type providers gain, while the utility of low-type providers and customers is unaffected.

It is interesting to point out that this occurs even though the expected surplus generated from high-quality trades decreases. Indeed, as customers attract an increasing mass of low-type providers in the mixed market and give them priority, the trading probability of high types declines. To understand why welfare increases despite this decline, we need to consider the effect of changes in σ on the level of entry by customers in the two markets: since entry is costly, any change in the probability of trade for providers should be assessed against the change in the total entry costs borne by customers.

The effect of σ on entry in the partial sorting region can be seen most clearly when we compare the equilibrium allocation at the boundaries of the parameter region where $\sigma = \sigma^S$ and $\sigma = 1$ (assuming that the fraction of low types in the population is sufficiently high: $\mu > \overline{\mu}(1)$). At the boundary where $\sigma = \sigma^S$, the measure of low-type providers in the mixed market is zero, so the allocation is of the pure sorting type. At the boundary where $\sigma = 1$, providers impose no negative externalities on each other in the meeting process, which means that the presence of high-type providers has no effect on the low-type providers' trading probability, given that they receive priority through the mechanism. We can then show that for $\sigma = 1$, the two following properties hold in equilibrium (see Appendix A.11): first, due to the high effectiveness of the priority rule, customers attract the same queue length of low-type providers in the mixed and in the low-quality market, thereby guaranteeing them the same probability of trade; second, competition for high-type providers is so fierce that customers' profits with them (gross of entry costs) are driven all the way down to zero.

Note that the queue length in the low-type market does not change with σ in the partial sorting region. This, together with the property that at $\sigma = 1$ the low-type queue length is the same in both active submarkets, implies that the allocation at $\sigma = 1$ can be viewed as if those customers and providers active in the low-quality market in a pure sorting equilibrium are split proportionally across two submarkets, whereas high-type providers are assigned to just one of them. Hence, when we move from $\sigma \leq \sigma^S$ to $\sigma = 1$, entry is unambiguously reduced by an amount equal to the mass of customers who, in the case of low σ, operate in the high-quality market. This reallocation leaves the trading probability of low types unchanged, while it reduces the trading probability of high types (see Proposition 5). Hence, both the level of trade (as measured by the average trading probability of providers) and the average quality of trades are lower at $\sigma = 1$ than at $\sigma \leq \sigma^S$. Yet, from Proposition 5, we know that welfare goes up. The key source of the welfare improvement thus lies in the reduction of the customers' entry costs.

Fig. 2 illustrates, for a numerical example, the effect of improvements in the meeting technology on entry across the different parameter regions. The left panel shows the case where $\mu > \overline{\mu}(1)$. For σ sufficiently low, entry is constant because the allocation in the pure sorting equilibrium does not depend on σ. When we enter the region of partial sorting, entry decreases monotonically with σ, as explained above. The right panel illustrates the case $\mu < \overline{\mu}(1)$, where at $\sigma = 1$, the mass of low types in the economy is insufficient to accommodate the demand coming from the mixed market. Under this specification, all three equilibrium types exist as σ is varied from 0 to 1. We see that entry is now non-monotonic in σ. As before, it is first constant and then decreasing in σ, up to the point where $\mu = \overline{\mu}(\sigma)$. For higher values of σ, the equilibrium is of the pure screening type, and entry increases with σ. To gain some intuition, note that the measure of low and high types in the single market remains constant when σ increases beyond the point where $\mu = \overline{\mu}(\sigma)$. As the congestion in the meeting technology becomes lower, the (more profitable) low-cost providers in the market are crowded out by high types less often. The market thus becomes more profitable and attracts a larger number of customers, which then increases competition for both types of providers. This drives up the utility gain attained by high-type providers—as before—as well as that attained by low-type providers, who are no longer in excess supply.

It is useful to briefly comment on the general welfare properties of the equilibrium allocation we characterized. It is clear that the equilibrium allocation does not maximize surplus subject to the feasibility constraints imposed by the meeting technology: since gains from trade are higher for the high- than for the low-type provider, surplus maximization would require a higher probability of trade for high types. This, however, violates the incentive constraint, which, as we saw, requires that the trading probability of low-quality providers is weakly higher than that of high-quality providers. The next result shows that even if incentive constraints, individual rationality constraints, and search frictions are taken into account, a social planner can always generate more surplus than the one attained in equilibrium.

Proposition 6. *Assume ϕ satisfies (13) with $\sigma < 1$ and $\overline{v} - \overline{c} > \underline{v} - \underline{c}$. The equilibrium allocation never maximizes the total surplus generated by trade subject to incentive compatibility, individual rationality, and feasibility constraints.*

The result is established by showing first that when the search equilibrium features pure or partial screening, the allocation can be improved upon by relaxing the priority given to low-type providers in the pooling market. This increases the trading probability of high-quality providers without changing the overall volume of trade and hence leads to an increase in the total gains from trade. The incentive compatibility of this change is assured by appropriately decreasing the utility gained by high types (always possible since, in equilibrium, we have $\bar{U} > 0$). At a fully separating equilibrium, instead, the increase in surplus is attained by reallocating customers from the low- to the high-quality market so that total entry is unchanged. It is easy to verify that, since this reallocation makes the queue lengths in the two markets closer to each other, it increases the total probability of trade. It also increases the average gains from trade and, hence, total surplus. Again, incentive compatibility can be assured by appropriately lowering the utility of high-quality providers.

The above result shows that it is always possible to increase total surplus by locally changing the equilibrium allocation so as to increase the high-type providers' trading probability at the margin. Of course, the surplus can be increased further by considering larger changes in the probabilities of trade as well as in customers' entry levels. A complete characterization of the optimal allocation (subject to the constraints in Proposition 6) is beyond the scope of this paper,[16] but it is easy to verify that the allocation with complete pooling—that is, a single market with no priority—and associated optimal entry level maximizes total surplus subject to feasibility and incentive constraints. This follows directly from the previous observations that incentive compatibility requires the high-type providers' trading probability to be less than or equal to that of low-type providers, while surplus maximization implies that the high type's trading probability should be maximal subject to this and the feasibility constraints. The pooling allocation is thus an optimum as long as the high-type providers' individual rationality constraint can be satisfied, that is, as long as the value of the maximal common payment from customers to providers is greater than or equal to $\bar{c}$. When the latter constraint is violated (this happens when the high-type provider's cost is sufficiently high and the fraction of high-type providers is sufficiently low), a planner maximizing total surplus needs to give high-quality providers a strictly lower trading probability than low-quality ones. How the high-type provider's trading probability is optimally reduced will then depend on the meeting technology. In particular, for high values of σ and moderate values of μ, the use of partial or full priority may be sufficient to reduce the high type's trading probability adequately. For sufficiently low values of σ, on the other hand, the planner may have to resort to separating out some providers and assigning them to another submarket. In either case, the differentiation of trading probabilities between the two types is less than what occurs in the search equilibrium. The inefficiency thus stems from 'an excessive level of separation' generated by the forces of competition.

4.2. Higher gains from trade for low types

Throughout the analysis, we assumed that $\overline{v} - \overline{c} \geq \underline{v} - \underline{c}$, i.e., there are higher gains from trade when contracting with a high-quality provider than with a low-quality provider. While this assumption seems natural for many applications, its main role in our analysis is to guarantee that the low type's incentive constraint is binding. In fact, up to Proposition 4, we only used the property $\overline{v} - \overline{c} \geq \underline{v} - \underline{c}$ to show such constraint is binding. This implies that the characterization presented in Proposition 4 holds for all parameter values, including $\overline{v} - \overline{c} < \underline{v} - \underline{c}$, as long as the low type's incentive constraint binds. If, on the other hand, this incentive constraint becomes slack, the properties of the equilibrium change. This happens when $\overline{v}$ is sufficiently close to $\underline{v}$ and, in particular, when the two values are the same.

The case $\overline{v} = \underline{v}$ (independent private values) has been studied by Eeckhout and Kircher (2010) and Cai et al. (2022), as already mentioned in the Introduction. Eeckhout and Kircher (2010) contrast the two extreme cases of bilateral meetings and urn-ball meetings, showing that equilibria feature pure sorting in the first case and pure screening in the second. Cai et al. (2022) complete the analysis of this environment by characterizing the equilibrium properties for a rich class of meeting technologies, including the one with stochastic capacity considered here. Like us, they show the possibility of partial sorting for intermediate values of σ, but the sorting pattern is reversed: high-cost providers are now the ones active in both submarkets, while low-cost providers only search in one of the two markets. Interestingly, their proofs do not rely on the fact that the uninformed party's valuation is the same for both types but only on the gains from trade with low-cost providers being higher than those with high-cost providers and on the incentive constraint being slack.[17] Hence, while Cai et al. (2022) focus on the independent value case, their analysis perfectly complements ours, as it allows us to obtain a characterization for the case $\overline{v} - \overline{c} < \underline{v} - \underline{c}$ when incentive constraints are slack. Appendix A.10 extends

[16] Davoodalhosseini (2019) examines the welfare properties of competitive search equilibria with adverse selection for the case of bilateral meetings. In Section 4, he provides a characterization of the surplus maximizing allocation for the case where valuations are binary (as we assume) and the meeting technology is efficient in the sense that the short side in a market is matched with probability one.

[17] In their setting, the privately informed agents are the buyers, while the principals are the sellers. To facilitate the comparison with the previous analysis, in this discussion, we keep referring to principals as customers and to agents with private information as providers.

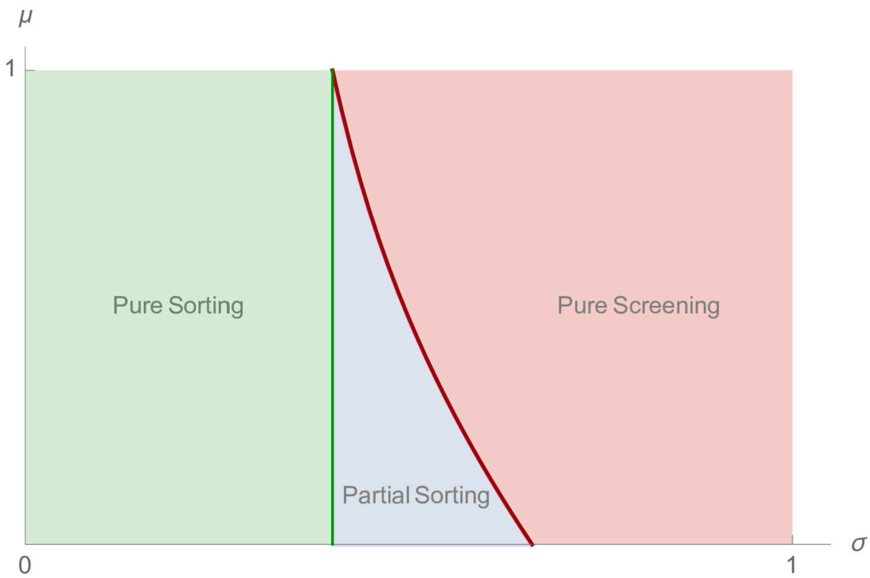

Fig. 3. Equilibrium market segmentation for the case $\overline{v} - \overline{c} < \underline{v} - \underline{c}$ and slack incentive constraints.

their results to our environment with interdependent values and free entry by providing the equilibrium conditions and characterizing the boundaries of the different equilibrium regions for the case where incentive constraints are slack.[18]

In Fig. 3, we illustrate the different parameter regions for this case. As we increase the value of σ, we move again from pure sorting to pure screening, this time transitioning through a region where high-cost providers are separated across two markets while low-cost providers are concentrated in one of them. As we move through the partial sorting region, an increased number of high-cost providers migrate to the mixed market. The larger the share of high-cost providers in the economy (the lower μ), the longer it takes to complete the transition and arrive at the pure screening region. As a result, the interval of values of σ for which we have partial sorting is larger when μ is smaller.

When incentive constraints are non-binding, the changes in the equilibrium outcome as the parameters μ and σ vary can be understood by taking into account the constrained efficiency property of the equilibrium. When $\sigma = 1$, the absence of crowding out means that all low-cost providers who have selected a given customer have priority in trade, while high-cost providers simply serve as insurance against the event that no low-cost provider is present. Recalling that, when all customers choose the same trading mechanism, the total number of meetings increases relative to the case where the market is segmented (keeping entry unchanged), it follows that overall trade is higher when compared to the case of segmented markets. If instead $\sigma < 1$, high-cost providers crowd out low-cost providers when both are present in the same submarket. Still, for σ not too low, attracting some high-cost providers in addition to the low-cost ones is beneficial, as it provides a hedge to customers when no low-cost provider shows up while only slightly lowering the probability of trade with low-cost providers, who are given priority in trade. For σ sufficiently low, on the other hand, the priority rule is too ineffective, so a higher level of surplus is attained if customers attract only one type of provider. While it is again true that welfare increases when the meeting technology improves, the mechanism leading to this is quite different from the one we saw in the previous sections: in the case of independent values, the volume and average surplus generated by trade increase, whereas, as we saw, both decrease when incentive constraints bind.

In markets with (sufficiently severe) adverse selection, incentive constraints bind and equilibrium outcomes are no longer determined by efficiency considerations but rather by the forces of competition among customers and these incentive constraints. This changes the logic behind partial sorting in our setting. Indeed, attracting some high types in the low-type market, as in the case we just discussed, would conflict with incentive compatibility. What becomes a profitable deviation instead, when one-to-many meetings are possible with limited crowding out, is to attract some low-quality providers on top of high-quality ones. This is to generate the possibility of trading with low-cost providers, exploiting the relatively long queue of the high-cost market. The profitability of the deviation is linked to the equilibrium feature whereby trade with low-quality providers remains more profitable for customers, even though a larger surplus is generated by trade with high-quality ones.

5. Conclusion

This paper studies the question of how recent technological innovations regarding the ability of counterparties to meet affect trading outcomes in decentralized markets with adverse selection. To address the issue, we examine a theoretical framework where

[18] Cai et al. (2022) consider a fixed measure of market participants on both sides, whereas we have free entry. This does not, however, affect the qualitative properties of the equilibrium. For any $K \in (0, \min\{\underline{v} - \underline{c}, \overline{v} - \overline{c}\})$, there is an associated measure γ of customers, such that the equilibrium outcome in our model is the same as when the measure of customers is exogenously fixed to γ.

uninformed customers offer general trading mechanisms in order to contract a provider for a service. The extent to which a customer in such markets can exploit competition between different providers offering the service depends on the meeting technology, in particular, the customer's capacity to meet multiple providers at the same time. We show that when customers have a limited ability to meet multiple providers, in equilibrium, markets are perfectly segmented, with low-quality providers searching for low-paying tasks and high-quality providers looking for high-paying tasks. As the meeting capacity increases, an increasing number of low-quality providers, as well as some customers, migrate to the submarket where high-quality providers search, while other customers leave the market altogether. The improvement of the meeting technology thus reduces both the average quality of trades and the total level of trades occurring in the market. Despite the reduction in quality and quantity of equilibrium trades, total welfare increases, as trading losses are overcompensated by the reduction in the number of customers entering the market and the corresponding savings in entry costs.

CRediT authorship contribution statement

Sarah Auster: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization.
Piero Gottardi: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

None.

Data availability

No data was used for the research described in the article.

Appendix A. Proofs

A.1. Mechanisms, trading probabilities, and transfers

A direct mechanism m is defined by the map:

$$\left(\underline{X}_m, \overline{X}_m, \underline{T}_m, \overline{T}_m\right) : \mathbb{N}^2 \to [0,1]^2 \times \mathbb{R}^2,$$

where the first argument is the number of low messages and the second argument is the number of high messages received by a customer from the providers he meets. The maps $\underline{X}_m(L,H), \overline{X}_m(L,H)$ describe the trading probabilities specified by mechanism m for providers reporting, respectively, to be of low and of high type. The associated transfers (unconditional on whether or not trade occurs) are given by the maps $\underline{T}_m(L,H), \overline{T}_m(L,H)$. For instance, $\underline{T}_m(L,H)$ is the transfer to a provider who reports to be of low type when $L-1$ other providers report to be of low type and H other providers report to be of high type. We say a mechanism m is feasible if:

$$\underline{X}_m(L,H)L + \overline{X}_m(L,H)H \leq 1, \forall (L,H) \in \mathbb{N}^2. \tag{14}$$

Let M denote the measurable set of feasible mechanisms. We assume that when matched with a customer, a provider does not observe how many other providers are actually matched with the same customer or their types.

Given a meeting technology, the queue lengths of low- and high-type providers, $\underline{\lambda}$ and $\overline{\lambda}$, induce a joint distribution over the number of low types and high types meeting a customer. We define

$$P(L,H;\underline{\lambda},\overline{\lambda}) \equiv P_n(\underline{\lambda}+\overline{\lambda})\frac{(L+H)!}{L!H!}\left(\frac{\underline{\lambda}}{\underline{\lambda}+\overline{\lambda}}\right)^L\left(\frac{\overline{\lambda}}{\underline{\lambda}+\overline{\lambda}}\right)^H$$

as the probability for a customer to meet L low types and H high types when the queue lengths of low- and high-type providers are, respectively, $\underline{\lambda}$ and $\overline{\lambda}$. On that basis, we determine the expected trading probabilities of the two types of providers induced by the mechanism:

$$\underline{x}_m(\underline{\lambda},\overline{\lambda}) = \sum_{L,H} P(L,H;\underline{\lambda},\overline{\lambda})\underline{X}(L+1,H),$$

$$\overline{x}_m(\underline{\lambda},\overline{\lambda}) = \sum_{L,H} P(L,H;\underline{\lambda},\overline{\lambda})\overline{X}(L,H+1).$$

Analogously, we can determine expected transfers $\underline{t}_m(\underline{\lambda},\overline{\lambda})$ and $\overline{t}_m(\underline{\lambda},\overline{\lambda})$.

A.2. Formal definition of a search equilibrium

We restrict out-of-equilibrium beliefs by the following conditions:

$$\underline{u}(m|\underline{\lambda}(m), \overline{\lambda}(m)) \leq \max_{m' \in M^{\beta}} \underline{u}(m'|\underline{\lambda}(m'), \overline{\lambda}(m')) \text{ holding with equality if } \underline{\lambda}(m) > 0, \tag{15}$$

$$\overline{u}(m|\underline{\lambda}(m), \overline{\lambda}(m)) \leq \max_{m' \in M^{\beta}} \overline{u}(m'|\underline{\lambda}(m'), \overline{\lambda}(m')) \text{ holding with equality if } \overline{\lambda}(m) > 0. \tag{16}$$

This specification of the beliefs for mechanisms that are not posted in equilibrium is standard in the literature of directed search, both with and without adverse selection (see Guerrieri et al. (2010) and Eeckhout and Kircher (2010), among others).

In our environment, conditions (15), (16) may not uniquely pin down the out-of-equilibrium beliefs $\underline{\lambda}(m), \overline{\lambda}(m)$ for $m \notin M^{\beta}$. To this end we postulate the following[19]:

i) if (15), (16) admit a unique solution, $\underline{\lambda}(m)$ and $\overline{\lambda}(m)$ are given by that solution;
ii) if (15), (16) admit no solution, we set $\underline{\lambda}(m)$ and/or $\overline{\lambda}(m)$ equal to $+\infty$ and $\pi(m|\underline{\lambda}(m), \overline{\lambda}(m)) = c$ for some $c \leq 0$;
iii) if (15), (16) admit multiple solutions, $\underline{\lambda}(m), \overline{\lambda}(m)$ are given by the solution for which the customer's payoff $\pi(m|\underline{\lambda}(m), \overline{\lambda}(m))$ is the highest.

Our formal definition of a directed search equilibrium is then given by:

Definition 2. A directed search equilibrium is a feasible and incentive compatible allocation, given by a measure β with support M^{β} and two maps $\underline{\lambda}, \overline{\lambda} : M \to \mathbb{R}^{+} \cup +\infty$, such that the following conditions hold:

- customer optimality: for all $m \in M$ such that $(m, \underline{\lambda}(m), \overline{\lambda}(m))$ satisfies incentive compatibility,

$$\pi(m|\underline{\lambda}(m), \overline{\lambda}(m)) \leq K, \text{ holding with equality if } m \in M^{\beta};$$

- provider optimality: for all $m \in M^{\beta}$,

$$\underline{u}(m|\underline{\lambda}(m), \overline{\lambda}(m)) \leq \max_{m' \in M^{\beta}} \underline{u}(m'|\underline{\lambda}(m'), \overline{\lambda}(m')) \text{ holding with equality if } \underline{\lambda}(m) > 0,$$

$$\overline{u}(m|\underline{\lambda}(m), \overline{\lambda}(m)) \leq \max_{m' \in M^{\beta}} \overline{u}(m'|\underline{\lambda}(m'), \overline{\lambda}(m')) \text{ holding with equality if } \overline{\lambda}(m) > 0;$$

- beliefs: for all $m \notin M^{\beta}$, $\underline{\lambda}(m)$ and $\overline{\lambda}(m)$ are determined by conditions i)-iii).

A.3. Proof of Proposition 1

We start by characterizing the set of tuples $(\underline{x}, \overline{x}, \underline{t}, \overline{t}, \underline{\lambda}, \overline{\lambda})$ that identifies the set of feasible and incentive compatible mechanisms in the original space.

Lemma 7. *For any $(\underline{x}, \overline{x}, \underline{t}, \overline{t}) \in [0,1]^2 \times \mathbb{R}^2$ and $\overline{\lambda}, \underline{\lambda} \in [0, \infty)$, there exists a feasible and incentive compatible mechanism m, such that:*

$$\underline{x}_m(\underline{\lambda}, \overline{\lambda}) = \underline{x}, \quad \overline{x}_m(\underline{\lambda}, \overline{\lambda}) = \overline{x}, \quad \underline{t}_m(\underline{\lambda}, \overline{\lambda}) = \underline{t}, \quad \overline{t}_m(\underline{\lambda}, \overline{\lambda}) = \overline{t},$$

if and only if:

$$\begin{aligned}
\underline{t} - \underline{x}\,\underline{c} &\geq \overline{t} - \overline{x}\,\underline{c}, \\
\overline{t} - \overline{x}\,\overline{c} &\geq \underline{t} - \underline{x}\,\overline{c}, \\
\overline{\lambda}\,\overline{x} &\leq \phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}), \\
\underline{\lambda}\,\underline{x} &\leq \phi(\overline{\lambda}, \underline{\lambda} + \overline{\lambda}), \\
\overline{\lambda}\,\overline{x} + \underline{\lambda}\,\underline{x} &\leq \phi(\underline{\lambda} + \overline{\lambda}, \underline{\lambda} + \overline{\lambda}).
\end{aligned}$$

Proof. Only if: We first show that for any feasible and incentive compatible mechanism m, expected trading probabilities and prices satisfy conditions (6)-(10). Let $\underline{x} = \underline{x}_m(\underline{\lambda}, \overline{\lambda}), \overline{x} = \overline{x}_m(\underline{\lambda}, \overline{\lambda})$ and $\underline{t} = \underline{t}_m(\underline{\lambda}, \overline{\lambda}), \overline{t} = \overline{t}_m(\underline{\lambda}, \overline{\lambda})$. Incentive compatibility of m trivially implies (6) and (7). Feasibility will imply the remaining conditions.

[19] See Auster and Gottardi (2019) for further discussion of these conditions. Similar specifications appear also in McAfee (1993), Eeckhout and Kircher (2010) and Guerrieri et al. (2010), among others.

Rather than using the queue lengths $\underline{\lambda}, \overline{\lambda}$, it will be convenient to work with the total queue length $\lambda = \underline{\lambda} + \overline{\lambda}$ and the fraction of low types $\mu = \underline{\lambda}/\lambda$. Let $\underline{X}_m^n(\mu)$ and $\overline{X}_m^n(\mu)$ denote, respectively, the trading probability for low- and high-type providers in a meeting with $n-1$ other providers according to mechanism m. From a provider's perspective there is a probability $Q_n(\lambda)$ of arriving at a customer who meets n providers. Consistency between $P_n(\lambda)$ and $Q_n(\lambda)$ requires $nP_n(\lambda) = \lambda Q_n(\lambda)$ (Eeckhout and Kircher, 2010). We can then write:

$$\underline{x}_m(\mu\lambda, (1-\mu)\lambda) = \sum_{n=1}^{+\infty} Q_n(\lambda)\underline{X}_m^n(\mu), \quad \overline{x}_m(\mu\lambda, (1-\mu)\lambda) = \sum_{n=1}^{+\infty} Q_n(\lambda)\overline{X}_m^n(\mu).$$

The feasibility condition, $\underline{X}_m(L,H)L + \overline{X}_m(L,H)H \leq 1, \forall L, H$, implies the following two inequalities:

$$\underline{X}_m(L,H) \leq 1/L, \quad \overline{X}_m(L,H) \leq 1/H.$$

Using these conditions, we can write[20]:

$$\begin{aligned}
\underline{X}_m^n(\mu) &= \sum_{L=1}^{n} \underline{X}_m(L, n-L)\frac{(n-1)!}{(L-1)!(n-L)!}\mu^{L-1}(1-\mu)^{n-L} \\
&\leq \sum_{L=1}^{n} \frac{1}{L}\frac{(n-1)!}{(L-1)!(n-L)!}\mu^{L-1}(1-\mu)^{n-L} \\
&= \frac{1}{\mu n}\sum_{L=1}^{n} \frac{n!}{L!(n-L)!}\mu^{L}(1-\mu)^{n-L} \\
&= \frac{1}{\mu n}\left(\sum_{L=0}^{n} \frac{n!}{L!(n-L)!}\mu^{L}(1-\mu)^{n-L} - (1-\mu)^n\right) \\
&= \frac{1}{\mu n}(1-(1-\mu)^n).
\end{aligned}$$

The condition $\underline{X}_m^n(\mu) \leq \frac{1}{\mu n}(1-(1-\mu)^n)$ allows us to show that any feasible mechanisms satisfy constraint (8):

$$\begin{aligned}
\mu\lambda\underline{x}_m(\mu\lambda, \lambda) &= \mu\lambda\sum_{1}^{+\infty} Q_1(\lambda)\underline{X}_m^n(\mu) \\
&= \mu\lambda\sum_{n=1}^{+\infty} P_n(\lambda)\frac{n}{\lambda}\underline{X}_m^n(\mu) \\
&\leq \sum_{n=1}^{+\infty} P_n(\lambda)(1-(1-\mu)^n) = \phi(\mu\lambda, \lambda).
\end{aligned}$$

Analogously, we can show $\overline{X}_n(\mu) \leq \frac{1}{(1-\mu)n}(1-\mu^n)$ and, hence, $(1-\mu)\lambda\overline{x}_m(\mu\lambda,\lambda) \leq \phi((1-\mu)\lambda,\lambda)$.

Finally, for each n and μ, the expected probability of trade in a meeting cannot exceed one, that is, $\mu n\underline{X}_m^n(\mu) + (1-\mu)n\overline{X}_m^n(\mu) \leq 1$. Hence:

$$\begin{aligned}
\mu\lambda\underline{x}(\mu\lambda,\lambda) + (1-\mu)\lambda\overline{x}(\mu\lambda,\lambda) &= \mu\lambda\sum_{n=1}^{+\infty} Q_n(\lambda)\underline{X}_m^n(\mu) + (1-\mu)\lambda\sum_{n=1}^{+\infty} Q_n(\lambda)\overline{X}_m^n(\mu) \\
&= \lambda\sum_{n=1}^{+\infty} Q_n(\lambda)(\mu\underline{X}_m^n(\mu) + (1-\mu)\overline{X}_m^n(\mu)) \\
&\leq \lambda\sum_{n=1}^{+\infty} Q_n(\lambda)/n = \sum_{n=1}^{+\infty} P_n(\lambda) = \phi(\lambda,\lambda)
\end{aligned}$$

If: Let $(\mu\lambda, (1-\mu)\lambda, \underline{x}, \overline{x}, \underline{t}, \overline{t})$ be a vector satisfying (6)-(10). Let $\alpha, \underline{\tau}, \overline{\tau} \in [0,1]$ and consider the following mechanism:

$$\begin{aligned}
\underline{X}_m(L,H) &= \underline{\tau}\left(\frac{1}{L+\alpha H}\right), \quad \underline{T}_m(L,H) = \underline{t}, \quad L \geq 1, H \geq 0 \\
\overline{X}_m(L,H) &= \overline{\tau}\left(\frac{\alpha}{L+\alpha H}\right), \quad \overline{T}_m(L,H) = \overline{t}, \quad L \geq 0, H \geq 1,
\end{aligned}$$

[20] Note that the number of low and high types in a meeting is distributed binomially (n is the number of trials, L is the number of successes).

with $\overline{X}_m(0,H) = 1/H$ if $\alpha = 0$. It can be easily verified that this mechanism is feasible and incentive compatible. We start by defining the trading probabilities that obtain when $\underline{\tau} = \overline{\tau} = 1$. For the low-type provider, this probability is:

$$\underline{y}(\alpha) \equiv \sum_{n=1}^{+\infty} Q_n(\lambda) \sum_{L=1}^{n} \frac{(n-1)!}{(L-1)!(n-L)!} \mu^{L-1}(1-\mu)^{n-L} \frac{1}{\alpha n + (1-\alpha)L}.$$

Notice that $\underline{y}(\alpha)$ is strictly decreasing in α, with $\underline{y}(0) = \phi(\mu\lambda, \lambda)/(\mu\lambda)$ and $\underline{y}(1) = \phi(\lambda,\lambda)/\lambda$. Analogously, we define for the high-type provider:

$$\overline{y}(\alpha) \equiv \sum_{n=1}^{+\infty} Q_n(\lambda) \left((1-\mu)^{n-1}\frac{1}{n} + \sum_{L=1}^{n-1} \frac{(n-1)!}{L!(n-1-L)!} \mu^L (1-\mu)^{n-1-L} \frac{\alpha}{\alpha n + (1-\alpha)L} \right).$$

The function $\overline{y}(\alpha)$ is strictly increasing on $[0,1]$ with

$$\begin{aligned}
\overline{y}(0) &= \sum_{n=1}^{+\infty} Q_n(\lambda)(1-\mu)^{n-1}\frac{1}{n} \\
&= \frac{1}{(1-\mu)\lambda} \sum_{n=1}^{+\infty} P_n(\lambda)(1-\mu)^n \\
&= \frac{1}{(1-\mu)\lambda} \left(\sum_{n=0}^{+\infty} P_n(\lambda)(1-\mu)^n - P_0(\lambda) \right) \\
&= \frac{1}{(1-\mu)\lambda} (\phi(\lambda,\lambda) - \phi(\underline{\lambda},\lambda)),
\end{aligned}$$

and $\overline{y}(1) = \phi(\lambda,\lambda)/\lambda$. The two functions satisfy the following condition:

$$\begin{aligned}
&\mu\lambda \underline{y}(\alpha) + (1-\mu)\lambda \overline{y}(\alpha) \\
&= \lambda \sum_{n=1}^{+\infty} \frac{Q_n(\lambda)}{n} \left(\mu^n + (1-\mu)^n + \sum_{L=1}^{n-1} \frac{n!}{L!(n-L)!} \mu^L (1-\mu)^{n-L} \left(\frac{L}{\alpha n + (1-\alpha)L} + \frac{\alpha(n-L)}{\alpha n + (1-\alpha)L} \right) \right) \\
&= \sum_{n=1}^{\infty} P_n(\lambda) \sum_{L=0}^{n} \frac{n!}{L!(n-L)!} \mu^L (1-\mu)^{n-L} \\
&= \phi(\lambda,\lambda)
\end{aligned}$$

We now want to show that for every vector $(\mu\lambda, (1-\mu)\lambda, \underline{x}, \overline{x}, \underline{t}, \overline{t})$ satisfying (6)-(10) we can find values for $\alpha, \underline{\tau}, \overline{\tau}$ such that $\underline{x}_m(\mu\lambda,\lambda) = \underline{x}, \overline{x}_m(\mu\lambda,\lambda) = \overline{x}$ and $\underline{t}_m(\mu\lambda,\lambda) = \underline{t}, \overline{t}_m(\mu\lambda,\lambda) = \overline{t}$. For the transfers this requirement is automatically satisfied, so we just need to show that by choosing $\alpha, \underline{\tau}, \overline{\tau}$ appropriately we can generate the trading probabilities $\underline{x}, \overline{x}$. Using the functions $\underline{y}(\alpha)$ and $\overline{y}(\alpha)$, we can write the expected trading probabilities of mechanism m as follows:

$$\underline{x}_m(\mu\lambda,\lambda) = \underline{\tau}\,\underline{y}(\alpha), \quad \overline{x}_m(\mu\lambda,\lambda) = \overline{\tau}\,\overline{y}(\alpha).$$

Since we can scale down trading probabilities arbitrarily by using $\underline{\tau}, \overline{\tau} \in [0,1]$, it suffices to show that there exists an α such that

$$\underline{y}(\alpha) \geq \underline{x} \quad \text{and} \quad \overline{y}(\alpha) \geq \overline{x}. \tag{17}$$

Condition (8) implies $\underline{x} < \phi(\mu\lambda,\lambda)/(\mu\lambda)$. We distinguish two cases.

- If $\underline{x} \leq \phi(\lambda,\lambda)/\lambda$, let $\alpha = 0$ so that $\underline{y}(\alpha) = \overline{y}(\alpha) = \phi(\lambda,\lambda)/\lambda$. Since $\overline{x} \leq \underline{x}$, both conditions in (17) are satisfied.
- If $\phi(\lambda,\lambda)/\lambda < \underline{x} \leq \phi(\mu\lambda,\lambda)/(\mu\lambda)$, let α be such that $\underline{y}(\alpha) = \underline{x}$. Then:

$$\overline{x} \leq \frac{1}{(1-\mu)\lambda}(\phi(\lambda,\lambda) - \mu\lambda\underline{x}) = \frac{1}{(1-\mu)\lambda}\left(\phi(\lambda,\lambda) - \mu\lambda\underline{y}(\alpha)\right) = \overline{y}(\alpha),$$

  where the first inequality follows from (10) and the last equality follows from the property $\mu\lambda\underline{y}(\alpha) + (1-\mu)\lambda\overline{y}(\alpha) = \phi(\lambda,\lambda)$, as demonstrated above.

Hence, for every vector $(\mu\lambda, (1-\mu)\lambda, \underline{x}, \overline{x}, \underline{t}, \overline{t})$ satisfying (6)-(10) there exist values of $\alpha, \underline{\tau}, \overline{\tau}$ that generate the desired parameters. □

Conditions (4) and (5) are analogous to the providers' optimality conditions in Definition 2 for the mechanisms posted in equilibrium and conditions (15) and (16) restricting out-of-equilibrium beliefs. Given Lemma 7, it is clear that the trading probabilities, transfers and queue length associated with any mechanism posted in equilibrium must satisfy conditions (5)-(10). To see that they must also solve P^{aux}, suppose they do not and consider a tuple $(\underline{x}, \overline{x}, \underline{t}, \overline{t}, \underline{\lambda}, \overline{\lambda})$ that satisfies constraints (5)-(10) and yields a higher value of $\overline{\lambda}(\overline{x}\,\overline{v}-\overline{t})+\underline{\lambda}(\underline{x}\,\underline{v}\underline{t})$. Then there exists a feasible and incentive-compatible mechanism m, which given queue lengths $\underline{\lambda},\overline{\lambda}$ generates $(\underline{x},\overline{x},\underline{t},\overline{t})$, satisfies conditions (15) and (16), and generates a strictly higher expected payoff. Hence, a deviation to mechanism m is profitable, which implies that the trading probabilities, transfers and queue length associated with an equilibrium mechanism must solve P^{aux}.

A.4. Solving P^{aux}: preliminaries

We will start by deriving some properties of the market utilities that need to be satisfied in equilibrium.

Lemma 8. *At a directed search equilibrium, we have:*

1.) $\underline{U}>\overline{U}$ *and* $\underline{U}-\overline{U}<\overline{c}-\underline{c}$*;*
2.) $\overline{U}>0$*;*
3.) $\underline{U}<\underline{v}-\underline{c}$ *and* $\overline{U}\leq(\overline{v}-\overline{c})/(\overline{v}-\underline{c})\underline{U}$.

Proof. Let $(\underline{x},\underline{t})$ and $(\overline{x},\overline{t})$ be pairs of expected trading probabilities and transfers associated to (possibly different) mechanisms chosen by low- and high-type providers in a given equilibrium. These values must then also be part of a solution of P^{aux}. Market utilities are therefore $\underline{U}=\underline{t}-\underline{x}\,\underline{c}$ and $\overline{U}=\overline{t}-\overline{x}\,\overline{c}$. The following properties must hold:

1a. $\underline{U}>\overline{U}$: the low-type incentive constraint (6) can be rewritten as $\overline{x}(\overline{c}-\underline{c})\leq\underline{U}-\overline{U}$. Since $\overline{x}\geq 0$, this inequality can only be satisfied if $\underline{U}\geq\overline{U}$. Suppose now that $\underline{U}=\overline{U}$ so that $\overline{x}=0$. Since, under any solution of P^{aux}, customers must make weakly positive profits with both types of provider,[21] we must have $\overline{t}=0$ and hence $\underline{U}=\overline{U}=0$. Given these market utilities, a customer's payoff is strictly increasing in both queue lengths, implying that no finite values of $\underline{\lambda},\overline{\lambda}$ can solve P^{aux}. Hence, $\underline{U}=\overline{U}$ are not be admissible equilibrium values.

1b. $\underline{U}-\overline{U}<\overline{c}-\underline{c}$: the high type incentive constraint (7) requires $\overline{t}-\overline{x}\,\overline{c}\geq\underline{t}-\underline{x}\overline{c}$, or $\underline{x}(\overline{c}-\underline{c})\geq\underline{U}-\overline{U}$. Since $\underline{x}\leq\phi(\underline{\lambda},\underline{\lambda}+\overline{\lambda})/\underline{\lambda}<1$, the inequality $\underline{x}(\overline{c}-\underline{c})\geq\underline{U}-\overline{U}$ can only be satisfied if $\underline{U}-\overline{U}<\overline{c}-\underline{c}$.

2. $\overline{U}>0$: Let $(\underline{\lambda},\overline{\lambda},\underline{x},\overline{x},\underline{t},\overline{t})$ describe a mechanism and associated queue lengths posted in equilibrium with $\overline{\lambda}\overline{x}>0$. To see that such a mechanism exists in all equilibria, consider a candidate equilibrium where this is not satisfied and, hence, high types do not trade. A customer could then deviate and offer a mechanism attracting only high types with terms of trade described by $(\overline{x},\overline{t})$ such that $\overline{t}/\overline{x}=c_H$ and $\overline{x}\leq\underline{U}/(\overline{c}-\underline{c})$. The first equality says that the mechanism yields zero payoff for high types and the second inequality says that the low-type incentive constraint (6) is satisfied. To satisfy the overall feasibility constraint (10), set $\underline{\lambda}=0$ and $\overline{\lambda}$ such that $\phi(\overline{\lambda},\overline{\lambda})/\overline{\lambda}=\overline{x}$. Notice that the value of $\overline{\lambda}$ solving this equation strictly decreases in $\overline{x}$ and tends to $+\infty$ as $\overline{x}\to 0$. The customer's payoff associated to this mechanism is given by $\overline{\lambda}(\overline{x}\overline{v}-\overline{t})=\phi(\overline{\lambda},\overline{\lambda})(\overline{v}-\overline{c})$. Since $k<\overline{v}-\overline{c}$, we can then choose $\overline{x}$ sufficiently small and hence $\overline{\lambda}$ sufficiently large such that $\phi(\overline{\lambda},\overline{\lambda})(\overline{v}-\overline{c})>k$ and the deviation is strictly possible.
Having shown that there is an equilibrium mechanism, described by $(\underline{\lambda},\overline{\lambda},\underline{x},\overline{x},\underline{t},\overline{t})$ with $\overline{\lambda}\overline{x}>0$, we want to prove $\overline{U}>0$. Towards a contradiction, assume $\overline{U}=0$ and consider an alternative mechanism $(\underline{\lambda},\overline{\lambda}',\underline{x},\overline{x}',\underline{t},\overline{t})$ with $\overline{\lambda}'>\overline{\lambda}$ and $\overline{x}'$ such that $\overline{\lambda}'\overline{x}'=\overline{\lambda}\overline{x}$ and $\overline{t}'=\overline{x}'\overline{c}$. Due to $\overline{\lambda}'\overline{x}'=\overline{\lambda}\overline{x}$, the customer's payoff associated to this mechanism, $\underline{\lambda}\underline{x}(\underline{v}-\underline{c})-\underline{\lambda}\underline{U}+\overline{\lambda}'\overline{x}'(\overline{v}-\overline{c})$, is the same as under the original mechanism. Moreover, since $\overline{x}'<\overline{x}$, this mechanism satisfies the low-type incentive compatibility constraint (6) and since

$$\underline{\lambda}\,\underline{x}+\overline{\lambda}'\overline{x}'=\underline{\lambda}\,\underline{x}+\overline{\lambda}\overline{x}\leq\phi(\underline{\lambda}+\overline{\lambda},\underline{\lambda}+\overline{\lambda})<\phi(\underline{\lambda}+\overline{\lambda}',\underline{\lambda}+\overline{\lambda}')$$

it satisfies the overall feasibility constraint (10) with strict inequality. An increase of $\overline{\lambda}'$ is then feasible for the customer, and it strictly increases her expected payoff. The customer thus has a profitable deviation when $\overline{U}=0$.

3a. $\underline{U}<\underline{v}-\underline{c}$: Suppose not, $\underline{U}\geq\underline{v}-\underline{c}$. Since, as shown in 1b. above, $\underline{x}<1$ whenever $\underline{\lambda}>0$, this implies that $\underline{x}(\underline{v}-\underline{c})-\underline{U}<0$; that is, a customer's payoff with each low-type provider is strictly negative. As a consequence, at any solution of P^{aux}, we have $\underline{\lambda}=0$. This in turn implies that the low types' market utility $\underline{U}$ must equal zero and therefore $\underline{U}<\underline{v}-\underline{c}$. A contradiction.

3b. $\overline{U}\leq(\overline{v}-\overline{c})/(\overline{v}-\underline{c})\underline{U}$: Suppose not, $\overline{U}>(\overline{v}-\overline{c})/(\overline{v}-\underline{c})\underline{U}$. As argued above, the low-type incentive constraint (6) can be written as $\overline{x}\leq(\underline{U}-\overline{U})/(\overline{c}-\underline{c})$. This implies that the payoff of a customer with each high-type provider is negative:

$$\overline{x}(\overline{v}-\overline{c})-\overline{U}\leq\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\overline{c})-\overline{U}=\frac{(\overline{v}-\overline{c})\underline{U}-(\overline{v}-\underline{c})\overline{U}}{\overline{c}-\underline{c}}<0.$$

All solutions of P^{aux} must therefore satisfy $\overline{\lambda}=0$, which in turn implies $\overline{U}=0$ and therefore $\overline{U}<(\overline{v}-\overline{c})/(\overline{v}-\underline{c})\underline{U}$. A contradiction. □

[21] If customers make losses with one type of provider, they can always set the respective queue length, $\underline{\lambda}$ or $\overline{\lambda}$, equal to zero.

A.5. Proof of Proposition 2

To prove the statement of the proposition, we start by establishing three properties of solutions of P^{aux}.

Lemma 9. *Let Assumption 1 be satisfied. At a solution of P^{aux}, the overall feasibility constraint (10) holds as equality.*

Proof. Suppose not and let $(\underline{\lambda},\overline{\lambda},\underline{x},\overline{x})$ be a solution of P^{aux}. Observe first that (10) can only be satisfied with strict inequality if $\overline{\lambda}>0$. Given $\overline{\lambda}>0$, we must have $\overline{x}=\frac{\underline{U}-\overline{U}}{\underline{c}-\overline{c}}$, as otherwise an increase in $\overline{x}$ would be feasible, incentive compatible and profitable. Substituting for $\overline{x}$ and $\underline{\lambda x}$ (the latter determined by (8) holding as equality), the customer's payoff can be written as

$$\hat{\pi}(\underline{\lambda},\overline{\lambda})=\phi(\underline{\lambda},\underline{\lambda}+\overline{\lambda})(\underline{v}-\underline{c})+\overline{\lambda}\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\overline{c})-\underline{\lambda}\underline{U}-\overline{\lambda}\overline{U},$$

Since the overall feasibility constraint is not binding at a solution of P^{aux}, the pair $(\underline{\lambda},\overline{\lambda})$ must be a maximizer of the function $\hat{\pi}$. Notice however that

$$\frac{\partial^2\hat{\pi}}{\partial\overline{\lambda}^2}(\underline{\lambda},\overline{\lambda})=\frac{\partial^2\phi}{\partial\overline{\lambda}^2}(\underline{\lambda},\underline{\lambda}+\overline{\lambda})>0.$$

The function $\hat{\pi}(\underline{\lambda},\overline{\lambda})$ does not have an interior maximimum. This implies that at a solution of $P^{aux\prime}$ the overall feasibility constraint (10) is satisfied with equality. □

Lemma 10 *(Auster and Gottardi, 2019). At any solution of P^{aux}, the low-type feasibility constraint (8) is satisfied with equality.*

Proof. See the proof of Lemma 3.3 in Auster and Gottardi (2019) and replace $1-e^{-\underline{\lambda}}$ with $\phi(\underline{\lambda},\underline{\lambda}+\overline{\lambda})$ and $1-e^{-\underline{\lambda}-\overline{\lambda}}$ with $\phi(\underline{\lambda}+\overline{\lambda},\underline{\lambda}+\overline{\lambda})$. □

Lemma 11. *At a solution of P^{aux} with $\overline{\lambda}>0$, the low-type incentive constraint (6) is satisfied as an equality.*

Proof. To prove the claim we proceed in two steps:

1. Given Lemma 9 and Lemma 10, we can write the customer's payoff as in (11). We start by proving that if the incentive constraint (6) were slack, customers would not find it optimal to attract both types of providers. In particular, we show that the objective function (11) has its unique maximum at $\underline{\lambda}=0$ and $\overline{\lambda}$ such that

$$\frac{\partial\phi}{\partial\overline{\lambda}}(\overline{\lambda},\overline{\lambda})(\overline{v}-\overline{c})=\overline{U}. \tag{18}$$

   To prove this claim, suppose, towards a contradiction, there is a local maximum with $\underline{\lambda}>0$ and consider the pair $(\underline{\lambda}-\varepsilon,\overline{\lambda}+\varepsilon)$ with $\varepsilon\in(0,\underline{\lambda})$. We then have

$$\begin{aligned}&\phi(\underline{\lambda},\underline{\lambda}+\overline{\lambda})[(\underline{v}-\underline{c})-(\overline{v}-\overline{c})]+\phi(\underline{\lambda}+\overline{\lambda},\underline{\lambda}+\overline{\lambda})(\overline{v}-\overline{c})-\underline{\lambda}\underline{U}-\overline{\lambda}\overline{U},\\ <\ &\phi(\underline{\lambda}-\underline{\lambda},\underline{\lambda}+\overline{\lambda})[(\underline{v}-\underline{c})-(\overline{v}-\overline{c})]+\phi(\underline{\lambda}+\overline{\lambda},\underline{\lambda}+\overline{\lambda})(\overline{v}-\overline{c})-\underline{\lambda}\underline{U}-\overline{\lambda}\overline{U}-\varepsilon(\underline{U}-\overline{U})\end{aligned}$$

   where the inequality follows from the facts that (i) ϕ is decreasing in the first argument, (ii) $\overline{v}-\overline{c}\geq\underline{v}-\underline{c}$ and (iii) $\overline{U}<\underline{U}$. Since the expression of the second line corresponds to the value of the objective function in (11) at $(\underline{\lambda}-\varepsilon,\overline{\lambda}+\varepsilon)$, the pair $(\underline{\lambda},\overline{\lambda})$ cannot be a local maximizer: by replacing some low types with high types the customer can strictly increase her payoff. Moreover, the maximum of the function in (11), if it exists, is attained at a point with $\underline{\lambda}=0$. It can be found by solving the one-variable optimization problem

$$\max_{\overline{\lambda}}\phi(\overline{\lambda},\overline{\lambda})\left(\overline{v}-\overline{c}\right)-\overline{\lambda}\overline{U}. \tag{19}$$

   Since $\phi(\overline{\lambda},\overline{\lambda})$ is concave in $\overline{\lambda}$ this problem has a unique solution, characterized by the first-order condition (18).
2. We are then left to consider the possibility of a separating equilibrium, where some customers attract low types and others attract high types. If incentive constraints are slack, customers attracting low-types make a profit equal to

$$\max_{\underline{\lambda}}\phi(\underline{\lambda},\underline{\lambda})(\underline{v}-\underline{c})-\underline{\lambda}\,\underline{U}$$

   while those attracting high types make a profit equal to (19). Since $\underline{U}>\overline{U}$ and $\underline{v}-\underline{c}\leq\overline{v}-\overline{c}$, the value of the latter is strictly greater than the value of the former. This means that customers cannot be indifferent between both markets, which then precludes such an equilibrium. □

Proposition 12. *Let Assumption 1 be satisfied. A directed search equilibrium exists.*

Proof. Consider the following modified (truncated) customer's auxiliary problem as follows. The customer chooses $\underline{\lambda}, \overline{\lambda}$ so as to maximize

$$\phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})(\underline{v}-\underline{c}) + [\phi(\underline{\lambda}+\overline{\lambda}, \underline{\lambda}+\overline{\lambda}) - \phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})](\overline{v}-\overline{c}) - \underline{\lambda U} - \overline{\lambda U}, \tag{$P^{aux,T}$}$$

subject to

$$[\phi(\underline{\lambda}+\overline{\lambda}, \underline{\lambda}+\overline{\lambda}) - \phi(\underline{\lambda}, \underline{\lambda}+\overline{\lambda})](\overline{c}-\underline{c}) \le \overline{\lambda}(\underline{U}-\overline{U}),$$
$$\Lambda \ge \underline{\lambda} \ge 0,$$
$$\Lambda \ge \overline{\lambda} \ge 0$$

for a given pair of market utilities, $\underline{U}, \overline{U}$, and a parameter Λ set at an arbitrarily large positive level. The problem differs from the auxiliary problem (as in the paper) only for the upper bound imposed on $\underline{\lambda}, \overline{\lambda}$.

Consider the following domain for market utilities:

$$\mathcal{U} = \left\{ \underline{U}, \overline{U} : \overline{v} - \overline{c} \ge \underline{U} \ge \overline{U} \ge 0 \right\},$$

a convex, compact set.

It is easy to verify that a solution of the above modified auxiliary problem exists for all $\underline{U}, \overline{U} \in \mathcal{U}$. Moreover, the solution is described by a map $\left(\underline{\lambda}, \overline{\lambda}\right) : \mathcal{U} \to [0, \Lambda]^2$, which is u.h.c. by the Maximum Theorem. The value of the customer's profits at a solution is then described by the map $\pi : \left(\underline{\lambda}, \overline{\lambda}\right) : \mathcal{U} \to [0, \overline{v}-\overline{c}]$, a continuous function.

Consider then the convex hull of this map: $co\left(\underline{\lambda}, \overline{\lambda}\right)(\underline{U}, \overline{U})$. This is u.h.c. and convex valued.

Next, we define another map $\left(\underline{U}, \overline{U}\right) : [0, \Lambda]^2 \times [0, \overline{v}-\overline{c}] \to \mathcal{U}$, that associates to any value of $\underline{\lambda}, \overline{\lambda}, \pi$ the solution of the following optimization problem:

$$\max_{\underline{U}, \overline{U}} \left(\underline{U} - \overline{U}\right)\left(\underline{\lambda} - \mu\left(\underline{\lambda} + \overline{\lambda}\right)\right) + (\underline{U} + \overline{U})(\pi - K)$$

It is immediate to verify that the solution of the above problem always exists, hence the map $\left(\underline{U}, \overline{U}\right)(\underline{\lambda}, \overline{\lambda}, \pi)$ is non empty. Furthermore, it is u.h.c. and convex valued.

Finally, consider the composite map:

$$\left(\underline{U}, \overline{U}\right) : \mathcal{U} \to \mathcal{U}$$

defined by $\left(\underline{U}, \overline{U}\right)(\underline{\lambda}, \overline{\lambda}, \pi)$ for all $\underline{\lambda}, \overline{\lambda} \in co(\underline{\lambda}, \overline{\lambda})\left(\underline{U}, \overline{U}\right)$, $\pi \in \pi\left(\underline{U}, \overline{U}\right)$. This map is non empty, u.h.c. and convex-valued, since these properties are satisfied by all composing functions, and its domain $\mathcal{U}$ is a compact, convex set. Hence, by Kakutani's fixed point theorem, a fixed point exists:

$$\underline{U}^*, \overline{U}^* \in \arg\max \left(\underline{U} - \overline{U}\right)\left(\underline{\lambda} - \mu\left(\underline{\lambda} + \overline{\lambda}\right)\right) + (\underline{U} + \overline{U})(\pi - K)$$
$$\text{for } \underline{\lambda}, \overline{\lambda} \in co(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right), \pi \in \pi\left(\underline{U}^*, \overline{U}^*\right)$$

It remains to establish that the fixed point pair of market utilities, $\underline{U}^*, \overline{U}^*$, and the associated values of the queue lengths $\underline{\lambda}^*, \overline{\lambda}^* \in co(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right)$, such that

$$\underline{U}^*, \overline{U}^* \in \arg\max \left(\underline{U} - \overline{U}\right)\left(\underline{\lambda}^* - \mu\left(\underline{\lambda}^* + \overline{\lambda}^*\right)\right) + (\underline{U} + \overline{U})\left(\pi\left(\underline{U}^*, \overline{U}^*\right) - K\right), \tag{20}$$

constitute a directed search equilibrium, that is:

$$\underline{\lambda}^* = \mu\left(\underline{\lambda}^* + \overline{\lambda}^*\right) \text{ and } \underline{\lambda}^* + \overline{\lambda}^* > 0 \tag{21}$$
$$\pi\left(\underline{U}^*, \overline{U}^*\right) = K$$

and the upper bound imposed on λ does not bind.

We prove the claim by contradiction, by considering in turn various ways in which the two above are violated:

1. $\underline{\lambda}^* = \overline{\lambda}^* = 0$ and hence $\pi\left(\underline{U}^*, \overline{U}^*\right) = 0$. At these values, the solution of (20) is given by the minimum value of $\underline{U} + \overline{U}$, that is $\underline{U}^* = \overline{U}^* = 0$, but then $(\underline{\lambda}, \overline{\lambda})(0,0) > 0$, a contradiction. Hence we must have $\underline{\lambda}^* + \overline{\lambda}^* > 0$.
2. $\underline{\lambda}^* \geq \mu\left(\underline{\lambda}^* + \overline{\lambda}^*\right)$, $\pi\left(\underline{U}^*, \overline{U}^*\right) \geq K$, with at least one of the two inequalities holding strict. At these values, the solution of (20) is a corner solution, given either by $\underline{U}^* = \overline{v} - \overline{c}, \overline{U}^* = 0$, in which case $(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right) = (0, \Lambda)$, a contradiction, or by $\underline{U}^* = \overline{U}^* = \overline{v} - \overline{c}$ in which case $(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right) = (0,0)$, also a contradiction.
3. $\underline{\lambda}^* \leq \mu\left(\underline{\lambda}^* + \overline{\lambda}^*\right)$, $\pi\left(\underline{U}^*, \overline{U}^*\right) \leq K$, with at least one of the two inequalities holding strict. At these values, the solution of (20) is a corner solution, given by $\underline{U}^* = \overline{U}^* = 0$, in which case $(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right) = (\Lambda, 0)$, a contradiction.
4. $\underline{\lambda}^* > \mu\left(\underline{\lambda}^* + \overline{\lambda}^*\right)$, $\pi\left(\underline{U}^*, \overline{U}^*\right) < K$. At these values, the solution of (20) is a corner solution, given either by $\underline{U}^* = \overline{v} - \overline{c}, \overline{U}^* = 0$, in which case $(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right) = (0, \Lambda)$, a contradiction, or by $\underline{U}^* = \overline{U}^* = 0$, in which case $(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right) = (\Lambda, 0)$, and $\pi\left(\underline{U}^*, \overline{U}^*\right) \simeq \underline{v} - \underline{c} > K$, a contradiction.
5. $\underline{\lambda}^* < \mu\left(\underline{\lambda}^* + \overline{\lambda}^*\right)$, $\pi\left(\underline{U}^*, \overline{U}^*\right) > K$. At these values, the solution of (20) is a corner solution, given by $\underline{U}^* = \overline{U}^* = \overline{v} - \overline{c}$ in which case $(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right) = (0,0)$, $\pi\left(\underline{U}^*, \overline{U}^*\right) = 0 < K$, also a contradiction.

The above argument shows that at a fixed point, the equilibrium conditions (21) hold. The last property to be shown is that $\underline{\lambda}^* < \Lambda, \overline{\lambda}^* < \Lambda$—that is, the constraint we imposed does not bind—for Λ sufficiently high, and hence $\underline{\lambda}^*, \overline{\lambda}^*$ are also a solution of the original auxiliary problem, given $\underline{U}^*, \overline{U}^*$, not only of the modified/truncated one.

Notice first that we must have $\underline{U}^* > 0, \overline{U}^* > 0$. If $\underline{U}^* = \overline{U}^* = 0$, as argued above we have $(\underline{\lambda}, \overline{\lambda})\left(\underline{U}^*, \overline{U}^*\right) = (\Lambda, 0)$, in which case (21) cannot hold, thus a contradiction. Suppose next $\underline{U}^* > \overline{U}^* = 0$. In this case $\underline{\lambda} = 0, \overline{\lambda} = \Lambda$ is a feasible choice for the modified auxiliary problem $P^{aux,T}$ and the value of the objective function, when $\underline{\lambda} = 0, \overline{\lambda} = \Lambda$ is larger than K, since it would equal $\overline{v} - \overline{c}$ (a customer would trade with a high-type provider with probability one and extract all the surplus), which is greater than K, thus contradicting the second equilibrium property in (21).

Having established that market utilities are strictly positive, $\underline{U}^* > 0, \overline{U}^* > 0$, observe that:

$$\phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}) + [\phi(\underline{\lambda} + \overline{\lambda}, \underline{\lambda} + \overline{\lambda}) - \phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda})] \leq 1$$

Hence the positive component of the payoff of a customer is bounded above by $\overline{v} - \overline{c}$. If the constraint on the queue length binds, the negative component of the customer's payoff is bounded below by $\Lambda\overline{U}^*$. Hence $\pi < \overline{v} - \overline{c} - \Lambda\overline{U}^*$. If the term on the r.h.s. were smaller than K we would have a contradiction. I do not think we can claim this property holds for any given Λ, no matter how large it is. However, we can consider a sequence of values of Λ, going to infinity, and examine the associated sequence of fixed points. We should be able to argue, on the basis of a similar argument as above that the sequence of values of $\underline{U}^*, \overline{U}^*$ at these fixed points, must be bounded away from 0 (if not, by the same argument as in the previous paragraph, as we approach the limit we would violate at least one of the two equilibrium properties). Given this, as we approach the limit, the value of the queue lengths must be strictly smaller than Λ, or we would violate the second property of (21). □

A.6. Proof of Lemma 3

Setting $\lambda = \underline{\lambda} + \overline{\lambda}$, for $\lambda > \underline{\lambda}$ the low-type incentive constraint (12) can be written as

$$\frac{1 + (1-\sigma)\lambda}{(1+\lambda)(1 + \sigma\underline{\lambda} + (1-\sigma)\lambda)} = \frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}} \tag{22}$$

For $\sigma > 0$, we can solve this condition for $\underline{\lambda}$ and write the low-type queue length as a function of the total queue length

$$\underline{\lambda}(\lambda) = \frac{(1 + (1-\sigma)\lambda)\left(1 - \frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}}(1+\lambda)\right)}{\sigma(1+\lambda)\frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}}} \tag{23}$$

The admissibility requirements $\underline{\lambda} \geq 0$ and $\underline{\lambda} \leq \lambda$ impose restrictions on the total queue length λ. When $\underline{\lambda} = 0$, the λ is such that

$$\frac{1}{1+\lambda} = \frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}}. \tag{24}$$

Let us call the solution of this equality $\overline{l}$. When $\underline{\lambda} = \lambda$, then λ is determined by

$$\frac{1}{1+\lambda}\left(1-\frac{\sigma\lambda}{1+\lambda}\right)=\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}. \tag{25}$$

Calling the solution of this equality $\underline{l}$, we thus require $\lambda \in [\underline{l}, \overline{l}]$.

Next, we can use (23) to substitute for $\underline{\lambda}$ in the customers' expected payoff (11) and obtain

$$\frac{1}{\sigma}\left(1-\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(1+\lambda)\right)(\underline{v}-\underline{c})+\lambda\left(\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\overline{c})-\overline{U}\right)-\underline{\lambda}(\lambda)\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\underline{c}) \tag{26}$$

with $\underline{\lambda}(\lambda)$ specified by (23). Notice that only the last term in (26) is non-linear in λ. Differentiation of $\underline{\lambda}(\lambda)$ yields

$$\underline{\lambda}'(\lambda)=-\frac{(1-\sigma)(1+\lambda)^2\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}+\sigma}{\sigma(1+\lambda)^2\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}}<0$$

and

$$\underline{\lambda}''(\overline{\lambda})=\frac{2}{(1+\lambda)\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}}>0$$

The last inequality implies concavity of the customers' expected payoff as written in (26) with respect to λ.

For $\sigma=0$ the left-hand side of (22) does not depend on $\underline{\lambda}$, as the trading probability in this market only depends on the overall queue length λ. It follows that there is at most one value of $\lambda>\underline{\lambda}$ at which (22) holds. Assuming that such a value of λ exists, the only free parameter is $\underline{\lambda}$. Since for a fixed λ, a customer's expected payoff

$$\phi(\underline{\lambda},\lambda)(\underline{v}-\underline{c})+(\phi(\lambda,\lambda)-\phi(\underline{\lambda},\lambda))(\overline{v}-\overline{c})-\underline{\lambda}\,\underline{U}-(\lambda-\underline{\lambda})\overline{U}$$

strictly decreases in $\underline{\lambda}$, the optimal value of $\underline{\lambda}$ is zero. Hence, on the domain $\lambda>\underline{\lambda}$, or equivalently $\overline{\lambda}>0$, the solution of P^{aux} is unique. □

A.7. Proof of Proposition 4

Pure sorting equilibrium Consider first the possibility of an equilibrium where high- and low-type providers search in two different submarkets. Let us start by describing the term of trade in the low-quality market. Since the high-type incentive constraint is not binding, the queue length in the low-quality market, $\underline{\lambda}_1$, is determined by the first-order condition of P^{aux} with respect to $\underline{\lambda}$ under the restriction $\overline{\lambda}=0$:

$$\underline{U}=\left(\frac{1}{1+\underline{\lambda}_1}\right)^2(\underline{v}-\underline{c}). \tag{27}$$

Substituting the market utility back into the customer's payoff, the free-entry condition pins down the queue length in the low-quality market:

$$\left(\frac{\underline{\lambda}_1}{1+\underline{\lambda}_1}\right)^2(\underline{v}-\underline{c})=K \quad\Rightarrow\quad \underline{U}=\left(\sqrt{\underline{v}-\underline{c}}-\sqrt{K}\right)^2 \tag{28}$$

Given $\underline{U}$, the queue length of high types $\overline{\lambda}_2$ and market utility $\overline{U}$ are determined by the low-type incentive constraint

$$\frac{1}{1+\overline{\lambda}_2}=\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}$$

and the customers' free-entry condition

$$\frac{\overline{\lambda}_2}{1+\overline{\lambda}_2}(\overline{v}-\overline{c})-\overline{\lambda}_2\overline{U}=K$$

The separating equilibrium exists if a customer cannot gain by posting a mechanism that attracts a mix of low- and high-type providers. In the proof of Proposition 3, Section A.6, we derived the bounds for the total queue length λ for which there exists a low-type queue length $\underline{\lambda}$ such that a priority mechanism satisfies the low-type incentive constraint with equality. We showed that the customer's payoff from such mechanism, (26), is differentiable and strictly concave in λ for all λ strictly above the lower bound for λ. At the lower bound, determined by (25), we have $\underline{\lambda}=\lambda$, which means that the associated mechanism attracts no high-type providers and the low-type incentive constraint imposes no restrictions. The customer's payoff as a function of λ has a discontinuity at this point, because the choice of the queue length of low-type providers is unconstrained. Since any optimal mechanism attracting a positive queue of high type providers—on and off path—satisfies the low-type incentive constraint with equality, the payoff function (26) provides an upper bound for the customer's payoff associated with mechanisms that attract both types of the provider.

Given the concavity of (26) and the continuity of (26) at the upper bound for λ where $\underline{\lambda} = 0$, attracting both types does not constitute a profitable deviation if the first derivative of (26) at the upper bound for λ, given by $\left(1 - \frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}\right) / \frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}$, determined by (24), is weakly positive. That is:

$$\frac{1}{\sigma}\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}\left((1-\sigma)(\overline{v}-\underline{c})-(\underline{v}-\underline{c})\right)+\left(\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\overline{c})-\overline{U}\right)+\left(\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}\right)^2(\overline{v}-\underline{c})\geq 0. \tag{29}$$

Notice that the first term on the left-hand side of (29) is strictly decreasing in σ. Moreover, this term is positive for σ small and approaches infinity as $\sigma \to 0$. Hence, the above condition is satisfied for σ sufficiently close to 0; equivalently, there is a threshold $\sigma^S > 0$ such that a separating equilibrium exists if and only if $\sigma \leq \sigma^S$.

We also want to prove that the threshold σ^S is strictly smaller than one. To establish this property, we show in what follows that for σ sufficiently close to 1 a profitable deviation exists at the upper bound for λ. To this end, it is useful to establish that market utilities in the candidate separating equilibrium satisfy the inequality

$$\underline{U}(\overline{v}-\overline{c}) > \overline{U}(\overline{v}-\underline{c}). \tag{30}$$

To show this, note first that, given the binding low-type incentive constraint (6), the customers' expected profits in the high-quality market can be written as

$$\overline{\lambda}\left(\underbrace{\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}}_{=\overline{x}}(\overline{v}-\overline{c})-\overline{U}\right).$$

Using the property that in the candidate separating equilibrium these profits are equal to $K > 0$, the following inequality holds:

$$\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\overline{c})-\overline{U} > 0$$

A simple manipulation of this inequality yields condition (30).

We then need to show that at $\sigma = 1$ inequality (29) is always violated. As argued above, given the concavity of (26), condition (29) is equivalent to the requirement that customers cannot profit by attracting a mix of high- and low-type providers.

Starting from a candidate separating equilibrium, consider a mechanism that attracts the same low-type queue length as in the low-quality market, $\underline{\lambda}_1$, and a queue $\overline{\lambda} > 0$ of high-type providers.[22] Since under $\sigma = 1$ there is no crowding out in the meeting process, this mechanism generates the same payoff with low-type providers as the candidate equilibrium mechanism attracting only low-type providers. Under the alternative mechanism, the trading probability for high-type providers is:

$$\overline{x} = \frac{1}{\overline{\lambda}}\left(\frac{\underline{\lambda}_1+\overline{\lambda}}{1+\underline{\lambda}_1+\overline{\lambda}} - \frac{\underline{\lambda}_1}{1+\underline{\lambda}_1}\right) = \frac{1}{(1+\underline{\lambda}_1+\overline{\lambda})(1+\underline{\lambda}_1)}. \tag{31}$$

Notice that for $\overline{\lambda} > 0$, we have $\overline{x} < 1/(1+\underline{\lambda}_1)^2 = \underline{U}/(\underline{v}-\underline{c})$ (see first-order condition (27)). The payoff the alternative mechanism generates with high-type providers is given by:

$$\overline{\lambda}\left(\overline{x}(\overline{v}-\overline{c})-\overline{U}\right) = \overline{\lambda}\left(\frac{1}{(1+\underline{\lambda}_1+\overline{\lambda})(1+\underline{\lambda}_1)}(\overline{v}-\overline{c})-\overline{U}\right) \tag{32}$$

$$= \overline{\lambda}\left(\frac{1+\underline{\lambda}_1}{1+\underline{\lambda}_1+\overline{\lambda}}\frac{\overline{v}-\overline{c}}{\underline{v}-\underline{c}}\underline{U}-\overline{U}\right), \tag{33}$$

where the second equality follows from (27). Since $\underline{U} > \overline{U}$ and $\overline{v}-\overline{c} \geq \underline{v}-\underline{c}$, the term in the bracket is strictly positive for values of $\overline{\lambda}$ belonging to a right neighborhood of zero. If $\underline{U}/(\underline{v}-\underline{c}) \leq (\underline{U}-\overline{U})/(\overline{c}-\underline{c})$, then the low-type incentive constraint (6) is satisfied for values of $\overline{\lambda}$ sufficiently close to zero (recall $\underline{x} < \underline{U}/(\underline{v}-\underline{c})$). Hence, in this case, for small values of $\overline{\lambda}$ the alternative mechanism is incentive compatible and constitutes a profitable deviation. Consider next the case $\underline{U}/(\underline{v}-\underline{c}) > (\underline{U}-\overline{U})/(\overline{c}-\underline{c})$. Since

$$\frac{1}{(1+\underline{\lambda}_1)^2} = \frac{\underline{U}}{\underline{v}-\underline{c}}, \tag{34}$$

given $\underline{U}/(\underline{v}-\underline{c}) > (\underline{U}-\overline{U})/(\overline{c}-\underline{c})$, we can pick $\overline{\lambda} > 0$ so as to satisfy

[22] This is a different deviation than the one considered in condition (29), but all we need to show is that there is a profitable deviation to *some* mechanism attracting both types. If this is true, then concavity of (26) implies that condition (29) is violated.

$$\frac{1}{(1+\underline{\lambda}_1)(1+\underline{\lambda}_1+\overline{\lambda})}=\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}. \tag{35}$$

Given this specification of $\overline{\lambda}$, the alternative mechanism satisfies the low-type incentive compatibility constraint and, by (30), generates a strictly positive profit with high-type providers, therefore constituting a profitable deviation when $\sigma=1$. By continuity, the same is true for σ sufficiently close to 1. Having established a profitable deviation exists at $\overline{\lambda}=0$, this point is not a maximum of (26); given the concavity of this function it follows its derivative at $\overline{\lambda}=0$ must be negative.

Partial sorting When $\sigma>\sigma^S$, no pure sorting equilibrium exists. Due to Proposition 12 and Lemma 3, we then know that either a partial sorting or pure screening equilibrium must exist.

In a partial sorting equilibrium, several conditions must be satisfied. First, conditions (27) and (28) pin down the queue length $\underline{\lambda}_1$ and market utility $\underline{U}$ in market 1 (recall $\overline{\lambda}_1=0$). Next, the customer's first-order condition of (26), pinning down the optimal total queue length $\lambda_2=\underline{\lambda}_2+\overline{\lambda}_2$ must hold:

$$-\frac{1}{\sigma}\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\underline{v}-\underline{c})+\left(\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\overline{c})-\overline{U}\right)-\underline{\lambda}'(\lambda_2)\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\underline{c})=0$$

with $\underline{\lambda}(\lambda_2)$ specified by (23). A simple manipulation of the above equality yields

$$\frac{1}{\sigma}\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\underline{v})-\underline{U}+\left(\frac{1}{1+\lambda_2}\right)^2(\overline{v}-\underline{c})=0. \tag{36}$$

The queue of low types $\underline{\lambda}_2=\underline{\lambda}(\lambda_2)$ is determined by the low-type incentive constraint

$$\frac{1}{\lambda_2-\underline{\lambda}_2}\left(\frac{\lambda_2}{1+\lambda_2}-\frac{\underline{\lambda}_2}{1+\sigma\underline{\lambda}_2+(1-\sigma)\lambda_2}\right)=\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}. \tag{37}$$

Finally, the customers' free-entry condition must hold

$$\frac{\underline{\lambda}_2}{1+\sigma\underline{\lambda}_2+(1-\sigma)\lambda_2}(\underline{v}-\underline{c})+\left(\frac{\lambda_2}{1+\lambda_2}-\frac{\underline{\lambda}_2}{1+\sigma\underline{\lambda}_2+(1-\sigma)\lambda_2}\right)(\overline{v}-\overline{c})-\underline{\lambda}_2\underline{U}-(\lambda-\underline{\lambda}_2)\overline{U}=K. \tag{38}$$

Whenever there is a solution $(\underline{\lambda}_2,\lambda_2,\overline{U})$ of conditions (36)-(38) with $\underline{\lambda}_2/\lambda_2<\mu$, we can find two measures β_1,β_2 of customers entering markets 1 and 2 and a measure of low types $l_1<\mu$ entering market 1 such that $l_1/\beta_1=\underline{\lambda}_1$, $(1-l_1)/\beta_2=\underline{\lambda}_2$, $(1-l_1+1-\mu)/\beta_2=\lambda_2$. This establishes that if for some σ, a partial sorting equilibrium exists for some μ, it also exists for higher values of μ. In the last part of the proof, we show that, assuming $\sigma>\sigma^S$, the equilibrium exists if μ is sufficiently close to one.

Pure screening Assume $\sigma>\sigma^S$ and consider the possibility of a pure screening equilibrium ($\underline{\lambda}_1=\overline{\lambda}_1=0$). Consistency with the population parameters requires that the fraction of low types in market 2 is μ, hence $\underline{\lambda}_2=\mu\lambda_2$. The remaining parameters of the candidate pure screening equilibrium, $\lambda_2,\underline{U},\overline{U}$, are then determined by the customers' first-order condition (36), by the low-type incentive constraint (37) and by the free-entry condition (38), as for the partial sorting equilibrium.

A pure screening equilibrium exists if customers have no incentives to deviate to the best mechanism that only attracts low types. Hence, K must be weakly larger than the payoff a customer can obtain when attracting only low types:

$$K\geq\max_{\lambda}\left(\frac{\lambda}{1+\lambda}(\underline{v}-\underline{c})-\lambda\underline{U}\right), \tag{39}$$

or equivalently

$$K\geq\left(\sqrt{\underline{v}-\underline{c}}-\sqrt{\underline{U}}\right)^2.$$

Notice that as $\mu\to1$, a customer's payoff in market 2 (the only active market in the pure screening equilibrium) converges to $\lambda_2/(1+\lambda_2)(\underline{v}-\underline{c})-\lambda_2\underline{U}$. For (39) not to be violated, λ_2 must converge to the maximizer of the right-hand side of (39). Hence, as $\mu\to1$, $1/(1+\lambda_2)^2\to\underline{U}/(\underline{v}-\underline{c})$. However, this implies that the right-hand side of (36) tends to

$$\frac{1}{\sigma}\frac{\underline{U}-\overline{U}}{\overline{c}-\underline{c}}(\overline{v}-\underline{v})+\left(\frac{\overline{v}-\underline{c}}{\underline{v}-\underline{c}}-1\right)\underline{U},$$

which is strictly positive. Hence, when $\mu\to1$, there is no solution to the system of equations characterizing the candidate equilibrium with pure screening. Note the above argument holds for all σ for which a separating equilibrium does not exist. This in turn implies that for every $\sigma>\sigma^S$, there exists a threshold $\overline{\mu}(\sigma)$ such that for all $\mu>\overline{\mu}(\sigma)$ there is an equilibrium with partial sorting. □

A.8. Proof of Proposition 5

Assume $\sigma > \sigma^S$ and $\mu > \overline{\mu}(\sigma)$ and consider the partial sorting equilibrium, which we have shown exists under these conditions. In this equilibrium, the low-type market utility is determined by (27) and thus independent of σ. Hence, a marginal increase in σ leads to no change in $\underline{U}$.

Consider next the high-type market utility $\overline{U}$ and let d$\sigma > 0$. Further, let the low-type incentive constraint (37) be called (IC) and the customer's free-entry condition (38) be called (FE). These two conditions, together with the first-order condition (36), determine the market utility $\overline{U}$, the total queue length in the mixed market λ_2, and the queue length of high types in the mixed market $\overline{\lambda}_2$. We will show that $\overline{U}$ is monotonically increasing in σ by differentiating (IC) and (FE), and using (36). The total derivative of (IC) is:

$$\frac{\partial(IC)}{\partial\lambda_2}d\lambda_2 + \frac{\partial(IC)}{\partial\overline{\lambda}_2}d\overline{\lambda}_2 + \frac{\partial(IC)}{\partial\overline{U}}d\overline{U} + \frac{\partial(IC)}{\partial\sigma}d\sigma = 0.$$

Solving for $d\lambda_2$ and substituting in

$$\frac{\partial(FE)}{\partial\lambda_2}d\lambda_2 + \frac{\partial(FE)}{\partial\overline{\lambda}_2}d\overline{\lambda}_2 + \frac{\partial(FE)}{\partial\overline{U}}d\overline{U} + \frac{\partial(FE)}{\partial\sigma}d\sigma = 0,$$

we obtain

$$\begin{aligned} 0 = &\left(\frac{\partial(FE)}{\partial\overline{U}} - \frac{\partial(FE)}{\partial\lambda_2}\frac{\partial(IC)/\partial\overline{U}}{\partial(IC)/\partial\lambda_2}\right)d\overline{U} + \left(\frac{\partial(FE)}{\partial\sigma} - \frac{\partial(FE)}{\partial\lambda_2}\frac{\partial(IC)/\partial\sigma}{\partial(IC)/\partial\lambda_2}\right)d\sigma \\ &+ \left(\frac{\partial(FE)}{\partial\overline{\lambda}_2} - \frac{\partial(FE)}{\partial\lambda_2}\frac{\partial(IC)/\partial\overline{\lambda}_2}{\partial(IC)/\partial\lambda_2}\right)d\overline{\lambda}_2. \end{aligned} \tag{40}$$

The first-order condition (36) can be rewritten as

$$\frac{\partial(FE)}{\partial\lambda_2} + \frac{\partial(FE)}{\partial\overline{\lambda}_2}\left(-\frac{\partial(IC)/\partial\lambda_2}{\partial(IC)/\partial\overline{\lambda}_2}\right) = 0 \tag{41}$$

so the last term in (40) is zero. Notice further that $\partial(IC)/\partial\overline{\lambda}_2 > 0, \partial(IC)/\partial\lambda_2 < 0$ and

$$\frac{\partial(FE)}{\partial\overline{\lambda}_2} = -\underbrace{\frac{\phi(\lambda_2 - \overline{\lambda}_2, \lambda_2)}{\partial\overline{\lambda}_2}}_{<0}[(\overline{v} - \overline{c}) - (\underline{v} - \underline{c})] + (\underline{U} - \overline{U}) > 0.$$

Condition (41) thus implies $\partial(FE)/\partial\lambda_2 < 0$. With regard to the term multiplying $d\overline{U}$ in (40), we observe

$$\underbrace{\frac{\partial(FE)}{\partial\overline{U}}}_{(-)} - \underbrace{\frac{\partial(FE)}{\partial\lambda_2}}_{(-)}\underbrace{\frac{\overbrace{\partial(IC)/\partial\overline{U}}^{(+)}}{\partial(IC)/\partial\lambda_2}}_{(-)} < 0.$$

Next, we consider the term in (40) multiplying $d\sigma$ and show

$$\underbrace{\frac{\partial(FE)}{\partial\sigma}}_{(-)} - \underbrace{\frac{\partial(FE)}{\partial\lambda_2}}_{(-)}\underbrace{\frac{\overbrace{\partial(IC)/\partial\sigma}^{(-)}}{\partial(IC)/\partial\lambda_2}}_{(-)} > 0.$$

To prove this we rewrite the latter inequality as

$$\frac{\partial(FE)/\partial\sigma}{\partial(FE)/\partial\lambda_2} < \frac{\partial(IC)/\partial\sigma}{\partial(IC)/\partial\lambda_2}$$

$$\frac{-\frac{\partial\phi(\lambda_2-\overline{\lambda}_2,\lambda_2)}{\partial\sigma}[(\overline{v}-\overline{c})-(\underline{v}-\underline{c})]}{\frac{\partial\phi(\lambda_2,\lambda_2)}{\partial\lambda_2}(\overline{v}-\overline{c}) - \underline{U} - \frac{\partial\phi(\lambda_2-\overline{\lambda}_2,\lambda_2)}{\partial\lambda_2}[(\overline{v}-\overline{c})-(\underline{v}-\underline{c})]} < \frac{-\frac{\partial\phi(\lambda_2-\overline{\lambda}_2,\lambda_2)}{\partial\sigma}}{\frac{\partial\phi(\lambda_2,\lambda_2)}{\partial\lambda_2} - \frac{\partial\phi(\lambda_2-\overline{\lambda}_2,\lambda_2)}{\partial\lambda_2}}$$

which in turn can be written as

$$\frac{\partial\phi(\lambda_2,\lambda_2)}{\partial\lambda_2}(\overline{v}-\overline{c}) - \underline{U} - \frac{\partial\phi(\lambda_2-\overline{\lambda}_2,\lambda_2)}{\partial\lambda_2}[(\overline{v}-\overline{c})-(\underline{v}-\underline{c})]$$

$$< \left(\frac{\partial \phi(\lambda_2, \lambda_2)}{\partial \lambda_2} - \frac{\partial \phi(\lambda_2 - \overline{\lambda}_2, \lambda_2)}{\partial \lambda_2} \right) [(\overline{v} - \overline{c}) - (\underline{v} - \underline{c})]$$

$$\Leftrightarrow \quad \frac{\partial \phi(\lambda_2, \lambda_2)}{\partial \lambda_2} (\underline{v} - \underline{c}) < \underline{U}$$

Given $\partial \phi(\lambda_1, \lambda_1)/\partial \lambda_1 (\underline{v} - \underline{c}) = \underline{U}$, concavity of $\phi(\lambda, \lambda) = \lambda/(1+\lambda)$ implies that the latter inequality is satisfied as long as $\lambda_1 < \lambda_2$. In this case, (40) has the form $(-)d\overline{U} + (+)d\sigma = 0$, which then implies $d\overline{U} > 0$.

To complete the argument, we thus need to show that the property $\lambda_1 < \lambda_2$ is satisfied in the partial sorting equilibrium. We start by showing that in order not to have a profitable deviation, λ_1 must satisfy

$$\frac{1}{1+\lambda_1} \left(1 - \frac{\sigma \lambda_1}{1+\lambda_1} \right) \geq \frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}}. \tag{42}$$

To see this, suppose that this inequality is violated and let $\varepsilon > 0$. Consider an alternative mechanism, described by $(\underline{\lambda}, \overline{\lambda}, \underline{x}, \overline{x}, \underline{t}, \overline{t})$, with $\underline{\lambda} = \lambda_1 - \varepsilon$, $\overline{\lambda} = \varepsilon$, $\underline{x}, \overline{x}$ determined by the priority rule (8) and overall feasibility condition (10) holding as equality, and transfers determined by the participation constraints. The difference in a customer's expected payoff between the alternative mechanism and the one in market 1 is

$$\left(\frac{\lambda_1}{1+\lambda_1} - \frac{\lambda_1 - \varepsilon}{1+\lambda_1 - \varepsilon} \right) [(\overline{v} - \overline{c}) - (\underline{v} - \underline{c})] + \varepsilon(\underline{U} - \overline{U}) > 0.$$

The mechanism thus yields a higher payoff than the one attracting only low types. It satisfies the low-type incentive constraint (6) if

$$\overline{x} = \frac{1 + (1-\sigma)(\underline{\lambda} + \overline{\lambda})}{(1 + (\underline{\lambda} + \overline{\lambda}))(1 + \sigma \underline{\lambda} + (1-\sigma)(\underline{\lambda} + \overline{\lambda}))} \tag{43}$$

$$= \frac{1 + (1-\sigma)\lambda_1}{(1+\lambda_1)(1 + \sigma(\lambda_1 - \varepsilon) + (1-\sigma)\lambda_1)} \tag{44}$$

$$= \frac{1}{1+\lambda_1} \left(1 - \frac{\sigma(\lambda_1 - \varepsilon)}{1 + \sigma(\lambda_1 - \varepsilon) + (1-\sigma)\lambda_1} \right) \tag{45}$$

$$\leq \frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}}. \tag{46}$$

If (42) is violated, then for $\varepsilon > 0$ sufficiently small, the above inequality is satisfied and a profitable deviation exists. Hence, condition (42) must be satisfied in a partial sorting equilibrium. Notice that in a partial sorting equilibrium, λ_2 belongs to the interval $(\underline{l}, \overline{l})$, where the bounds $\underline{l}$ and $\overline{l}$ are, respectively, defined by (25) and (24). By definition of the lower bound $\underline{l}$ and (42), we have:

$$\frac{1}{1+\lambda_1} \left(1 - \frac{\sigma \lambda_1}{1+\lambda_1} \right) \geq \underbrace{\frac{1}{1+\underline{l}} \left(1 - \frac{\sigma \underline{l}}{1+\underline{l}} \right)}_{= \frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}}} \tag{47}$$

As can be verified, the term $\frac{1}{1+\lambda}\left(1 - \frac{\sigma\lambda}{1+\lambda}\right)$ is strictly decreasing λ.[23] Hence, inequality (47) implies $\lambda_1 \leq \underline{l} < \lambda_2$. From $\lambda_1 < \lambda_2$, it then follows $d\overline{U} > 0$.

Since in a partial sorting equilibrium, $\underline{U}$ is constant in σ and $\overline{U}$ is strictly increasing in σ, it follows that an increase in σ constitutes a pareto improvement. The fact that $\underline{U}$ is constant and $\overline{U}$ is increasing in σ also implies that a high type's trading probability, $\overline{x} = (\underline{U} - \overline{U})/(\overline{c} - \underline{c})$, is decreasing in σ (see (37)). □

A.9. Proof of Proposition 6

Suppose first the equilibrium features pure or partial sorting so that there is a market, say market 2, where both types trade. Consider then a marginal change in the allocation where we increase the high-type providers' trading probability, $d\bar{x}_2 > 0$, and reduce the trading probability of low types (by $d\bar{x}_2 \bar{\lambda}_2/\underline{\lambda}_2$) so that the customers' total trading probability in that market is unchanged (entry does not vary). Note that such change is always feasible (by slightly relaxing the priority to low types in the mechanism) and is incentive-compatible if the transfer to high-type providers is adjusted so that $\bar{U}$ decreases by $d\bar{x}_2(\bar{c} - \underline{c})$. The level of profits of customers in the pooling market then changes by

$$d\pi = \bar{\lambda}_2(\bar{v} - \bar{c})d\bar{x}_2 - \underline{\lambda}_2 \left(\frac{\bar{\lambda}_2}{\underline{\lambda}_2} \right) d\bar{x}_2(\underline{v} - \underline{c}) + \bar{\lambda}_2 d\bar{x}_2(\bar{c} - \underline{c}),$$

[23] The first derivative of $\frac{1}{1+\lambda}\left(1 - \frac{\sigma\lambda}{1+\lambda}\right)$ with respect to λ is given by $-\frac{1+\sigma+(1-\sigma)\lambda}{1+\lambda} < 0$.

which is always strictly greater than the variation in the total utility gains of high types, equal to $-\bar{\lambda}_2 d\bar{x}_2(\bar{c}-\underline{c})$. This establishes that total surplus is strictly higher at the modified allocation.

Next, we turn our attention to the case where the equilibrium features pure sorting, so high and low types trade in separate markets. As before, we consider an increase in the trading probability of high-type providers, $d\bar{x} > 0$, now attained by increasing the mass of customers present in the high-type market. The mass in the low-type market is then reduced by the same amount so that entry is unchanged. It can be verified that the reduction that is induced in the trading probability of low types is

$$d\underline{x} = -\left(\frac{1+\bar{\lambda}}{1+\underline{\lambda}}\right)^2 \frac{(1-\mu)}{\mu}\left(\frac{\underline{\lambda}}{\bar{\lambda}}\right)^2.$$

The total trading probability changes by:

$$d\bar{x}(1-\mu) + d\underline{x}\mu = d\bar{x}(1-\mu)\left[1-\left(\frac{\underline{\lambda}/(1+\underline{\lambda})}{\bar{\lambda}/(1+\bar{\lambda})}\right)^2\right] > 0,$$

where the positive sign follows from the fact that in the pure sorting equilibrium, we have $\bar{\lambda} > \underline{\lambda}$. The share of trades with high-type providers increases as well, thus establishing that total surplus increases. □

A.10. Slack incentive constraints

In this Section, we present the equilibrium conditions as a solution of P^{aux} for the case where $\overline{v}-\overline{c} < \underline{v}-\underline{c}$ and incentive constraints are slack. From Cai et al. (2022) (Proposition 2), we can restrict attention to three possibilities: pure sorting, pure screening, and partial sorting with two submarkets where the low-quality provider is active in just one of the two markets. To simplify notation, we define $\underline{\Delta} := \underline{v}-\underline{c}$ and $\overline{\Delta} := \overline{v}-\overline{c}$.

Pure sorting Assuming that the incentive constraints of both types are non-binding, the queue length in the low- and high-quality markets, $\underline{\lambda}$ and $\bar{\lambda}$, are determined by the first-order condition of P^{aux} with respect to $\underline{\lambda}$ (resp. $\bar{\lambda}$) under the restriction $\bar{\lambda}=0$ (resp. $\underline{\lambda}=0$). As shown in the proof of Proposition 4, for the low-quality market, this yields the following queue length and market utility:

$$\lambda_L := \frac{\sqrt{K}}{\sqrt{\underline{\Delta}}-\sqrt{K}}, \quad \underline{U} = \left(\sqrt{\underline{\Delta}}-\sqrt{K}\right)^2. \tag{48}$$

Likewise, for the high-quality market, we have

$$\lambda_H := \frac{\sqrt{K}}{\sqrt{\overline{\Delta}}-\sqrt{K}}, \quad \overline{U} = \left(\sqrt{\overline{\Delta}}-\sqrt{K}\right)^2. \tag{49}$$

Pure screening In the case of pure screening, P^{aux} has to be solved by a pair $(\underline{\lambda}, \lambda)$ such that $\underline{\lambda}/\lambda = \mu$. Letting $m(\lambda) := \phi(\lambda, \lambda)$, the first-order conditions with respect to $\underline{\lambda}$ and λ of maximizing (11) are given by:

$$\phi_{\underline{\lambda}}(\underline{\lambda}, \lambda)(\underline{\Delta}-\overline{\Delta}) - (\underline{U}-\overline{U}) = 0, \tag{50}$$

$$\phi_{\lambda}(\underline{\lambda}, \lambda)(\underline{\Delta}-\overline{\Delta}) + m'(\lambda)\overline{\Delta} - \overline{U} = 0. \tag{51}$$

Using the free-entry condition,

$$\phi(\underline{\lambda}, \lambda)\underline{\Delta} + (m(\lambda)-\phi(\underline{\lambda}, \lambda))\overline{\Delta} - \underline{\lambda}\underline{U} - (\lambda-\bar{\lambda})\overline{U} = K \tag{52}$$

and $\underline{\lambda} = \mu\lambda$, this pins down the candidate equilibrium values $(\underline{\lambda}, \lambda, \underline{U}, \overline{U})$.

Partial sorting For the partial sorting equilibrium, we are looking for a tuple $(\lambda_1, \underline{\lambda}_2, \lambda_2, \underline{U}, \overline{U})$ of queue lengths and market utilities, where $(\underline{\lambda}_2, \lambda_2)$ are the low-type and total queue lengths in the pooling market, while λ_1 is the queue length in the market where only high-type providers search. The market utility $\overline{U}$ and the queue length λ_1 are the same as in the pure sorting equilibrium, described by (49). The queue lengths in the pooling market and the utility of the low-quality providers are determined by (50), (51), and the free entry condition (52).

Parameter regions To get the parameter regions, consider first the boundary between the pure sorting and partial sorting regions. At the boundary, the low- and high-type market utilities, $\underline{U}$ and $\overline{U}$, are given by the solution for the sorting equilibrium, specified in (48) and (49). Substituting for these market utilities in the right-hand side of (50) and evaluating it at $\underline{\lambda} = \lambda = \lambda_L$ (the queue length in the low-type market in the pure sorting region), the right-hand side turns positive if

$$\phi_{\underline{\lambda}}(\underline{\lambda}, \lambda)|_{\underline{\lambda}=\lambda=\lambda_L}(\underline{\Delta} - \overline{\Delta}) - (\underline{U} - \overline{U}) > 0 \iff \sigma > \underbrace{\frac{\underline{\Delta} - \sqrt{\underline{\Delta}\overline{\Delta}}}{(\sqrt{\underline{\Delta}} + \sqrt{\overline{\Delta}})(\sqrt{\underline{\Delta}} - \sqrt{K})}}_{\equiv \hat{\sigma}^S}.$$

For $\sigma > \hat{\sigma}^S$, the system of equations pinning down the utilities and queue lengths in the candidate partial sorting equilibrium are such that $\underline{\lambda}_2 < \lambda_2$. The allocation is compatible with the distribution of provider types in the population if and only if $\mu < \underline{\lambda}_2/\lambda_2$ (see the argument in Section A.7). This pins down the boundary with the pure screening equilibrium. Given $\sigma > \hat{\sigma}^S$, if $\mu < \underline{\lambda}_2/\lambda_2$, the candidate equilibrium features partial sorting; if $\mu \geq \underline{\lambda}_2/\lambda_2$, the candidate equilibrium features pure screening.

Incentive compatibility The described allocation constitutes an equilibrium if the low-cost provider has no incentives to imitate the high-cost provider. Given the candidate allocation, this condition can be easily checked by calculating the trading probability of the high-cost provider. When $\sigma \leq \sigma^S$, incentive compatibility requires

$$\frac{m(\lambda_H)}{\lambda_H} \leq \frac{\underline{U} - \overline{U}}{\overline{c} - \underline{c}} \iff \frac{(\sqrt{\underline{\Delta}} - \sqrt{K})^2 - (\sqrt{\overline{\Delta}} - \sqrt{K})^2}{\overline{c} - \underline{c}} \geq \frac{\sqrt{\overline{\Delta}} - \sqrt{k}}{\sqrt{\overline{\Delta}}}$$

When $\underline{v} = \overline{v}$, we have $\overline{c} - \underline{c} = \underline{\Delta} - \overline{\Delta}$. After some manipulation, it can be easily seen that the above inequality simplifies to $\underline{\Delta} > \overline{\Delta}$, which is always satisfied. The verification of equilibrium existence in the other two regions is analogous: using the candidate market utilities and queue lengths, incentive compatibility is satisfied if and only if the trading probabilities of high-quality providers fall below $(\underline{U} - \overline{U})/(\overline{c} - \underline{c})$.

A.11. The case $\sigma = 1$

Proposition 13. *Assume $\sigma = 1$, $\phi(\underline{\lambda}, \underline{\lambda} + \overline{\lambda}) = \frac{\underline{\lambda}}{1+\underline{\lambda}}$. A pure screening equilibrium exists and is unique if and only if*

$$\mu \leq \underbrace{\frac{\sqrt{K(\underline{v} - \underline{c})}}{\overline{v} - \underline{v} + \sqrt{K(\underline{v} - \underline{c})}}}_{=\overline{\mu}(1)}. \tag{53}$$

If (53) is violated, there exists a partial sorting equilibrium satisfying the following properties:

- *low-type queue lengths are the same in both active submarkets;*
- *customers make zero profits with high-type providers.*

Proof. With $\sigma = 1$, we have $\phi(\underline{\lambda}, \lambda) = 1/(1 + \underline{\lambda})$. Notice that the proof of Lemma 10 continues to apply, so at any solution of P^{aux} the low-type feasibility constraint (8) binds. However, since Assumption 1 is not satisfied, the feasibility constraint (10) does not necessarily hold as equality at a solution of P^{aux}. Nevertheless, we can show that for any mechanism that does not satisfy the overall feasibility constraint (10), there is an alternative mechanism that yields the same expected payoff for the customer and satisfies (10). To see this, consider a mechanism $(\underline{x}, \overline{x}, \underline{t}, \overline{t}, \underline{\lambda}, \overline{\lambda})$ such that

$$\underline{\lambda}\,\underline{x} + \overline{\lambda}\,\overline{x} < \phi(\underline{\lambda} + \overline{\lambda}, \underline{\lambda} + \overline{\lambda}).$$

For a customer, this mechanism generates a payoff equal to

$$\underline{\lambda}(\underline{x}(\underline{v} - \underline{c}) - \underline{U}) + \overline{\lambda}(\overline{x}(\overline{v} - \overline{c}) - \overline{U}).$$

If $\overline{x}(\overline{v} - \overline{c}) - \overline{U} > 0$, then a marginal increase in $\overline{\lambda}$ is feasible and strictly improves the customer's payoff (since the low-type feasibility (8) is binding, the high-type feasibility constraint (9) is slack). If $\overline{x}(\overline{v} - \overline{c}) - \overline{U} = 0$, then the customer is indifferent between the mechanism $(\underline{x}, \overline{x}, \underline{t}, \overline{t}, \underline{\lambda}, \overline{\lambda})$ and a mechanism $(\underline{x}, \overline{x}, \underline{t}, \overline{t}, \underline{\lambda}, \overline{\lambda}')$, where $\overline{\lambda}'$ is such that the overall feasibility constraint (10) is satisfied with equality.

Following this argument, we can consider a modified problem where customers are restricted to post mechanisms that satisfy the overall feasibility constraint (10) (no rationing). For this case our existence proof (Proposition 12) applies. Moreover, since for each mechanism that features rationing, there is another one that does not and generates a weakly higher payoff, this implies that also in the original environment, there is an equilibrium where customers only post mechanisms that satisfy the overall feasibility constraint (10). For this equilibrium all arguments in the proofs of Lemma 3 and Proposition 4 apply. Hence, in the equilibrium where customer post mechanisms that satisfy the overall feasibility constraint (10), there are at most two active submarkets, one that attracts only low types and one that attracts also the high type.

To characterize these and other equilibria, we distinguish the cases according to whether or not customers can make positive profits with high-type providers. Incentive compatibility for low-type providers implies that the maximal incentive compatible trading probability for high-type providers is $x^{max} = (\underline{U} - \overline{U})/(\overline{c} - \underline{c})$. When $x^{max}(\overline{v} - \overline{c}) - \overline{U} = 0$, any incentive compatible mechanism yields a weakly negative profit with high types. In terms of market utilities, this condition can be written as $\overline{U}(\overline{v} - \underline{c}) = \underline{U}(\overline{v} - \overline{c})$.

We start by considering the case where customers cannot make positive profits with high-type providers: $\overline{U}(\overline{v} - \underline{c}) = \underline{U}(\overline{v} - \overline{c})$. An optimal mechanism then satisfies $\overline{\lambda} = 0$ or $\overline{x} = \overline{x}^{max}$ or both. Since profits with high types are zero, any equilibrium mechanism must maximize the payoff with low-type providers and thus solve the problem:

$$\max_{\underline{\lambda}} \quad \frac{\underline{\lambda}}{1+\underline{\lambda}}(\underline{v} - \underline{c}) - \underline{\lambda}\,\underline{U}$$

The optimal low-type queue length satisfies the first-order condition:

$$\frac{1}{(1+\underline{\lambda})^2}(\underline{v} - \underline{c}) = \underline{U}.$$

Together with the free-entry condition, $\underline{\lambda}/(1+\underline{\lambda})(\underline{v} - \underline{c}) - \underline{\lambda}\underline{U} = K$, we can solve for the market utility

$$\underline{U} = (\sqrt{\underline{v} - \underline{c}} - \sqrt{K})^2,$$

and the optimal queue length of low-type providers

$$\underline{\lambda}^* = \frac{\sqrt{K}}{\sqrt{\underline{v} - \underline{c}} - \sqrt{K}}.$$

The market utility for high-type providers is then determined by our initial condition $\overline{U}(\overline{v} - \underline{c}) = \underline{U}(\overline{v} - \overline{c})$ and given by:

$$\overline{U} = \frac{\overline{v} - \overline{c}}{\overline{v} - \underline{c}}(\sqrt{\underline{v} - \underline{c}} - \sqrt{K})^2.$$

The overall feasibility constraint (10) can be written as

$$\overline{x} \leq \frac{1}{\lambda - \underline{\lambda}^*}\left(\frac{\lambda}{1+\lambda} - \frac{\underline{\lambda}^*}{1+\underline{\lambda}^*}\right) = \frac{1}{(1+\underline{\lambda}^*)(1+\lambda)}.$$

For mechanisms with $\lambda > \underline{\lambda}^*$, we have $\overline{x} = \overline{x}^{max} = (\underline{U} - \overline{U})/(\overline{c} - \underline{c})$, so the previous inequality can be written as:

$$\frac{1}{1+\lambda} \geq \frac{\underline{v} - \underline{c} - \sqrt{K(\underline{v} - \underline{c})}}{\overline{v} - \underline{c}}. \tag{54}$$

Consistency with the population parameters then requires that there is at least one market with

$$\underline{\lambda}^*/\lambda \leq \mu.$$

This inequality is compatible with (54) if and only if

$$\mu \geq \frac{\sqrt{K(\underline{v} - \underline{c})}}{\overline{v} - \underline{v} + \sqrt{K(\underline{v} - \underline{c})}}. \tag{55}$$

We now show that if (55) is satisfied with strict inequality, there exists a partial sorting equilibrium in which customers post mechanisms that satisfy the overall feasibility constraint (10). The pair of queue lengths in the first market is $(\underline{\lambda}^*, 0)$ and in the second market is $(\underline{\lambda}^*, \lambda - \underline{\lambda}^*)$, where λ solves (54) with equality. Since (55) is satisfied as a strict inequality, the fraction of low types in the second market $\mu_2 \equiv \underline{\lambda}^*/\lambda$ is strictly smaller than the population ratio μ. Hence, we can find two values, β_1 and β_2, that describe the measures of customers visiting, respectively, the first and second submarket, such that

$$\frac{\beta_1}{\beta_1+\beta_2} \cdot 1 + \frac{\beta_2}{\beta_1+\beta_2}\mu_2 = \mu, \tag{56}$$

$$\frac{\beta_2}{\beta_1+\beta_2}(1-\mu_2) = 1 - \mu. \tag{57}$$

Note that if condition (55) holds as an equality, this system of equalities is solved by $\beta_1 = 0$, hence the first submarket is inactive (pure screening).

If instead condition (55) holds as a strict inequality, we can construct additional, payoff-equivalent equilibria where customers post mechanisms that satisfy the overall feasibility constraint as a strict inequality and thus feature rationing. In particular, there is an equilibrium where all customers post the same mechanism so that all providers visit a single submarket. In this case, consistency with the population parameters requires that the fraction of low-type providers in the single market is equal to the population ratio μ, hence the overall queue length λ satisfies $\underline{\lambda}^*/\underline{\lambda} = \mu$. Given that there is a unit measure of providers in the economy, the measure of customers entering the market is simply $1/\lambda$. Provided condition (55) holds as a strict inequality, the overall feasibility condition (54) is then satisfied as a strict inequality. Hence, high types are rationed in meetings where no low-type providers are present.

Evidently, in this case, we can construct other equilibria with a set of M markets, each attracting a measure β of customers, such that

$$\frac{1}{\beta(M)}\int_m \frac{\underline{\lambda}^*(m)}{\underline{\lambda}(m)}\mathrm{d}\beta(m)=\mu.$$

If condition (55) is not satisfied, there must exist an equilibrium where customers post mechanisms that satisfy the overall feasibility constraint (10) and make strictly positive profits with high-type providers. Market utilities thus satisfy $\overline{U}(\overline{v}-\underline{c}) > \underline{U}(\overline{v}-\overline{c})$. It is in fact easy to see that customers never find it optimal to post a mechanism that features rationing when they can make positive profits with high-type providers: any mechanism, $(\underline{x}, \overline{x}, \underline{t}, \overline{t}, \underline{\lambda}, \overline{\lambda})$, that satisfies the overall feasibility constrain (10) as a strict inequality can be improved upon by either increasing $\overline{\lambda}$ if $\overline{x} = \overline{x}^{max}$ or $\overline{x}$ if $\overline{x} < \overline{x}^{max}$ or both. Such change does not affect the payoff that is generated with the low types but strictly increases a customer's payoff with the high type.

Having observed that any optimal mechanism satisfies the overall feasibility constraint (10), it can be proven that there is no solution of P^{aux} with $\overline{\lambda} = 0$. For any mechanism that attracts only the low type, there is an alternative mechanism that attracts the same queue length of low-type providers and a positive queue of high-type providers, which generates the same expected profit with the low type and a strictly positive profit with the high type, as we show in the proof of Proposition 4, Section A.7, conditions (31)-(35). Hence, given $\overline{U}(\overline{v}-\underline{c}) > \underline{U}(\overline{v}-\overline{c})$, there is no solution of P^{aux} with $\overline{\lambda} = 0$. Moreover, since the solution of P^{aux} is unique on the domain with $\overline{\lambda} > 0$, it follows that any equilibrium where customers make strictly positive profits with high-type providers is of the 'pure screening' type: all customers and providers trade in a single submarket according to a mechanism that features no rationing. □

References

Akerlof, G.A., 1970. The market for "lemons": quality uncertainty and the market mechanism. Q. J. Econ. 84 (3), 488–500.
Albrecht, J.W., Gautier, P.A., Vroman, S.B., 2014. Efficient entry with competing auctions. Am. Econ. Rev. 104 (10), 3288–3296.
Auster, S., Gottardi, P., 2019. Competing mechanisms in markets for lemons. Theor. Econ. 14 (3), 927–970.
Auster, S., Gottardi, P., Wolthoff, R.P., 2022. Simultaneous search and adverse selection.
Batra, H., Michaud, A., Mongey, S., 2023. Online job posts contain very little wage information. Technical report. National Bureau of Economic Research.
Burdett, K., Judd, K.L., 1983. Equilibrium price dispersion. Econometrica 51, 955–969.
Cai, X., Gautier, P., Wolthoff, R., 2022. Meetings and mechanisms. Int. Econ. Rev.
Chang, B., 2018. Adverse selection and liquidity distortion. Rev. Econ. Stud. 85 (1), 275–306.
Davoodalhosseini, S.M., 2019. Constrained efficiency with adverse selection and directed search. J. Econ. Theory 183, 568–593.
Eeckhout, J., Kircher, P., 2010. Sorting versus screening - search frictions and competing mechanisms. J. Econ. Theory 145, 1354–1385.
Gale, D., 1992. A Walrasian theory of markets with adverse selection. Rev. Econ. Stud. 59 (2), 229–255.
Guerrieri, V., Shimer, R., Wright, R., 2010. Adverse selection in competitive search equilibrium. Econometrica 78 (6), 1823–1862.
Hall, R.E., Krueger, A.B., 2012. Evidence on the incidence of wage posting, wage bargaining, and on-the-job search. Am. Econ. J. Macroecon. 4 (4), 56–67.
Inderst, R., Müller, H.M., 2002. Competitive search markets for durable goods. Econ. Theory 19 (3), 599–622.
Lester, B., Shourideh, A., Venkateswaran, V., Zetlin-Jones, A., 2019. Screening and adverse selectionin frictional markets. J. Polit. Econ. 127 (1), 338–377.
Lester, B., Visschers, L., Wolthoff, R., 2015. Meeting technologies and optimal trading mechanisms in competitive search markets. J. Econ. Theory 155, 1–15.
McAfee, R.P., 1993. Mechanism design by competing sellers. Econometrica, 1281–1312.
Moen, E.R., 1997. Competitive search equilibrium. J. Polit. Econ. 105, 385–411.
Shimer, R., 2005. The assignment of workers to jobs in an economy with coordination frictions. J. Polit. Econ. 113 (5), 996–1025.
Williams, B., 2021. Search, liquidity, and retention: screening multidimensional private information. J. Polit. Econ. 129 (5), 1487–1507.