# Cream Skimming and Information Design in Matching Markets

Romanyuk, Gleb and Smolin, Alexey

Harvard University, University of Bonn

14 May 2018

# Cream Skimming and Information Design in Matching Markets

Gleb Romanyuk, Alex Smolin*

May 14, 2018

**Abstract**

Short-lived buyers arrive to a platform over time and randomly match with sellers. The sellers stay at the platform and sequentially decide whether to accept incoming requests. The platform designs what buyer information the sellers observe before deciding to form a match. We show full information disclosure leads to a market failure because of excessive rejections by the sellers. If sellers are homogeneous, then coarse information policies are able to restore efficiency. If sellers are heterogeneous, then simple censorship policies are often constrained efficient as shown by a novel method of calculus of variations.

*Romanyuk: Harvard University, 1805 Cambridge St, Cambridge, MA 02138, USA (email: gleb.romanyuk@gmail.com); Smolin: University of Bonn, Lennestraße 37, 53113, Bonn, Germany (email: alexey.v.smolin@gmail.com). We thank John Asker, Coeditor, and three anonymous referees for their productive suggestions. This paper builds on Romanyuk's dissertation chapter written at Harvard University, Department of Economics; he is indebted to Susan Athey, Drew Fudenberg, Greg Lewis, Tomasz Strzalecki, as well as to Chiara Farronato and Andrei Hagiu for their guidance and support. We are grateful to Andrey Fradkin, Ben Golub, Divya Kirti, Sergei Kovbasyuk, Jeffrey Picel, as well as seminar participants at Harvard University and MIT for helpful discussions. The views expressed in this article are those of the authors and do not necessarily reflect those of the employing organizations.

# 1 Introduction

A primary objective for matching platforms, such as platforms for ride hailing (Uber, Lyft), accommodation rental (Airbnb, Craigslist), and freelance labor (Upwork, TaskRabbit) is to facilitate value-creating transactions. Providing detailed information to the market participants about the goods or trading partners is a way to ensure that they identify and pursue the most valuable matches. Matching platforms spend significant resources eliciting match-relevant information from users and incorporating it into the platform's design.[1] Sometimes, however, the relevant information is only partially revealed to the users. As an example, Uber does not show its drivers the passenger destination until after the driver has picked up the passenger. Why do such platforms choose not to fully disclose all relevant information? On what does an optimal information policy depend?

We study the information intermediation problem in the context of a dynamic matching market. Buyers and sellers try to match on a platform over time. A stream of buyers contact sellers, who then review buyer information and choose whether or not to accept them. Buyers are short lived and indifferent to with whom seller they match. Sellers are long lived and derive heterogeneous payoffs from being matched with different buyers. The utilities are nontransferable. Importantly, the market has capacity constraints—if a seller accepts a buyer, the seller becomes unavailable for a fixed time. The platform designs a stationary and non-discriminatory information policy that governs what buyer information is revealed to sellers before they accept buyer requests. It does so to maximize a weighted average of a buyer and seller steady-state surplus.

We analyze the setting in Section 3. First, we show in Proposition 1 that for an arbitrary information policy, the steady state equilibrium exists and is essentially unique. Intuitively, the more sellers available in the steady state, the lower the rate at which they are matched with the buyers. It reduces their continuation value and increases their frequency

[1] Tadelis and Zettelmeyer (2015) show that in wholesale automobile auctions, information disclosure helps match heterogeneous buyers to cars of varying quality. Einav, Farronato, and Levin (2016) name eliciting and aggregating the dispersed user information as a key objective of peer-to-peer platforms.

of acceptance. This results in less sellers being available and leads to unique equilibrium characterization.

Second, we show in Proposition 2 that the full disclosure outcome is inefficient: There exists a seller's strategy profile that generates higher surplus on both market sides. The market failure is driven by *cream-skimming* seller behavior: Each seller tries to keep his schedule open by rejecting low-value matches in order to increase his individual chances of getting high-value matches. These rejections impose negative externalities on both sides of the market. On the buyer side, the set of matches that create positive surplus for the sellers is distinct from the set of matches that create positive surplus for the buyers. When the sellers make the acceptance decision, they do not internalize buyer surplus. On the seller side, by rejecting a buyer, a seller remains available in the marketplace and competes with other sellers for incoming requests. As a result, the other sellers face fewer valuable buyers and realize lower profits.

Third, we study the platform-optimal information policies that are constrained efficient. We highlight that an information policy determines not only a seller's stage payoff, but also his continuation value. An optimal information policy must balance the positive effect of information disclosure on the match quality and the negative effect on the match rate. The effect of information disclosure on the match quality is straightforward. Holding the match rate fixed, each seller benefits from more details about buyers because it allows him to make a more informed decision. However, information disclosure can exacerbate cream-skimming behavior and decrease the match rate because greater transparency increases return on search.

We find that if sellers are homogeneous, information control can fully restore efficiency—the platform can implement any Pareto efficient outcome that delivers positive surplus to sellers. Theorem 1 demonstrates that the optimal policies are partially informative and take a simple form. They send only two possible signals, which can be interpreted as recommendations to accept or reject the buyer request. The signal likelihoods are chosen to achieve a given outcome.

If sellers are heterogeneous, then it is generally impossible to fully restore efficiency. Moreover, binary recommendation policies are less effective because different sellers respond differently to the same information. We focus on a linear payoff environment and show how the form of the optimal information policies depends on details of the seller type distribution in Theorem 2. If its density is log-concave, or the density is unimodal and the platform maximizes only the matching rate, then a *censorship* policy is optimal. Such a policy pools more attractive buyers and fully reveals less attractive buyers. If the seller density is weakly decreasing, then, interestingly, the full disclosure policy *is* optimal.

We discuss how our analysis can inform the design of digital marketplaces in Section 4. We observe that most platforms actively engage in information design by disclosing only selected information categories to sellers. We highlight how such information selection can avoid inefficient cream skimming and create a balanced market. Finally, we discuss the role that strategic buyers and flexible pricing schemes can play in the optimal platform design.

**Related literature.** The role of information in matching and peer-to-peer markets is an active area of economic research. Chade (2006) highlights an adverse selection problem leading to the acceptance curse—a negative inference about the accepting party when it is privately informed. In contrast, we focus on cream skimming driven by moral hazard. Hoppe, Moldovanu, and Sela (2009) show that the option to disclose information may lead to wasteful, costly signaling that can offset the benefits of improved matching. Coles, Kushnir, and Niederle (2013) study the case of costless but restricted signalling and show that it generally improves the match rate but can decrease the receiver's welfare. Levin and Milgrom (2010) show that in online advertising markets, standardization, or "conflation," helps to address at least two problems associated with excessive targeting: adverse selection and reduced competition. Ostrovsky and Schwarz (2010) study information disclosure in a static matching market and show that equilibrium information policies provide no incentives for early-contracting unraveling. Lauermann (2012) highlights the cream-skimming behavior in search markets and the resulting inefficiencies caused by information provision which is also the focus of our paper.[2]

Our work differs from the theory of centralized matching because the market participants must inspect potential matches to identify the valuable ones. Anderson and Smith (2010) show how learning and reputation concerns can preclude positive assortative matching in dynamic markets. Ashlagi, Jaillet, and Manshadi (2013) and Akbarpour, Li, and Oveis Gharan (2017) study the role of uncertainty in centralized matching markets and show that waiting for the market to thicken improves matching in kidney exchanges because individual agents do not fully incorporate the benefits of waiting. We emphasize the opposite effect: individual sellers wait too long because they do not fully incorporate the buyer loss from rejections.

At the same time, this paper contributes to the rapidly growing literature on information design in dynamic settings (e.g., Ely (2017); Che and Hörner (2018); Smolin (2017); Orlov, Skrzypacz, and Zryumov (2018)).[3] In contrast to the existing studies, we investigate stationary information policies with heterogeneous and forward-looking audience. In a related paper, Kovbasyuk and Spagnolo (2017) study information policies in a dynamic market with competitive prices and myopic players. They analyze the effect of past experience disclosure and highlight that revealing negative experiences can exclude too many sellers from the market. Notably, in solving for an optimal policy in the case of homogeneous sellers, we relate to the efficiency approach of Smolin (2017).[4]

[2]In comparison to standard search models, our market features the congestion externality but not the thick market externality.

[3]Kamenica and Gentzkow (2011) and Bergemann and Morris (2016) outline the information design paradigm in static settings.

[4]This approach has proven useful in other information design problems as well (e.g., Ely and Szydlowski (2017)).

## 2 Model

**Matching Market.** Three parties are involved in the matching process: sellers, buyers, and the platform itself. Time is continuous.

The unit mass of sellers always stays on the platform; they never leave or arrive. At each moment in time, a seller is either available or busy. An available seller is sequentially presented with buyer requests and decides whether to accept or reject them in order to maximize his total payoff. Busy sellers do not receive any requests. When an available seller accepts a buyer, he becomes busy for a fixed time $\tau$. The higher the value of $\tau$, the fewer buyers the seller can match with during the same period of time.

Buyers gradually arrive to the platform at a flow rate $\beta$ so that within time interval $dt$ a deterministic mass $\beta dt$ of buyers arrives. Each new buyer contacts one of the available sellers. The seller is chosen uniformly at random from the pool of available sellers. If the buyer is accepted, he stays while the match lasts, for the time $\tau$; otherwise, he leaves the platform.

**Assumption 1.** (Market Capacity) *It is possible for sellers to accept all buyers: $\beta\tau < 1$.*

The capacity assumption simplifies the exposition. Relaxing it requires more notation to deal with either automatic rejections or queues. Figure 1 illustrates the matching process.

**Match Heterogeneity.** There are two dimensions of heterogeneity in the market. First, each seller has a heterogeneous match payoff across buyers. Second, different sellers have different payoff functions. Let $x$ denote a buyer type, with the interpretation that $x$ is comprised of buyer characteristics observed by the platform.[5] The buyer types are distributed over a compact convex set $X \subseteq \mathbb{R}$ according to a distribution function $F$ with full support. Let $y$ denote a seller type, with the interpretation that $y$ is comprised of seller characteristics

[5] Buyer type $x$ captures the payoff-relevant information the platform elicits from the buyer, whether passively from the buyer's cookies and queries or actively by asking questions. For example, on Uber, $x$ would incorporate the rider's destination; on Airbnb, $x$ would incorporate the guest's race, age, and gender.

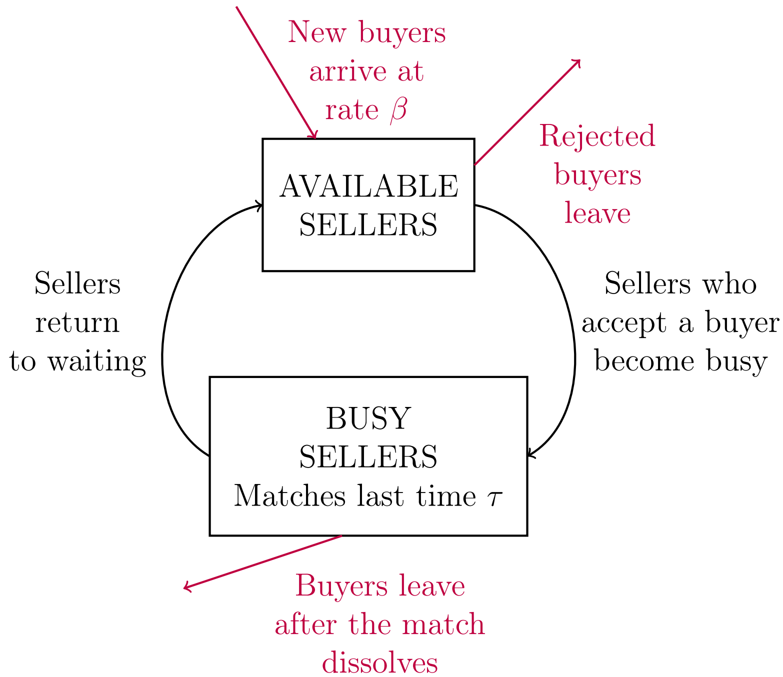


Figure 1: Matching process. Buyers arrive at exogenous rate $\beta$ and contact available sellers. If rejected, a buyer leaves the platform. If accepted, the buyer forms a match which lasts for time $\tau$. After that time elapses, the buyer leaves the platform, and the seller returns to waiting.

unobserved by the platform.[6] The seller types are distributed over a compact convex set $Y \subseteq \mathbb{R}$ according to a distribution function $G$ with full support that admits a differentiable density $g$. The seller profit for one match is $\pi(x, y)$, the buyer utility is $u > 0$.[7]

**Assumption 2.** (Profit Variability) *Profits $\pi(x, y)$ are continuous and for almost all $y \in Y$ there exist $x_1, x_2 \in X$ such that $\pi(x_1, y) \leqslant 0$ and $\pi(x_2, y) > 0$.*

In other words, all sellers can make profit or losses on a match, whereas all buyers benefit. This condition allows us to focus on sellers' screening incentives. An important special case of a linear profit function is studied in Section 3.4.

Each seller seeks to maximize an average profit flow.[8] Each buyer is essentially non-

[6]Seller type $y$ captures payoff-relevant information that, for whatever reason—costly, unethical, etc.—the platform cannot elicit from the sellers. For example, on Uber, $y$ would incorporate the driver's preference for long rides and traffic; on Airbnb, $y$ would incorporate the host's preference for his guest's age, gender, socio-economic status, race, etc.

[7]Many results generalize to the case of heterogeneous buyer utility.

[8]The results immediately generalize to the case where the sellers have discount rate $r$ by replacing $\tau$ with

strategic as it does not make any decisions. We focus solely on the information design and assume the player utilities are nontransferable. We discuss the potential role of pricing schemes in Section 4.

**Information Design.** Before the matching process begins, the platform designs and commits to a stationary and non-discriminatory information policy that governs what buyer information is disclosed to sellers. Formally, the platform observes buyer type $x$ and sends a signal about it to the seller. Let $S = \Delta(X)$ be the set of all posterior distributions over $X$. An *information policy* $\lambda \in \Delta(S)$ is a probability distribution of posteriors.[9] The interpretation is that $s \in S$ is an induced posterior of a seller and $\lambda(S')$ is the fraction of buyers with signals $S' \subset S$.[10] The set of possible information policies is

$$\left\{\lambda \in \Delta(S)\colon \int s \, \mathrm{d}\lambda(s) \sim F\right\}. \tag{1}$$

When a buyer of type $x$ requests an available seller, the platform sends a signal to the seller according to $\lambda$. The seller knows the platform's choice of $\lambda$ and forms posterior beliefs via Bayes' rule. The full disclosure policy, denoted by $\lambda^{FD}$, perfectly reveals buyer type $x$ to the sellers. The no disclosure policy fully conceals $x$. We say an information policy $\lambda'$ is *coarser* than $\lambda''$ if $\lambda'$ is a Blackwell garbling of $\lambda''$.

**Steady State.** The matching process is the dynamic system in which sellers become repeatedly busy and available. Let $\alpha(y) \geqslant 0$ be the *acceptance rate*, the fraction of buyers accepted by type-$y$ sellers. Let $\rho(y)$ be the *utilization rate,* the fraction of type-$y$ sellers who are busy. Denote the average utilization rate by $\bar{\rho} \triangleq \int_Y \rho(y) dG(y)$. As the total mass of sellers is 1, $\bar{\rho}$ is also the mass of busy sellers.

The seller's strategy is a function $\sigma(\cdot, y)\colon S \to [0, 1]$ that for every seller of type $y$ maps

$\tau_r = (1 - e^{-r\tau})/r$.

[9] $\Delta(S)$ is the set of Borel probability distributions with the weak-* topology on $\Delta(X)$.

[10] Equivalently, we could model an information policy as an arbitrary statistical experiment informative about $x$.

a signal to the probability of accepting the request. The seller's acceptance rate is then:

$$\alpha(y) = \int \sigma(s, y) \mathrm{d}\lambda(s). \tag{2}$$

In a steady state, the flow of sellers who form new matches should be equal to the flow of sellers who return to waiting. The flow of new matches is equal to the product of the buyer flow that type-$y$ sellers receive $\beta\frac{(1-\rho(y))g(y)}{1-\bar{\rho}}$ and type-$y$ sellers' acceptance rate $\alpha(y)$. The flow of returning sellers is equal to $g(y)\rho(y)/\tau$, because the mass of busy type-$y$ sellers is $g(y)\rho(y)$ and matches last time $\tau$. Hence, a steady state is accounted by the following equation:

$$\beta\frac{g(y)(1-\rho(y))}{1-\bar{\rho}}\alpha(y) = \frac{g(y)\rho(y)}{\tau}, \quad \forall y \in Y. \tag{3}$$

In a steady state, buyers request an available seller at a Poisson rate denoted by $\beta_A$. On one hand, an individual available seller faces a stochastic request process such that the probability of a new request over the time interval $dt$ is $\beta_A dt + o(dt)$. On the other hand, the available sellers jointly face the deterministic arrival process of buyers, where over time interval $dt$ the mass $\beta dt$ of buyers arrives. Uniform assignment of buyers across sellers results in a particularly simple form of the relationship between the two rates (Myerson (2000)):

$$\beta_A \triangleq \frac{\beta}{1-\bar{\rho}}. \tag{4}$$

A steady-state equilibrium is a market outcome in which the sellers take the request rate $\beta_A$ as given and optimize independently, and the busy-available seller flows balance.

**Definition 1.** The pair $(\sigma, \bar{\rho})$, where $\sigma\colon S \times Y \to [0,1]$ is the sellers' strategies and $\bar{\rho} \in [0,1]$ is the mass of busy sellers, constitutes a *steady-state equilibrium* if the following hold:

1. (Steady state) The average utilization rate $\bar{\rho}$ is given by (2) and (3). The request rate $\beta_A$ is given by (4).

2. (Optimality) For all $y$ and every type-$y$ seller, $\sigma(\cdot, y)$ is an optimal strategy given the

request rate $\beta_A$ and an information policy $\lambda$.

The platform's objective is to maximize a weighted average of the buyer surplus and the joint seller profits in a steady-state equilibrium. For a given steady-state equilibrium, define by $V(y)$ the seller-$y$ profits, by $V = \int_Y V(y)\mathrm{d}G(y)$ the joint seller profits, and by $U$ the buyers' utility. The platform then formally maximizes

$$\mathcal{J}(\gamma) = \gamma U + (1-\gamma)V. \tag{5}$$

The general objective $\mathcal{J}(\gamma)$ includes as special cases maximization of total surplus at $\gamma = 1/2$, maximization of the joint seller profits at $\gamma = 0$ and maximization of buyer surplus at $\gamma = 1$. As $\gamma$ spans the interval $[0, 1]$ the platform-optimal allocations span constrained-Pareto-efficient allocations.

**Assumption Discussion.** Before proceeding with the main analysis, we highlight two important features of our model that are motivated by the stylized facts about the online platforms.

First, no two buyers simultaneously contact the same seller. Such coordination frictions in matching markets have been extensively studied in the theoretical literature (Burdett, Shi, and Wright (2001); Kircher (2009)), and digital platforms now usually have good technological means of resolving the simultaneity-driven friction. Instead, we focus on the matching friction that pertains to preference heterogeneity and screening.[11]

Second, the buyers make a single search attempt and pick the seller uniformly at random. The single search attempts can be attributed to rejection intolerance and capture an aspect of real-life matching markets in that rejections are costly to buyers (wasted time, search efforts, bidding costs).[12] It allows us to not have an endogenous distribution of buyer types.

[11] Fradkin (2015) finds that on Airbnb, seller screening causes around 41% of all failed matches, whereas simultaneity only around 12% (the rest is attributed to unavailable vacancies).

[12] Fradkin (2015) reports that on Airbnb, an initial rejection decreases by 51% the probability that the guest eventually books any listing.

The uniform assignment is plausible in settings in which the buyers are indifferent across the sellers or the search costs are prohibitively high. It allows us to analyze the match-quality and match-rate tradeoff in a cleaner model, in which any seller faces the same request rate, $\beta_A$, and the same distribution of buyer types, $F$. Altogether, this implies that the buyers are effectively non-strategic and allows us to focus on the sellers' side of the market. We discuss possible implications of buyers' strategic behavior in Section 4.

# 3 Analysis

We proceed with studying how the platform's information policy affects the equilibrium market outcome. We show that for an arbitrary information policy, an equilibrium exists and is essentially unique. We then show that full disclosure policy aggravates cream skimming and produces an inefficient market outcome. We then characterize optimal information policies in the cases of homogeneous sellers and vertically differentiated sellers.

## 3.1 Equilibrium Existence and Uniqueness

For a given information policy, each seller is facing a dynamic decision problem of maximizing an average profit flow. He is presented with a sequence of buyers arriving at a Poisson rate $\beta_A$. For each buyer request, the seller observes the signal realization $s$ and chooses whether to accept the request (see Figure 2). With a small abuse of notation, let $\pi(s, y) \triangleq \mathbb{E}\left[\pi\left(x, y\right) \mid s\right]$ be the type-$y$ seller's flow profit if he accepts a buyer with signal $s$ and let $V(y)$ be the seller's value function, an average profit flow when he acts optimally.[13] An opportunity cost of accepting is equal to $\tau V(y)$, and the seller optimization problem corresponds to the following Bellman equation

[13]For example, if a seller earns \$1 on each buyer, and the time interval between accepting a pair of consequent buyers is 2, then $V(y) = 1/2$.

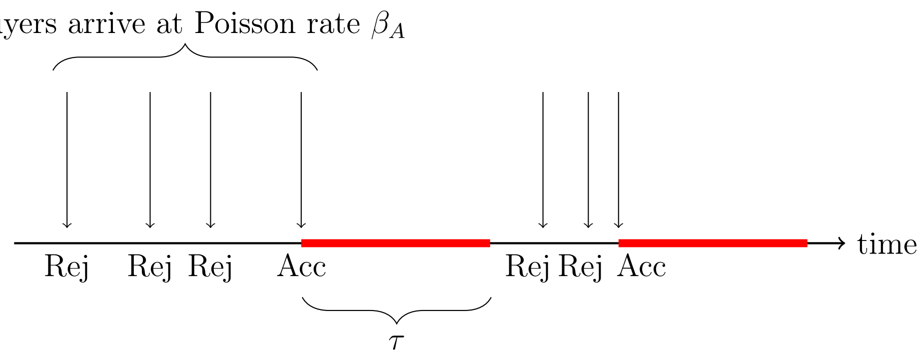


Figure 2: Seller's dynamic optimization problem with screening and waiting. An available seller receives buyer requests at a Poisson rate $\beta_A$. If a request is accepted, the seller becomes busy for time $\tau$, during which he does not receive new requests.

$$V(y) \quad = \quad \beta_A \int \max\{0, \pi(s,y) - \tau V(y)\} \mathrm{d}\lambda(s). \tag{6}$$

It is immediate that an optimal seller strategy is cutoff with him accepting all buyers with $s$ such that $\pi(s,y) > \hat{\pi}$ and rejecting all buyers with $s$ such that $\pi(s,y) < \hat{\pi}$. The optimal threshold $\hat{\pi}$ naturally depends on the information policy.

**Proposition 1.** (Existence and Uniqueness) *For an arbitrary information policy $\lambda$, a steady-state equilibrium exists and is unique up to the acceptance of marginal buyers: if $(\sigma, \bar{\rho})$ and $(\sigma', \bar{\rho}')$ are two steady-state equilibria, then (1) $\bar{\rho} = \bar{\rho}'$, (2) for any $y \in Y$, $\sigma(\cdot, y)$ and $\sigma'(\cdot, y)$ coincide except on the set $\{s \colon \pi(s,y) = \tau V(y)\}$, and (3) the surpluses coincide, $V(y) \equiv V'(y)$, $U = U'$.*

To prove this result, we show that for an arbitrary acceptance rate profile $\alpha(y)$, there is a unique steady-state value of the average utilization rate $\bar{\rho}$. The equilibrium existence and uniqueness roughly follows from the continuity and monotonicity of the reaction curves of $\alpha$ in $\bar{\rho}$ and $\bar{\rho}$ in $\alpha$. On one hand, if the average utilization rate $\bar{\rho}$ increases, then the request traffic to each available seller increases. As a result, sellers become pickier, threshold $\hat{\pi}$ increases, and the acceptance rate $\alpha(y)$ decreases. On the other hand, if $\alpha(y)$ increases, sellers become less available, and $\bar{\rho}$ strictly decreases. The equilibrium values of $\alpha$ and $\bar{\rho}$ then determine the sellers' strategies up to the acceptance of marginal buyers and, in turn, the sellers' surpluses. The buyer surplus is uniquely determined by the utilization rate as well.

## 3.2 Market Failure Under Full Disclosure

In this subsection, we establish that full disclosure leads to a Pareto inefficient market outcome. The market failure is driven by the sellers' decentralized decision-making and their capacity constraints.

We first define Pareto efficiency in our setting. An outcome $O = (V(\cdot), U)$ is a combination of seller profit profile and buyer surplus. We say that a market outcome is *feasible* if there is a seller strategy profile under full disclosure that achieves it. A feasible outcome $O$ is Pareto efficient if there is no other feasible $O'$ such that $V'(y) \geqslant V(y)$ for all $y$ and $U' \geqslant U$, and at least one seller type or buyers are strictly better off. The Pareto frontier is the set of all Pareto-efficient outcomes. An outcome $O$ is *implementable* if there is an information policy that induces it in a steady state equilibrium.

Imagine that the platform starts with full disclosure as its default information policy. Write $\alpha^{FD}(y), V^{FD}(y), U^{FD}$ for equilibrium acceptance rates, seller profits, and consumer surplus. Correspondingly, write $\alpha^{\sigma}(y)$, $V^{\sigma}(y)$, $U^{\sigma}$ for these parameters when a given strategy profile $\sigma$ is played.

**Proposition 2.** (Market Failure) *There exists a strategy profile $\sigma$ under which all market participants match more often and are strictly better off than in a full-disclosure equilibrium. That is for all $y \in Y$:*

$$\begin{aligned} \alpha^{\sigma}(y) &> \alpha^{FD}(y), \\ V^{\sigma}(y) &> V^{FD}(y), \\ U^{\sigma} &> U^{FD}. \end{aligned}$$

A full proof is given in the Appendix. The basic intuition for the coordination problem is *cream skimming*: A seller keeps his schedule open by rejecting low-value matches in order to increase his own chances of getting high-value matches. As a result, in equilibrium, sellers spend considerable time waiting for high-value matches. Collectively, this behavior is suboptimal because some low-value matches have to be accepted to maximize the joint seller surplus. The market feature that gives rise to the coordination failure is that collectively, the sellers are not capacity constrained, whereas individually, they *are* capacity constrained. Notably, unlike price competition that benefits buyers and improves market efficiency, seller competition for better buyers hurts buyers and decreases market efficiency.

It is instructive to highlight the inefficiencies on the seller and the buyer sides of the market as outlined in the introduction. Each seller rejects low-value matches to increase his chances of getting high-value matches. Such rejections directly harm the buyers because the sellers do not internalize their surplus. Such rejections hurt other sellers as well, because the seller remains available on the marketplace and competes with other sellers for incoming buyers.

Can information design mitigate the market failure? To answer this question, it is important to understand competing effects of information provision. Define the match rate $m$ as the number of matches formed on the platform over a unit time interval. Higher $m$ implies that more buyers are served, and the buyer surplus is proportional to it, $U = mu$. Similarly, define the seller average match quality by $q$ so that $V = mq$. Evidently, both surpluses increase in the match rate $m$. However, the match rate and the match quality are in conflict, and the information disclosure affects them through three different channels.

First, information provision has a positive effect on the seller match quality. From an individual seller's point of view, more information increases his set of attainable payoffs. Holding the match rate fixed, he individually benefits from more information about buyers. Second, information provision reduces the match rate by allowing the sellers to reject less valuable buyers. The platform can limit information and induce sellers to accept more buyers. This effect is reminiscent of static persuasion settings. Third, information provision reduces the match rate through increasing sellers' return on search. With returns on search, the sellers reject more often, and the match rate goes down.

The match quality effect motivates the platform to provide more information, while the effects on match rate work in the opposite direction. The interaction of these effects depends on the primitives of the economic environment and shapes the preferred information policy. In the next section, we show that with homogeneous sellers, information design can fully mitigate the market failure and achieve an efficient allocation.

## 3.3 Optimal Policy with Homogeneous Sellers

We start with the case of homogeneous sellers so that $Y$ is a singleton or, equivalently:

$$\pi(x,y) = \pi\left(x\right). \tag{7}$$

It is instructive to think about what seller strategies are Pareto efficient in this environment. It is straightforward to see that any efficient strategy is cutoff—all sellers should accept buyers that generates profits above some threshold. Intuitively, any given acceptance rate $\alpha$ results in the same buyer utility and a cutoff strategy maximizes the seller surplus among all strategies with the same acceptance rate. Conversely, a cutoff monotonically affects the acceptance rate and hence corresponds to some efficient outcome. We show that information control allows the platform to restore efficiency and bring the payoffs to the Pareto frontier.

**Theorem 1.** (Homogeneous Sellers) *Suppose the sellers are homogeneous. Then, for any Pareto-efficient outcome* $(U,V)$ *with* $V \geqslant 0$*, there is an information policy that sends only two signals that implement it.*

The proof follows the efficiency approach of Smolin (2017).[14] Consider an arbitrary seller

[14]The efficiency result also relates to that of Bergemann, Brooks, and Morris (2015). They show that segmentation of a monopolistic market can achieve every feasible combination of consumer and producer surplus

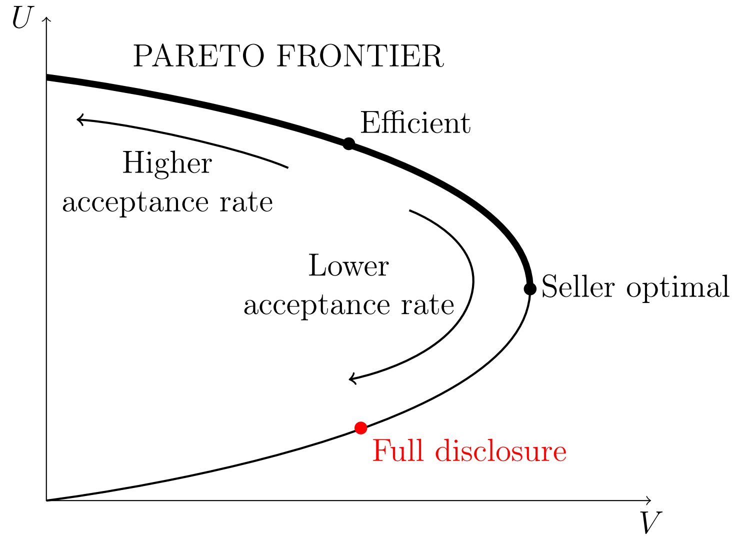


Figure 3: A sketch of a set of feasible and implementable outcomes in the case of homogeneous sellers. $V$ denotes the seller surplus; $U$ denotes buyer surplus. Any point on the Pareto frontier (thick solid line) can be implemented by some information policy. The full-disclosure outcome is suboptimal.

strategy that under full disclosure achieves a Pareto-efficient outcome with positive seller surplus. Consider an information policy that sends two signals and effectively recommends an action according to the strategy. Under the recommendation policy, there is only one signal with a recommendation to accept, so there is only one type of acceptable buyer. Hence, all profitable buyers are the same, and there is no return to cream skimming. For any Pareto-efficient outcome that delivers a positive surplus to sellers, they are willing to obey the recommendations.

Theorem 1 characterizes at once the range of optimal information policies corresponding to different platform objectives as any point on the Pareto frontier maximizes $\gamma U + (1-\gamma)V$ for some $\gamma \in [0, 1]$. Figure 3 illustrates the result. Note that by Proposition 2, the full-disclosure equilibrium outcome is outside of the Pareto frontier.

Now, note that in the case of homogeneous sellers, Pareto-efficient allocations do not depend on exact parameters $\beta$ and $\tau$. We have an immediate corollary.

**Corollary 1.** *Suppose the sellers are homogeneous. The optimal information policy $\lambda^*$ depends only on the objective weights $\gamma$ and not on the buyer traffic intensity $\beta$ or the task* *length $\tau$.*

This corollary implies that the same information policy would be optimal if sellers became available immediately after accepting a buyer, $\tau = 0$. The independence from $\beta$ and $\tau$ is driven by the coarse structure of an optimal policy. The rate of arrival to the available sellers, $\beta_A$, matters to sellers only to the extent that a higher request rate magnifies the opportunity cost of accepting. With a binary policy, this opportunity cost is nil because all profitable buyers are the same.

In the next section, we study the setting with heterogeneous sellers and show that the optimal policy is richer and does depend on $\beta$ and $\tau$.

## 3.4 Optimal Policy with Heterogeneous Sellers

We proceed with studying the setting in which the sellers feature vertical payoff heterogeneity. We characterize the optimal information policy and show that, unlike in the case of homogeneous sellers, the Pareto frontier is generally not attainable.

Formally, consider the *linear payoff environment* in which the seller profit is a linear function of buyer and seller types: $X = [0, \overline{x}], Y = [0, \overline{y}]$, and

$$\pi(x, y) = y - x. \tag{8}$$

We interpret $x$ as the level of difficulty of the buyer's task and interpret $y$ as the seller's skill level. In this way, buyers are vertically differentiated in the eyes of the sellers, and high-type sellers can profitably match with more buyers than can low-type sellers. Note that by the profit variability assumption, $\overline{y} \leqslant \overline{x}$.

Similar to the case of a homogeneous sellers, it is easy to see that the Pareto efficient strategies are cutoff—all seller types should accept buyers below a threshold which can now be type-dependent. At the same time, choosing the cutoff profile to maximize the total seller surplus should account for the congestion externalities due to uniform assignment. In general, this profile depends on the details of seller and buyer type distributions. However, it is clear that the efficient cutoffs are generically different from their types and vary across them.

Finding an optimal information policy in this environment is a challenging problem. First, the sellers' preferences have unobserved heterogeneity and, as discussed above, Pareto efficient strategies differ across sellers. This precludes the effective use of the recommendation policies that provided a solution in the case of homogeneous sellers. Second, the sellers are forward looking, and the information policy determines not only their stage payoff but also their continuation values. As a result, a seller's decision depends not only on the posterior mean of the current buyer type but on the general shape of the information policy.

To solve the problem, we develop a variational approach in which we map information policies to a class of convex functions and then use the calculus of variation to recover a structure of optimal information policies. First, we follow Gentzkow and Kamenica (2016) and show that optimizing $\mathcal{J}$ with respect to information policies is equivalent to optimizing it with respect to a class of convex functions. Namely, for a given information policy $\lambda$, denote the posterior mean of $x$ conditional on signal $s$ by $z(s) \triangleq \int_X x \, \mathrm{d}s(x)$, the induced distribution of $z(s)$ by $H$, and define the *spread function* $\Lambda\colon [0,\infty) \to \mathbb{R}_+$ as

$$\Lambda(z) \triangleq \int_0^z H(\zeta) \, \mathrm{d}\zeta. \tag{9}$$

In the linear payoff environment, $\Lambda$ is proportional to the option value of rejecting a buyer with expected cost $z$. Let $\overline{\Lambda}$ be the spread function under full disclosure, $\overline{\Lambda}(z) = \int_0^z F(\zeta) \, \mathrm{d}\zeta$, and let $\underline{\Lambda}$ be the spread function under no disclosure, $\underline{\Lambda}(z) = \max\{0, z - \mathbb{E}[x]\}$. The argument of Gentzkow and Kamenica (2016) establishes a one-to-one correspondence between the set of admissible posterior mean distributions and the following set,

$$\mathcal{L} \triangleq \left\{\Lambda\colon [0,\infty) \to \mathbb{R}_+, \text{ convex, and } \underline{\Lambda} \leqslant \Lambda \leqslant \overline{\Lambda}\right\}. \tag{10}$$

As a result, maximization over information policies is equivalent to maximization over admissible spread functions $\Lambda \in \mathcal{L}$.

We then build on the analysis of Section 3.1 and show that for any information policy, the seller's optimal strategy has a simple cutoff structure: a type-$y$ seller accepts all buyers with expected difficulty below the cutoff $\hat{z}(y)$ and rejects all buyers above it. The cutoff profile has a simple structure—it is continuously increasing in $y$ and satisfies the following equation

$$\frac{y - \hat{z}(y)}{\tau \beta_A(\Lambda)} = \Lambda(\hat{z}(y)). \tag{11}$$

The $\beta_A(\Lambda)$ is an equilibrium request rate that is linked to the seller behavior by (3) and (4).

The condition (11) allows easy illustration of cream skimming in Figure 4. Under cream skimming, a type-$y$ seller accepts buyers with expected difficulty $z$ not below his type $y$ but below a cutoff $\hat{z}(y) < y$. The cutoff can be seen as a projection of the type reflected against the spread function along the trajectory with the slope inversely proportional to the request rate $\beta_A$ and the task length $\tau$. The cutoff is, hence, increasing with type—higher seller types accept more buyers, cream skimming less.[15] The higher the request rate, the flatter the slope and the greater the extent of cream skimming. Similarly, the lower the task length, the steeper the slope and the lower the extent of cream skimming. In the limit case as $\tau \to 0$, the slope is vertical, there is no cream skimming, $\hat{z}(y) = y$, and the problem can be seen as a static persuasion of Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017).

Next, to solve for an optimal spread function, we refer to the calculus of variations and consider effects of small spread-function variations on the platform's objective. Namely, we consider variations $\varepsilon h : [0, \infty) \to \mathbb{R}_+$ such that $\Lambda + \varepsilon h \in \mathcal{L}$. We use the equilibrium conditions and calculate a variational derivative of the platform objective at $\varepsilon = 0$:

$$\frac{d\mathcal{J}}{d\varepsilon}(0) = \int_0^{\overline{y}} h(\hat{z}(y))\, \nu(y) \left( -\frac{g'(y)}{g(y)} \Psi + (1-\gamma) \frac{\beta}{\overline{\nu}} \right) g(y)\, \mathrm{d}y + h(\hat{z}(\overline{y}))\, \nu(\overline{y})\, g(\overline{y})\, \Psi. \tag{12}$$

[15] Thus, the equilibrium exhibits a kind of negative stochastic assorting in terms of Chade (2006).

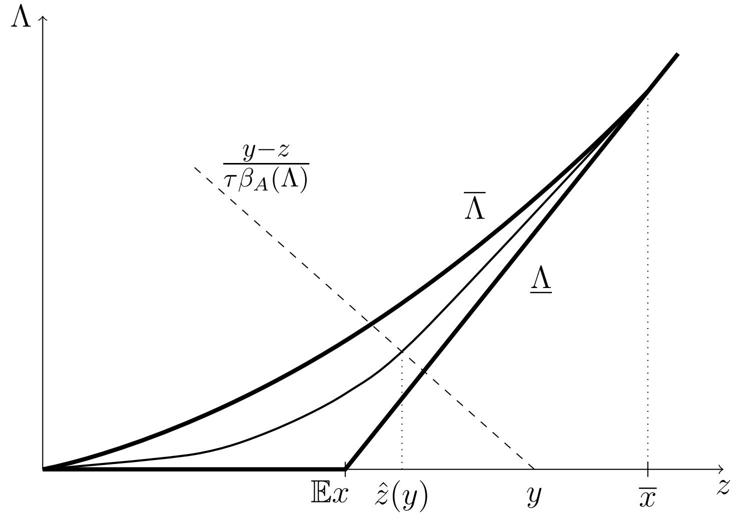


Figure 4: Admissible spread functions $\Lambda$ and equilibrium acceptance cutoffs $\hat{z}\,(y)$.

Here, $\nu\,(y) \triangleq 1 - \rho\,(y)$, $\bar{\nu} \triangleq 1 - \bar{\rho}$ are the equilibrium availability rates under the policy $\Lambda$, and $\Psi > 0$ is a positive term that depends on the objective weight $\gamma$ and the aggregate equilibrium behavior. The integral term captures the relative distribution of the platform's surplus across the seller types, whereas the stand-alone term captures its aggregate level.

The behavior of the integrand expression shows the direction in which $\mathcal{J}$ can be locally improved. If $g$ is decreasing, then the integrand is positive everywhere. Any positive admissible variation increases $\mathcal{J}$, pushing the spread function towards full disclosure. If $g$ is log-concave, then the integrand crosses zero from below at most once.[16] The same thing happens if $g$ is unimodal and the platform maximizes the matching rate, $\gamma = 1$. In these cases, an improving deviation pushes towards more disclosure in the upper tail of $x$ and less disclosure in the lower tail. We show that these deviations lead to censorship policies.

**Definition.** An information policy $\lambda$ is a *lower-censorship policy* if for some $\hat{x} \in [0,1]$, $\lambda$ fully reveals all $x > \hat{x}$ and pools all $x < \hat{x}$.

**Theorem 2.** (Heterogeneous Sellers) *In the linear payoff environment:*

1. *If the seller distribution $g$ is log-concave, then a lower-censorship policy is optimal;*

[16] The class of log-concave distributions includes normal, logistic, exponential, and uniform distributions, as well as their truncations.

2. *If the seller distribution g is weakly decreasing, then a full disclosure policy is optimal;*

3. *If the platform maximizes only the matching rate, $\gamma = 1$, and the seller distribution g is unimodal then a lower-censorship policy is optimal.*

Optimal disclosure depends on the shape of seller type distribution because the optimal policy should satisfy incentive constraints for the sellers. Suppose there are two types of sellers, professionals and amateurs, and two types of tasks, easy and hard. Professionals accrue relatively higher value from tasks than do amateurs; but assume that under full disclosure, both professionals and amateurs accept only easy tasks. Limiting information may cause opposite reactions from them. Consider an information policy that pools easy tasks with hard tasks. A professional's average profit from the pooled task is positive, and so he starts accepting hard tasks. Conversely, an amateur's average profit from the pooled task is negative, and so he starts rejecting even easy tasks. Because the platform does not observe the seller type, the optimal information policy should strike a balance between professionals and amateurs and thus depends on their relative population sizes. If there are sufficiently many professionals, then no disclosure is optimal; if there are sufficiently many amateurs, then full disclosure is optimal.

Theorem 2 extends this binary intuition to the case of many seller types. It shows that for a broad class of type distributions, lower-censorship policies are optimal. These policies pool attractive buyers with infra-marginal buyers and fully reveal the less attractive buyers. They effectively provide no useful information for low-type sellers but bring a lot of informational value to high-type sellers. Under these policies, the worst sellers are excluded from the platform, the medium sellers accept only the pooled buyers, and the best sellers screen among the less attractive buyers. In this way, the platform optimally trades off the buyer and seller surpluses.

# 4 Applications

Strategic information coarsening is relevant in practice and can be found in the designs of many platforms. In this section, we discuss how our model relates to ride-hailing, accommodation, and freelance agencies, and how our analysis can guide their design.

- *Ride-Hailing Platforms.* Sellers and buyers are, respectively, drivers and passengers. When idle, drivers receive requests from passengers. The driver type captures his home location, preference for long rides, and tolerance for congestion. The rider type captures the destination and ride history. Drivers do not like short rides or rides to remote neighborhoods.[17] Passengers do not like waiting.[18] The leading agency Uber applies fixed pricing per mile and per minute that can depend only on aggregate multipliers, such as surge in demand.[19] The passenger destination is fully concealed and represents the platform information coarsening.[20]

- *Accommodation Rental Platforms.* Sellers and buyers are, respectively, hosts and guests. Hosts are capacity constrained: once a room is booked, the host cannot accept a better guest. The host type captures his preference for age, race, personality, and daily schedule. The guest type captures his gender, age, socio-economic status, and lifestyle. In the leading agency Airbnb, every host sets a price that applies for all guests, but he may prefer to reject a guest whom he expects to be a bad fit. The host can adopt the "Instant Book" feature, with which the host commits to accepting all guests filtered by a few characteristics (e.g., pets, smoking, infants). The choice of filters acts as an implicit information coarsening.

[17] http://www.forbes.com/sites/harrycampbell/2015/03/24/just-how-far-is-your-uber-driver-willing-to-take-you.

[18] Cohen, Hahn, Hall, Levitt, and Metcalfe (2016) estimate that an increase in wait time of 1 minute decreases the probability of requesting a ride by 1.7 percent.

[19] The surge price coefficient depends on the numbers of riders and nearby drivers who have their Uber app open at that moment: https://newsroom.uber.com/upfront-fares-no-math-and-no-surprises/.

[20] Uber drivers can use the destination filter twice a day. If used, the platform filters passenger requests and offers only those aligned with the driver's preferred destination: http://www.ridesharingdriver.com/uber-destination-filter-get-passengers-heading-your-direction/.

- *Freelance Labor Platforms.* Sellers and buyers are, respectively, service providers and clients. The service providers are constrained in the number of tasks they can do per week. The provider's type captures his skills and work ethic. The client type captures the task category and difficulty, and the client's expertise and location. In the leading agency TaskRabbit, the service providers set an hourly rate that applies to all tasks in the same category. Pulling many tasks into a single category is a form of information coarsening.

Our results shed light on why limiting buyer information can improve performance of these platforms. In an example of ride-hailing platforms, if drivers have too detailed information about passenger requests, they cream skim. This results in low acceptance rates and a suboptimal demand–supply fit. The platform's policy of hiding the passenger destination from drivers can be used to force drivers to accept more rides and avoid the inefficient cream-skimming behavior.

At the same time, we show that providing some information is often beneficial. In particular, we argue that censorship policies that reveal the least valuable buyers and pool the most valuable buyers with intermediate ones can be particularly useful. In the case of ride-hailing platforms, a censorship policy can be implemented as follows. The best driver destinations are pooled with intermediate destinations. The worst destinations, like those to remote places, are flagged and accompanied with a precise description. In equilibrium, the pooled destinations are served by all drivers active on the platform. The flagged destinations are served only by the most willing drivers, depending on the description details. This policy optimally balances between the platform's objectives—it provides sufficient information and surplus to the drivers but insures a high matching rate for the riders.

Before concluding, we discuss two considerations we abstracted away from but that might be important and should be taken into account in some applications: strategic buyers and optimal pricing.

**Strategic Buyers.** If buyers have heterogeneous match quality and search for better matches, then the platform might face the information design possibilities on both sides of the market. Disclosure to buyers can play several roles, depending on their strategic options.

First, the platform can use the disclosure to direct buyers to the sellers with higher acceptance rates. Imagine buyers observe seller acceptance rates imperfectly. Then they request selective sellers too often compared to what would be efficient. The platform can pool selective and non-selective sellers to shift buyer flow towards non-selective sellers.

Second, the platform can use information disclosure to direct buyers to the sellers who value them the most. Imagine that the prices do not reflect the seller match payoffs perfectly. Then, the information policy that pools matches that are less valuable for buyers and more valuable for sellers with those that are, conversely, more valuable for buyers and less valuable for sellers, will improve overall efficiency.[21] In particular, one consider a natural extension of our model in which the buyers have differential tastes over sellers. Our analysis can inform this setting as well. If both sides of the market are homogeneous, then the argument analogous to that of Section Section 3.3, can be used to establish that the platform can again fully restore efficiency. If the sides are heterogeneous, then the analysis is more complicated but we expect the variational approach still being useful, even if more tedious. Moreover, we conjecture that the lower-censorship policies would continue to play an important role in analysis.

**Optimal Pricing.** Throughout the paper, we assumed the participants' payoffs are non-transferable. In practice, a platform may want to complement the information design with the pricing design. Such complementation is straightforward if the pricing is uniform and takes the form of a fixed amount a buyer should pay to a seller if he is accepted. In this case,

[21] Rayo and Segal (2010) study the last situation in a static match model when both buyers and sellers derive heterogeneous payoffs from matches. They find that in the case of the uniform distribution of receiver reservation values, the optimal information policy pools "non-ordered prospects," where two prospects are non-ordered in the sense described above.

the platform should simply perform backward induction optimization. For any given price, the platform can calculate the matching payoffs and find the optimal information policy and its corresponding payoffs by using our analysis. In the first step, the platform can then choose the price optimal.

In fact, if the sellers are homogeneous, then the uniform pricing is sufficient to implement an efficient outcome that maximizes joint surplus. Indeed, Theorem 1 establishes that any Pareto-efficient outcome can be implemented by some information policy as long as the sellers receive some rents. If the efficient outcome delivers positive rents to the sellers, then no additional price is required. Otherwise, the platform can set the price to subsidize the sellers and give them sufficient rent to participate in the platform. Then, by the same argument as in the main analysis, there exists a recommendation policy that implements the outcome.

If the sellers are heterogeneous, then Theorem 2 shows that the optimal information policy need not even achieve Pareto frontier, so that flexible pricing rules can possibly outperform the uniform pricing. However, most matching platforms adopt simple pricing. There could be several explanations. First, consumers appreciate transparent and simple pricing, both because it is easier to assess the cost of a ride and because complex pricing can feel like price gouging and prompt the risk of price-discrimination litigation. Second, flexible pricing is a non-trivial development task that requires significant resources. Relatedly, it may be impractical to condition on every variable relevant in the transportation market. There is some complexity threshold for the designer beyond which increasing the price complexity is not justified.[22]

As a final remark, we note that often the platforms accompany uniform pricing with a commission proportional to the price of each match. In these cases, maximizing a profit flow is equivalent to maximizing a matching rate. Our results then point out that censorship policies can be useful to bring the most revenue to the platform.

[22] See Segal (2007) for the study of communication requirements of social choice rules. It demonstrates that in some social choice problems, the space of prices that must be discovered is prohibitively large.

# 5 Conclusion

Strategic information disclosure is arguably less intrusive than direct control—and often more palatable than sophisticated pricing schemes. When the two sides of the market are not symmetric in their preferences for match quality and match rate, the more patient and selective side of the market tends to cream skim. In this paper, we provide a case for using information design as a way to mitigate the adverse effects of cream skimming and improve platform performance. Specifically, we highlight the use of censorship policies as a way to balance surpluses on both sides of the market.

For most of this paper, we isolated information design from other tools available to platforms. Flexible pricing, real-time auctions, recommender systems, and queue management can also help improve matching market performance. We believe that studying interaction between these schemes can complement our current insights and constitutes a promising avenue for further research.

# 6 Appendix

**Lemma 1.** *For any information policy $\lambda$ and an equilibrium acceptance rate profile $\alpha(y)$:*

- *the average utilization rate $\bar{\rho} \in [0,1]$ is a solution to*

$$1 = \int \frac{1}{1-\bar{\rho}+\beta\tau\alpha(y)} \mathrm{d}G(y); \tag{13}$$

- *the solution $\bar{\rho}$ exists and is unique; it increases in $\beta$, $\tau$, and whenever $\alpha(y)$ is increased on a set of strictly positive measure with respect to $G$;*

- *the utilization rate $\rho(y)$ increases whenever $\alpha(y')$ is increased on a set of strictly positive measure with respect to $G$.*

*Proof.* By (3):

$$1 - \rho(y) = \frac{1}{1 + \tau\beta\alpha(y)/(1-\bar{\rho})}. \tag{14}$$

Integrate both sides with respect to $G(y)$ and rearrange to obtain (13). Its right-hand side is continuously and strictly increasing in $\bar{\rho}$. Evaluated at $\bar{\rho} = 0$, it equals $\int \frac{dG(y)}{1+\beta\tau\alpha(y)} \leqslant 1$; evaluated at $\bar{\rho} = 1$, it equals $\int \frac{dG(y)}{\beta\tau\alpha(y)} \geqslant 1$, by the aggregate capacity assumption. Therefore, a solution exists and is unique. The monotonicity of $\bar{\rho}$ follows. Finally, consider a perturbation of increasing $\alpha$ on some positive measure of $y' \neq y$. By the argument above, $\bar{\rho}$ increases so, by (14), $\rho(y)$ increases. □

**Lemma 2.** *For any given information policy $\lambda$ and utilization rate $\bar{\rho}$, the sellers optimal strategies are cutoff:*

$$\sigma(s,y) = \begin{cases} 1, & \pi(s,y) > \hat{\pi}(y), \\ 0, & \pi(s,y) < \hat{\pi}(y). \end{cases}$$

*The cutoff function $\hat{\pi}(y)$ is unique and for all types $\hat{\pi}(y) > 0$ and $\alpha(y) < 1$.*

*Proof.* Fix $y \in Y$. By (6), the optimal seller strategy is such that all signals from $S_a(y) \triangleq \{s\colon \pi(s,y) > \tau V(y)\}$ are accepted and all signals from $S_r(y) \triangleq \{s\colon \pi(s,y) < \tau V(y)\}$ are rejected. The seller is indifferent between accepting and rejecting signals from $S_m(y) \triangleq \{s\colon \pi(s,y) = \tau V(y)\}$. The cutoff for a type-$y$ seller is $\hat{\pi}(y) = \tau V(y)$. Using (6), the cutoff of the optimal strategy is a solution to

$$\hat{\pi}(y) = \tau\beta_A \int \max\{0, \pi(s,y) - \hat{\pi}(y)\}\,\lambda(\mathrm{d}s). \tag{15}$$

The solution is unique because the left-hand side of (15) is strictly increasing in $\hat{\pi}(y)$ while the right-hand side is decreasing in $\hat{\pi}(y)$. By the variability assumption, $\hat{\pi}(y) > 0$. Finally, under the full disclosure policy, $\pi$ is continuously distributed, so $\alpha(y) < 1$. □

**Proof of Proposition 1.** Step 1. Existence. Consider the correspondence $\psi\colon [0,1] \rightrightarrows [0,1]$ which maps $\bar{\rho}$ to a set of "reaction" $\bar{\rho}$'s by the following procedure. First, find the unique cutoff function $\hat{\pi}$ from (15) with $\beta_A = \beta/(1-\bar{\rho})$. The cutoff $\hat{\pi}$ does not pin down the acceptance rate profile $\alpha(y)$ uniquely because marginal signals can have positive probability under $\lambda$. The acceptance rates consistent with $\hat{\pi}(y)$ are integrable functions such that

$$\lambda(S_a(y)) \leqslant \alpha(y) \leqslant \lambda(S_a(y)) + \lambda(S_m(y)), \quad \forall y \in Y.$$

Denote by $\mathcal{A}$ the set of $\alpha$ consistent with $\hat{\pi}$. For any $\alpha \in \mathcal{A}$, find $\bar{\rho}$ as shown in Lemma 1. Going over all $\mathcal{A}$ will produce the set $\psi(\bar{\rho})$.

By construction, $\mathcal{A}$ is convex. By Lemma 1, $\bar{\rho}$ is continuous in $\alpha$, so $\psi(\bar{\rho})$ is a closed interval in $[0,1]$. Moreover, $\psi$ is upper hemicontinuous because $\hat{\pi}$ is continuous in $\beta_A$ according to (15). By Kakutani's theorem, $\psi$ has a fixed point, which is a solution.

Step 2. Uniqueness. Towards a contradiction, suppose $\bar{\rho}$ and $\bar{\rho}'$, $\bar{\rho} > \bar{\rho}'$ are two distinct fixed points of $\psi$. By uniformity of assignment, $\beta_A > \beta_A'$. By (15) and the variability assumption, $\hat{\pi}(y) > \hat{\pi}'(y)$ for any $y \in Y$. Thus, whenever $\pi(s,y) \geqslant \hat{\pi}(y)$ we also have $\pi(s,y) > \hat{\pi}'(y)$. But this means that $S_a'(y) \supseteq S_a(y) \cup S_m(y)$. Therefore, $\alpha'(y) \geqslant \lambda(S_a') \geqslant \lambda(S_a \cup S_m) \geqslant \alpha(y)$ for all $y$. By Lemma 1, this implies $\bar{\rho}' \geqslant \bar{\rho}$, a contradiction. Moreover, since $\bar{\rho}' = \bar{\rho}$, $U = U'$, and by Lemma 2, the strategy cutoffs coincide $\hat{\pi}'(\cdot) = \hat{\pi}(\cdot)$. Hence, the seller surpluses coincide as well $V(\cdot) = V'(\cdot)$.

**Proof of Proposition 2.** Rearranging (6),

$$\begin{aligned} V(y) &= \frac{\beta_A}{1+\tau\beta_A\alpha(y)} \int \pi(s,y)\sigma(s,y)\mathrm{d}\lambda(s) = \\ &= \frac{\beta}{1-\bar{\rho}+\beta\tau\alpha(y)} \int \pi(s,y)\sigma(s,y)\mathrm{d}\lambda(s). \end{aligned} \tag{16}$$

By Lemma 2, strategy $\sigma^{FD}$ prescribes jobs in $\{x\colon \pi(x,y) > \tau V(y)\}$ to be accepted with probability 1 and induces acceptance rates $\alpha^{FD}(y) < 1$. Consider a deviation $\sigma$ that induces uniformly higher acceptance rates $\alpha(y) = \alpha^{FD}(y) + \Delta$ for some small $\Delta > 0$ for all $y \in Y$. As $Y$ is compact, such deviation exists and strictly increases the buyer surplus, $U > U^{FD}$.

Moreover, by (13), $1-\bar{\rho}+\beta\tau\alpha(y)=1-\bar{\rho}^{FD}+\beta\tau\alpha^{FD}(y)$ for all $y$. Yet, the integral strictly increases, because $\hat{\pi}(y)>0$, and so $V(y)>V^{FD}(y)$.

**Proof of Theorem 1.** Take any Pareto-efficient pair $O=(V,U)$. Since $O$ is feasible, there is a seller strategy $\sigma$ that induces $O$. Consider an information policy $\hat{\lambda}$ that recommends this strategy with $s_a$ being the recommendation to accept and $s_r$ the recommendation to reject. We need to check that the sellers would follow the recommendations.

From (6) we have $v(s_a)=\pi(s_a)-\tau V$, $v(s_r)=0$, and $V=\beta_A(\pi(s_a)-\tau V)\hat{\lambda}(s_a)$. The incentive constraints require that $\pi(s_a)\geqslant\tau V$ and $\pi(s_r)\leqslant\tau V$. For the former,

$$\tau V=\frac{\tau\beta_A\hat{\lambda}(s_a)}{1+\tau\beta_A\hat{\lambda}(s_a)}\pi(s_a)\leqslant\pi(s_a).$$

For the latter, recall that $O$ is Pareto efficient, hence $\sigma$ accepts all profitable jobs. This implies that

$$\pi(s_r)\leqslant 0\leqslant\tau V.$$

**Proof of Theorem 2.** The sequence of lemmas that follow support and formalize the argument outlined in the main text.

**Lemma 3.** *$\ell\in\mathcal{L}$ if and only if there is $\lambda\in\Delta(S)$ such that $\Lambda(\cdot,\lambda)=\ell$.*

*Proof.* The proof is standard and omitted. Kolotilin et al. (2017) and Gentzkow and Kamenica (2016) describe it in details. □

**Lemma 4.** *In the linear payoff environment, for any information policy $\lambda$, the seller's optimal strategy has a cutoff form such that a type-$y$ seller accepts all buyers with expected difficulty $z<\hat{z}(y)$ and rejects all buyers with $z>\hat{z}(y)$. For a given $\beta_A$, the seller payoff $V(y)$ and the cutoff $\hat{z}(y)$ satisfy:*

$$V(y)=\frac{y-\hat{z}(y)}{\tau}=\beta_A\Lambda\left(\hat{z}\left(y\right)\right). \tag{17}$$

*As a result, $\hat{z}(y)$ is increasing in $y$ and decreasing in $\tau\beta_A$.*

*Proof.* The first steps of the proof of Proposition 1 show that in the case of arbitrary $X$ and $Y$, the seller optimal strategy has a cutoff form. They also show that the minimal acceptable profit is $\hat{\pi}(y) = \tau V(y)$ and, moreover, (15) holds. In the linear payoff environment, $\pi(s, y) = y - z(s)$ and $\hat{\pi}(y) = y - \hat{z}(y)$ for some $\hat{z}$. Plug it back into (15) and obtain

$$y - \hat{z}(y) = \tau\beta_A \int_0^{\hat{z}(y)} (\hat{z}(y) - z)\mathrm{d}H(z) = \tau\beta_A \int_0^{\hat{z}(y)} H(z)dz = \tau\beta_A\Lambda(\hat{z}(y)),$$

where the second equality follows by the integration by parts. $\square$

**Lemma 5.** *In the linear payoff environment, for any information policy $\lambda$ and for almost all $y$,*

$$\alpha(y) = H\left(\hat{z}\left(y\right)\right) = \Lambda'(\hat{z}(y)).$$

*The acceptance rate $\alpha(y)$ is increasing in $y$ and decreasing in $\tau\beta_A$ for any $y \in Y$. Moreover,*

$$\hat{z}'(y) = \frac{1}{1 + \tau\beta_A\alpha\left(y\right)} = \nu\left(y\right).$$

*Proof.* By Lemma 4, the accepted jobs are those with $z < \hat{z}(y)$. Using the definition of $\Lambda$ from (9), the probability of accepting $\alpha(y)$ must lie between $H(\hat{z}(y)-)$ and $H(\hat{z}(y)+)$. The points of discontinuity of $H$ constitute a set of measure zero and seller types have full support on $[0, 1]$. It follows that $\alpha(y) = \Lambda'(\hat{z}(y))$ for almost all types $y$.

By (17), $\hat{z}(y)$ is increasing in $y$. Since $\Lambda$ is convex, $\alpha(y)$ is increasing in $y$. Same argument reveals that $\alpha(y)$ is decreasing in $\tau\beta_A$.

Finally, using Lemma 4, we have $y - \hat{z}(y) = \tau\beta_A\Lambda(\hat{z}(y))$. Differentiating that expression on both sides with respect to $y$ and substituting $\Lambda'(\hat{z}(y))$ with $\alpha(y)$, the result follows. $\square$

Consider an initial information policy $\Lambda_0 \in \mathcal{L}$, and a deviation $h : [0, 1] \to [0, 1]$. A deviation is *admissible* if $\Lambda_0 + h \in \mathcal{L}$. Since $\mathcal{L}$ is convex, $\Lambda_0 + \varepsilon h \in \mathcal{L}$ for all $\varepsilon \in [0, 1]$. For a functional $\mathcal{I} \colon \mathcal{L} \to \mathbb{R}$ we will denote the value of $\mathcal{I}$ under information policy $\Lambda_0 + \varepsilon h$ by

$\mathcal{I}(\varepsilon)$. The main step in the analysis is to find the variational derivative with respect to $\varepsilon$ at zero:

$$\mathcal{J}'(0) = \gamma U'(0) + (1-\gamma)V'(0).$$

**Lemma 6.** *In the linear payoff environment, for any admissible deviation $h$:*

$$U'(0) = -u\beta\Psi_1\left(\int_0^{\overline{y}} h(\hat{z}(y))\,\nu(y)\,g'(y)\,\mathrm{d}y - h(\hat{z}(\overline{y}))\,\nu(\overline{y})\,g(\overline{y})\right),$$
$$V'(0) = \Psi_2 U'(0) + \frac{\beta}{\overline{\nu}}\int h(\hat{z}(y))\,\nu(y)\,g(y)\,\mathrm{d}y,$$

*where $\Psi_1 = 1/\int_0^{\overline{y}}\left(\nu^2(y) - \tau\nu'(y)V(y)\right)g(y)\,\mathrm{d}y > 0$ and $\Psi_2 = \frac{\tau}{u\overline{\nu}}\int \nu(y)V(y)\mathrm{d}G(y) > 0$.*

*Proof.* Consider a variation $\varepsilon h(z)$. The resulting matching rate is by definition and Lemma 5:

$$m(\varepsilon) = \frac{1}{\tau}\int_0^1 1-\nu(y,\varepsilon)\,\mathrm{d}G(y) = \frac{1}{\tau}\left(1-\int \hat{z}'(y,\varepsilon)dG(y)\right) = \frac{1}{\tau}\left(1-\hat{z}(\overline{y},\varepsilon)g(\overline{y}) + \int g'(y)\hat{z}(y,\varepsilon)\mathrm{d}y\right),$$

so

$$m'(\varepsilon) = \frac{1}{\tau}\int_0^{\overline{y}} g'(y)\,\hat{z}'_\varepsilon(y,\varepsilon)\,\mathrm{d}y - g(\overline{y})\,\hat{z}'_\varepsilon(\overline{y},\varepsilon). \tag{18}$$

By Lemma 4,

$$\tau V(y;\varepsilon) = y - \hat{z}(y,\varepsilon) = \tau\beta_A(\varepsilon)\left(\Lambda_0+\varepsilon h\right)(\hat{z}(y,\varepsilon)). \tag{19}$$

Finding $U'(0)$: The surplus derivative is $U'(0) = um'(0) = u\overline{\rho}'_\varepsilon(0)/\tau$. Differentiate the second equation in (19) with respect to $\varepsilon$, evaluate at 0 and collect terms:

$$\begin{aligned}
\hat{z}'_\varepsilon(y,0) &= -\left(1+\tau\beta_A\alpha(y)\right)^{-1}\left(\tau\beta'_A(0)\,\Lambda_0(\hat{z}(y,0)) + \tau\beta_A(0)\,h(\hat{z}(y,0))\right)\\
&= -\nu(y,0)\left(\tau\beta'_A(0)\,\Lambda_0(\hat{z}(y,0)) + \tau\beta_A(0)\,h(\hat{z}(y,0))\right)\\
&= -\nu(y,0)\,\tau\beta\left(\frac{\overline{\rho}'_\varepsilon(0)\,\Lambda_0(\hat{z}(y,0))}{\overline{\nu}^2(0)} + \frac{h(\hat{z}(y,0))}{\overline{\nu}(0)}\right).
\end{aligned} \tag{20}$$

where the second line follows by Lemma 5 and the last line follows as $\beta_A(\varepsilon) = \beta/\overline{\nu}(\varepsilon)$.

Substituting (20) into (18) and omitting the $\varepsilon = 0$ argument in the function we obtain,

$$\overline{\rho}'_\varepsilon = -\int \nu(y)\,\tau\beta\left(\frac{\overline{\rho}'_\varepsilon \Lambda_0(\hat{z}(y))}{\overline{\nu}^2} + \frac{h(\hat{z}(y))}{\overline{\nu}}\right) g'(y)\,\mathrm{d}y + \nu(\overline{y})\,\tau\beta\left(\frac{\overline{\rho}'_\varepsilon \Lambda_0(\hat{z}(y))}{\overline{\nu}^2} + \frac{h(\hat{z}(y))}{\overline{\nu}}\right) g(\overline{y}).$$

Rearranging,

$$\overline{\rho}'_\varepsilon = -\tau\beta\Psi_1\left(\int_0^{\overline{y}} h(\hat{z}(y))\,\nu(y)\,g'(y)\,\mathrm{d}y - h(\hat{z}(\overline{y}))\,\nu(\overline{y})\,g(\overline{y})\right)$$

where

$$\begin{aligned}
\Psi_1 &\triangleq \left(\overline{\nu} + \tau\beta_A\int_0^{\overline{y}} \nu(y)\,\Lambda(\hat{z}(y))\,g'(y)\,\mathrm{d}y - \tau\beta_A\nu(\overline{y})\,\Lambda(\hat{z}(\overline{y}))\,g(\overline{y})\right)^{-1} \\
&= \left(\overline{\nu} + \tau\int_0^{\overline{y}} \nu(y)\,V(y)\,g'(y)\,\mathrm{d}y - \tau\nu(\overline{y})\,V(\overline{y})\,g(\overline{y})\right)^{-1} \\
&= \left(\int_0^{\overline{y}} \nu(y)\,g(y) + \tau\nu(y)\,V(y)\,g'(y)\,\mathrm{d}y - \tau\nu(\overline{y})\,V(\overline{y})\,g(\overline{y})\right)^{-1} \\
&= \left(\int_0^{\overline{y}} (\nu(y) - \tau\nu'(y)\,V(y) - \tau V'(y)\,\nu(y))\,g(y)\,\mathrm{d}y\right)^{-1} \\
&= \left(\int_0^{\overline{y}} \left(\nu^2(y) - \tau\nu'(y)\,V(y)\right) g(y)\,\mathrm{d}y\right)^{-1}
\end{aligned}$$

with the last line following from differentiating the first equation in (19) with respect to $y$: $\tau V'(y) = 1 - \nu(y)$. The term $\Psi_1$ is positive because $\nu'(y) < 0$.

Finding $V'[0]$: Differentiate the first equation in (19) with respect to $\varepsilon$ and evaluate at zero. Using (20), we have:

$$\tau V'_\varepsilon(y,0) = -\hat{z}'_\varepsilon(y,0) = \nu(y,0)\,\tau\left(\beta'_A[0]\,\Lambda_0(\hat{z}(y)) + \beta_A[0]\,h(\hat{z}(y))\right).$$

Integrating over seller types:

$$
\begin{aligned}
V'(0) &= \int V'_{\varepsilon}(y,0)\,\mathrm{d}G(y) \\
&= \beta'_A(0)\int \nu(y)\,\Lambda_0(\hat{z}(y))\,\mathrm{d}G(y) + \beta_A\int \nu(y)\,h(\hat{z}(y))\,\mathrm{d}G(y) \\
&= m'(0)\,\frac{\tau\beta_A}{\bar{\nu}}\int \nu(y)\,\Lambda_0(\hat{z}(y))\,\mathrm{d}G(y) + \beta_A\int \nu(y)\,h(\hat{z}(y))\,\mathrm{d}G(y) \\
&= m'(0)\,\frac{\tau}{\bar{\nu}}\int \nu(y)\,V(y)\,\mathrm{d}G(y) + \frac{\beta}{\bar{\nu}}\int \nu(y)\,h(\hat{z}(y))\,\mathrm{d}G(y).
\end{aligned}
$$

The second part of the result follows. □

**Lemma 7.** $\lambda \in \Delta(S)$ *is* $x^*$*-lower-censorship if and only if the corresponding spread function* $\Lambda$ *has the following form:*

$$
\Lambda(z) = \begin{cases} \underline{\Lambda}(z), & z \in [0, \mathbb{E}[x \mid x < x^*]], \\ \overline{\Lambda}(x^*) + F(x^*)(z - x^*), & z \in (\mathbb{E}[x \mid x < x^*], x^*), \\ \overline{\Lambda}(z), & z \in [x^*, \overline{x}]. \end{cases} \tag{21}
$$

*Proof.* Denote by $z^*$ the expected value of $x$ on the pooled part of $X$, $z^* = \mathbb{E}[x \mid x > x^*]$. If $\lambda$ is the $x^*$-lower-censorship, then a straightforward application of the spread-function definition obtains (21).

Conversely, suppose $\Lambda$ has a form given in (21). Let $H(z) = \Lambda'(z)$, where $\Lambda'$ is the right derivative. Function $H(z)$ then equals to $F(z)$ for $z \leqslant x^*$, to $F(x^*)$ for $x^* < z < z^*$, and to 1 for $z \geqslant z^*$. Such distribution corresponds to a $x^*$-lower-censorship. □

Now let's return to the proof of Theorem 2. Use Lemma 6 to obtain the expression for the variational derivative of the objective function with respect to $\varepsilon$:

$$\begin{aligned}\mathcal{J}'(0) &= \gamma U'(0) + (1-\gamma) V'(0) \\ &= \int_0^{\overline{y}} h(\hat{z}(y))\, \nu(y) \left(-g'(y)\,\Psi + (1-\gamma)\frac{\beta}{\overline{\nu}} g(y)\right) \mathrm{d}y + h(\hat{z}(\overline{y}))\, \nu(\overline{y})\, g(\overline{y})\, \Psi \\ &= \int_0^{\overline{y}} h(\hat{z}(y))\, \nu(y) \left(-\frac{g'(y)}{g(y)}\Psi + (1-\gamma)\frac{\beta}{\overline{\nu}}\right) g(y)\, \mathrm{d}y + h(\hat{z}(\overline{y}))\, \nu(\overline{y})\, g(\overline{y})\, \Psi \qquad (22)\end{aligned}$$

where $\Psi \triangleq (\gamma + \Psi_2 (1-\gamma))\, u\beta\Psi_1 > 0$.

If $g$ is weakly decreasing, then $g'/g \leqslant 0$, and so the integrand in (22) is positive for all $y \in Y$. Therefore, for any $\Lambda_0$ and any $h \geqslant 0$, $\mathcal{J}'(0) > 0$. The only spread function that cannot be improved like that corresponds to full disclosure so full disclosure is optimal.

If $g$ is log-concave, then $g'(y)/g(y)$ is decreasing. Then for any $\Lambda_0$, the integrand in (22) crosses zero from below at most once. If the integrand is always positive, then, again, full disclosure is optimal. If the integrand does cross zero once, then we will show that a lower censorship is optimal. Towards the contradiction, suppose that $\Lambda_0$ is not a lower censorship. We will construct a platform-profitable deviation $h$ as follows. Denote by $y^*$ the zero of the integrand in (22), and denote by $z^* = \hat{z}(y^*, \Lambda_0)$ the optimal cutoff of the $y^*$-type seller. Let $\tilde{\Lambda} \in \mathcal{L}$ be the lower censorship that passes through the point $(z^*, \Lambda_0(z^*))$, as in (21). Essentially, $\tilde{\Lambda}$ is the policy that discloses as little as possible on the left tail and as much as possible on the right tail but still passes through $(z^*, \Lambda_0(z^*))$. This lower censorship is unique because there is a unique line tangent to $\overline{\Lambda}$ and passing through $(z^*, \Lambda_0(z^*))$. We have $\tilde{\Lambda}(z) \leqslant \Lambda_0(z)$ for $z < z^*$ and $\tilde{\Lambda}(z) \geqslant \Lambda_0(z)$ for $z > z^*$; as $\Lambda_0$ is not a lower censorship, for some $z$ one of these inequalities is strict. Consider a deviation $h = \tilde{\Lambda} - \Lambda_0$. We have

$$\begin{aligned}
\mathcal{J}'(0) &= \int_0^{\overline{y}} h(\hat{z}(y))\,\nu(y)\left(-\frac{g'(y)}{g(y)}\Psi + (1-\gamma)\frac{\beta}{\overline{\nu}}\right) g(y)\,\mathrm{d}y + h(\hat{z}(\overline{y}))\,\nu(\overline{y})\,g(\overline{y})\,\Psi \\
&= \int_0^{y^*} \underbrace{h(\hat{z}(y))}_{\leqslant 0}\,\nu(y)\underbrace{\left(-\frac{g'(y)}{g(y)}\Psi + (1-\gamma)\frac{\beta}{\overline{\nu}}\right)}_{\leqslant 0} g(y)\,\mathrm{d}y \\
&\quad + \int_{y^*}^{\overline{y}} \underbrace{h(\hat{z}(y))}_{\geqslant 0}\,\nu(y)\underbrace{\left(-\frac{g'(y)}{g(y)}\Psi + (1-\gamma)\frac{\beta}{\overline{\nu}}\right)}_{\geqslant 0} g(y)\,\mathrm{d}y + \underbrace{h(\hat{z}(\overline{y}))}_{\geqslant 0}\,\nu(\overline{y})\,g(\overline{y})\,\Psi.
\end{aligned}$$

By the argument above, one of the inequalities is strict, and so $\mathcal{J}'(0) > 0$, contradicting the optimality of $\Lambda_0$.

The same logic applies if the platform maximizes only the matching rate, $\gamma = 1$, and the seller distribution $g$ is unimodal.

# References


Akbarpour, M., S. Li, and S. Oveis Gharan (2017): "Thickness and Information in Dynamic Matching Markets," *Working Paper.*

Anderson, A. and L. Smith (2010): "Dynamic Matching and Evolving Reputations," *Review of Economic Studies*, 77, 3–29.

Ashlagi, I., P. Jaillet, and V. H. Manshadi (2013): "Kidney Exchange in Dynamic Sparse Heterogenous Pools," *Working Paper.*

Bergemann, D., B. Brooks, and S. Morris (2015): "The Limits of Price Discrimination," *American Economic Review*, 105, 921–957.

Bergemann, D. and S. Morris (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

Burdett, K., S. Shi, and R. Wright (2001): "Pricing and Matching with Frictions," *Journal of Political Economy*, 109, 1060–1085.

Chade, H. (2006): "Matching with Noise and the Acceptance Curse," *Journal of Economic Theory*, 129, 81–113.

Che, Y.-K. and J. Hörner (2018): "Optimal Design for Social Learning," *Quarterly Journal of Economics*, forthcoming.

Cohen, P., R. Hahn, J. Hall, S. Levitt, and R. Metcalfe (2016): "Using Big Data to Estimate Consumer Surplus: The Case of Uber," *Working Paper*.

Coles, P., A. Kushnir, and M. Niederle (2013): "Preference Signaling in Matching Markets," *American Economic Journal: Microeconomics*, 5, 99–134.

Einav, L., C. Farronato, and J. Levin (2016): "Peer-to-Peer Markets," *Annual Review of Economics*, 8, 615–635.

Ely, J. C. (2017): "Beeps," *American Economic Review*, 107, 31–53.

Ely, J. C. and M. Szydlowski (2017): "Moving the Goalposts," *Working Paper*.

Fradkin, A. (2015): "Search Frictions and the Design of Online Marketplaces," *Working Paper*.

Gentzkow, M. and E. Kamenica (2016): "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, 106, 597–601.

Hoppe, H. C., B. Moldovanu, and A. Sela (2009): "The Theory of Assortative Matching Based on Costly Signals," *Review of Economic Studies*, 76, 253–281.

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Kircher, P. (2009): "Efficiency of Simultaneous Search," *Journal of Political Economy*, 117, 861–913.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a Privately Informed Receiver," *Econometrica*, 85, 1949–1964.

Kovbasyuk, S. and G. Spagnolo (2017): "Memory and Markets," *Working Paper*.

Lauermann, S. (2012): "Asymmetric Information in Bilateral Trade and in Markets: An Inversion Result," *Journal of Economic Theory*, 147, 1969–1997.

Levin, J. and P. Milgrom (2010): "Online Advertising: Heterogeneity and Conflation in Market Design," *American Economic Review*, 100, 603–607.

Myerson, R. B. (2000): "Large Poisson Games," *Journal of Economic Theory*, 94, 7–45.

Orlov, D., A. Skrzypacz, and P. Zryumov (2018): "Persuading the Principal to Wait," *Working Paper*.

Ostrovsky, M. and M. Schwarz (2010): "Information Disclosure and Unraveling in Matching Markets," *American Economic Journal: Microeconomics*, 2, 34–63.

Rayo, L. and I. Segal (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.

Segal, I. (2007): "The Communication Requirements of Social Choice Rules and Supporting Budget Sets," *Journal of Economic Theory*, 136, 341–378.

Smolin, A. (2017): "Dynamic Evaluation Design," *Working Paper*.

Tadelis, S. and F. Zettelmeyer (2015): "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment," *American Economic Review*, 105, 886–905.