Quality Versus Fit: Market Design and Externalities on Multidimensional Matching Platforms*

Sam Flanders†
Asia School of Business

Melati Nungsari‡
Asia School of Business

December 3, 2017

Abstract

This paper studies externalities in one-to-one matching markets when agents have preferences over multidimensional types by utilizing a minimal search setting where agents are either high or low quality and have an idiosyncratic per-match "fit" shock. Applications include dating and marriage, job search, and school choice. It identifies a novel source of externalities that does not exist in the one-dimensional models focused on in the previous literature, but is endemic in multidimensional settings, appearing in both search and frictionless matching models so long as nontransferabilities are present. Agents match too aggressively on traits where preferences are homogeneous across agents (*quality*), and too little on traits where preferences are heterogeneous across agents (*fit*). This effect is decomposed into a *intermatch externality* – when you match to someone, you impose a cost on the rest of the market by removing them from it, and a *intramatch externality* – you don't account for your partner's payoffs when choosing partners. Given these generic externalities, we provide a survey of instruments a matching platform could use to improve surplus, analyzing for each the efficiency properties of the solution and its ease of implementation under a variety of assumptions. Having the platform act as a middle-man to make the transfers that agents cannot make directly is an obvious way to achieve first best, but may be difficult to implement. Utilizing two-part tariffs can improve efficiency somewhat, and is easy to implement, but cannot address the inefficient fit/quality tradeoff. Splitting the platform along quality can achieve first best in some settings, but will forego increasing returns to scale if they exist in the model. Censoring agents' choice sets (curating the set of partners they can see on the platform) is only effective in some settings when studied in isolation, but when combined with two-part tariffs can achieve first best in any setting.

JEL classification: D21, C78, D83

Keywords: Externalities, Pricing, Monopolist, Matching, Search, Two-Sided Platforms

*We'd like to thank Asia School of Business for their support throughout this project. All mistakes are our own.

†E-mail: samuel.flanders@asb.edu.my. Address: Address: Asia School of Business in Collaboration with MIT Sloan, Sasana Kijang, 2 Jalan Dato' Onn, Kuala Lumpur, 50480, Malaysia. Phone: (+6013)-207-3225.

‡E-mail: melati.nungsari@asb.edu.my. Address: Asia School of Business in Collaboration with MIT Sloan, Sasana Kijang, 2 Jalan Dato' Onn, Kuala Lumpur, 50480, Malaysia. Phone: (+6013)-290-3225.

1 Introduction

Consider a simple matching model with two newly minted doctors and two residency programs, each with one opening. Both the doctors and the programs have preferences over *quality*. Program A is ranked first while Program B is ranked second, and student a is ranked first while student b is ranked second. However, there is also another form of variation, *fit*. Student a and program B are specialized in gynaecology, while student b and program A are specialized in nephrology. An obvious question for society, a medical association, or a for-profit matching platform that serves the medical labor market is the following: which students should go to which programs? More formally, if each potential pairing generates a match surplus, how do we maximize the total surplus in the market? For simplicity, let's say that each agent's quality contributes an additively separable payoff to the match surplus. Then, to maximize total surplus, we should assign a to B and b to A since the total quality in the market is same regardless, but a has a much better fit with B, and b has a much better fit with A. However, if salaries for each program are non-negotiable, then for appropriate relative weights on quality preference – that is, if both programs and students have a strong preference for prestige – the stable assignment will see a matching to A and b matching to B. The lost surplus relative to assignment $\{aB,bA\}$ arises from wedges between the preferences of the individuals in the market and those of the planner–matching externalities.

In this paper, we study these efficiency issues in two-sided one-to-one matching models – markets with two sides, each having a preference over the other, where each agent searches for a single partner. These models can be used to study marriage and dating, job search, and the assignment of students to schools, among other things. There is a rich literature on externalities in these markets, but it focuses on stylized one-dimensional preference models, typically assuming vertical preferences, where every agent agrees on the ranking of all agents on the other side of the market. This paper extends the literature by developing a very general model of matching over multiple partner traits that is simple enough to illustrate a novel source of externalities that, while absent in one-dimensional models, is endemic to matching models with multi-dimensional preferences. In particular, we show that tradeoffs between traits embodying quality and traits embodying fit generate externalities in a wide variety of settings. By quality, we mean that preferences over the trait are homogeneous, with all agents agreeing on the preference ordering of potential partners. Examples of quality variables in job search include prestige, ranking, and salary. Examples in dating and marriage include income, attractiveness, and status. By contrast, a trait embodying fit exhibits heterogeneous preferences – different agents have different, idiosyncratic preference orderings. In job search, location and specialization can be framed as fit variables. Examples in dating and marriage include location, tastes, and personal chemistry. Notably, horizontal preferences á la Hotelling (1929) and Salop (1979) are a common form of fit in a variety of applications. Having established these externalities, we study a variety of possible instruments to improve efficiency.

We focus on a search model of matching, where agents meet potential partners (who they can accept or reject) via a Poisson process. We allow the meeting rate to be either constant, termed *constant returns to matching* (CRM), or to be proportional to the number of agents on the platform, termed *linear returns to matching* (LRM). CRM illustrates a search market where the bottleneck on search is the agent's ability to evaluate parters – no matter how many people an agent meets, they can only evaluate a fixed number (say, 10) in a day. LRM illustrates a search market where the bottleneck on search is the availability of partners, implying that a larger platform then means more opportunities to match. CRM and LRM together exemplify most matching markets in real-life. We allow for both of the common search friction specifications, where time discounting and search costs that impose a cost for each meeting. For simplicity, we model quality via a binary vertical trait θ that determines the quality payoff recieved by their partner. Agents are either Studs (i.e. the high type, abbreviated H) or Duds (i.e. the low type, abbreviated L), where $\theta_H > \theta_L$. Fit is modeled as an idiosyncratic match shock $\psi \sim F(.)$ where every possible pairing has an associated mutual fit parameter. Agents permanently leave the market when they match.

When two individuals both accept a match, they generate a match surplus. If the agents can costlessly bargain over this surplus, we term it a *transferable utility (TU)* matching market. As with the Coase Theorem in one-sided markets, TU tends to resolve matching externalities[1]. However, like the Coase Theorem, TU is a very strong assumption, one that can be undermined by a variety of frictions, such as transaction costs, social norms against certain forms of transfers, and commitment problems, among many others. By contrast, *nontransferable utility* (NTU) assumes that the division of the match surplus is exogenous. Intermediate partially transferable utility settings are difficult to study, so papers in this literature typically focus on one or the other – TU for low friction markets and NTU for higher friction markets. In this paper, we'll assume NTU, focusing on matching markets like dating and marriage, where social norms limit transfers and the infrequency of negotiated dating and marriage contracts creates commitment problems. Some labor market applications include public sector jobs, where salaries are often non-negotiable, and entry level professional jobs that feature high levels of wage compression. For simplicity, we focus on a market with symmetric distributions of agents across the two sides.

While we focus on a search model for tractability, these externalities are highly general. As illustrated in the first paragraph, they extend to frictionless environments like Gale and Shapley (1962). We decompose the externalities into to components: we call the first the *intermatch externality* and the second the *intramatch externality*. The intermatch externality corresponds to the cost a matching agent imposes by taking their partner off the market. In our setting, this means that agents will match too aggressively on quality and too little on fit. This result is intuitive but not trivial: the fact that all agents agree on the ranking of the homogenous traits causes a strong intermatch effect – that is, if I am able to obtain the best potential match on the quality dimension, this means that I am removing that person from the market and thus preventing other agents from meeting that person. With more heterogeneous preferences, by contrast, my ideal match is less likely to be someone else's ideal match, so the intermatch effect

[1] Though some still persist in search models – see Shimer and Smith (2001).

is weaker or nonexistent. For example, if all agents' preferences are uncorrelated, pursuing my ideal match will not, in expectation, adversely affect the pool of potential matches for other matches. The intramatch externality arises when an agent's share of the surplus from a match isn't proportional to the overall match surplus. Pairings where one agent gets an unusually high payoff while the other gets an unusually low payoff may be rejected by the low payoff agent despite a good overall match surplus. In our model, this again leads to inefficiently aggressive matching on quality – an individual's match utility depends only on their match's type, while total match surplus depends on both, so matching to a Dud rather than a Stud is relatively more costly to the agent than to the planner. By contrast, fit payoffs are symmetric, creating no wedge between the agent's and the planner's incentives.[2]

Given these externalities, we reframe the matching market as a strategic platform operated by either a social planner or a profit-seeking monopolist and evaluate a variety of potential instruments to improve efficiency. In each case, we'll allow the monopolist to perfectly price discriminate with respect to Studs and Duds via fixed fees, so the monopolist's problem will reduce to the planner's problem of maximizing total surplus. We'll study three classes of instruments: pricing alone, splitting the platform, and censoring search. For pure pricing instruments and split platforms, we assume the platform can observe agent quality. For censored search we assume they can observe both fit and quality.

For pure pricing, we'll study two cases; the first is a two-part tariff. while per-interaction pricing[3] can resolve some externality issues, and in some special cases can achieve first-best, resolving the inefficient tradeoff between quality and fit requires the platform to make studs less picky when matching to duds and more picky when matching to studs, while per-interaction prices can only make an agent uniformly more picky or less picky. Second, we consider a more ambitious pricing scheme: match dependent pricing. If the platform can combine per-interaction pricing with pricing that depends both on own and match quality, the platform can generically achieve first-best. This can be framed as the platform implementing TU matching for its customers. However, this pricing scheme may be difficult to implement in practice.

For separate platforms, we show that, under CRM, creating separate platforms for Studs and Duds generates more surplus than a single platform. Externalities arise from matching on quality at the expense of fit, so partitioning agents by quality and forcing them to match solely on fit improves surplus. This parallels the separation result of Damiano and Li (2007). However, under LRM (and low search costs) a single platform can always outperform separate platforms. LRM implies that there are increasing returns to scale for platform size, and appropriate choice of acceptance regions will ensure that this effect dominates the externality cost. Examples of this include "elite" dating platforms like DateHarvardSQ and BeautifulPeople.

Finally, we consider a novel instrument: censored search: many matching platforms (eHarmony,

[2]Note that heterogeneous preferences do not require the symmetric payoffs we've assumed. However, symmetric payoffs are a logical consequence of horizontal preferences (if I'm close to you, you're also close to me). We briefly consider the case where the fit payoffs for the two matching agents are independent.

[3]A price charged for every draw an agent receives.

Chemistry.com) explicitly restrict their users' choice sets, only allowing them to see a small, curated subset of the market each day. Others disallow or limit search and filtering over the full set of potential partners (Tinder), forcing users to search via a stream of draws controlled by the platform. In either case, we can frame this as the platform limiting users' search. One common justifications for this is lowering search costs by excluding partners you'd never match with. Another is that the platform acts as an expert middleman that is better able to assess match quality than you are. However, our analysis of matching externalities suggests a third explanation: some matches may be individually rational for their participants, but may generate negative externalities that make them unattractive to the platform. The platform can censor these draws to improve total surplus. Note that there is a clear limitation to this instrument because it can't directly induce agents to accept matches they don't want. By way of analogy, we can think of type I and type II errors in this setting – the platform commits a "type I error" when a match is rejected that ought to be accepted. A "type II error" occurs when a match is accepted when it should not be. Censored search can eliminate "type II errors", but not "type I"–at least not directly.[4] We show that, under some parameterizations of the model, censored search can improve total surplus and achieve first best, but in other cases there are no "type II errors" to correct and censored search has no marginal value for the platform. Our analysis of censored search relates to previous papers that have studied the potential benefits of restricting the cardinality of the choice set in matching problems (Halaburda et al. (2016)). Our work is distinct in that it focuses on censoring specific draws rather than the number or rate of draws overall. We conclude this analysis by studying censored search with per-interaction pricing. As mentioned above, censored search is limited by the fact that it can only address "type II errors". However, positive per-interaction costs act as additional search costs, forcing agents to be less selective. This brings a larger proportion of draws into the region where censored search has traction – matches that the agents will accept if presented to them. By making agents extremely unselective, all errors become "type II errors", and censored search can generically achieve first best.

Summarizing, we contribute to the literature by characterizing previously unstudied externalities that are endemic to matching problems with multi-dimensional preferences and evaluating a variety of instruments a platform can use to eliminate these inefficiencies. We also provide the first analysis of the value of censored search in eliminating externalities on a matching platform.

The outline of the paper proceeds as follows. Section 2 reviews the literature. Section 3 provides the model setup. Section 4 studies the intermatch and intramatch externalities and provides two simulations that relax some of the assumptions made in the model. Section 5 provides a formal analysis of the efficiency properties of the six instruments discussed above, and Section 6 concludes, providing some qualitative discussion on the real-world feasibility of these instruments.

[4]Excluding "type II errors" changes the agent's optimization problem, which can lead to them accepting matches they would have rejected otherwise, but this second order effect is only effective in some cases.

2 Literature Review

This paper combines three literatures: search, matching, and pricing, The matching literature can be traced back to the seminal paper by Gale and Shapley (1962). Applications of matching theory to real-life markets have included the kidney exchange market (Roth, Sönmez, and Ünver (2005), Roth, Sönmez, and Ünver (2007), students to public school matching (Abdulkadiroğlu, Pathak, and Roth (2005), Abdulkadiroğlu et al. (2005)), and the job search literature (Bulow and Levin (2006)). In terms of utility specifications, the first two applications focus on non-transferable utility, which is similar to our paper. Bulow and Levin (2006), however, model transferable utility, where wages are taken as prices. Aside from the utility specification, another issue of interest when studying matching models is *assortation*. Positive assortation in equilibrium implies that high types match with other correspondingly high types, whereas negative assortation means that high types match with low types. Assuming non-transferable utility, Becker (1973) proved conditions under which positive assortation occurs. As in many other matching papers, we also assume that matching is random, exogenous, and non-targeted.

Previous papers in the literature have focused on matching where agents have preferences over only one dimension. For example, in Burdett and Coles (1997), agents are vertically-differentiated with characteristics summarized into a single number, defined as agent's *pizazz*, which takes a value in the interval $[0, 1]$. Burdett and Coles analyze the matching equilibrium by studying steady-state conditions and find a rich set of equilibria. Agents match in assortative partitions in equilibrium, where the interval $[0, 1]$ for both types of agents is broken up into distinct pieces with agents in the same piece matching randomly with each other.

The search literature, on the other hand, started with the study of job search in the labor market (McCall (1970), Burdett (1978), Mortensen and Pissarides (1994), Rogerson, Shimer, and Wright (2005)). It has since expanded to include other applications as well, including the marriage and dating market (e.g. Cornelius (2003)). Applications of search theory to the dating market usually assumes non-transferable utility, as this mirrors real-life more closely than the assumption of transferable utility. In particular, Burdett and Wright (1998) and Adachi (2003) both study properties of equilibria in a search models with non-transferable utility.

This paper incorporates pricing theory in a world with externalities and a profit-maximizing monopolist. We study two cases: the first is when the monopolist can only operate one platform (perhaps due to high fixed or operating costs), and the second is when monopolist has the ability to operate two platforms that cater to the different types. This is in contrast to Rocher and Tirole (2010), who study platform competition with two-sided markets. Pricing in search markets has also been studied by other authors (e.g. Bester (1994)) in the context of a buyer-seller relationship where buyers are looking to buy a good from a price-posting seller. However, models such as Bester (1994) do not have the complexity of externalities which affect agent's optimal matching behavior in equilibrium, as in our model.

The closest paper related to the pricing framework in this paper Bloch and Ryder (2000), where they study the provision of matching services in a model of two-sided search. Agents are distributed on the unit interval and their utility is equal to the index of their mate. Bloch and Ryder (2000) find that in a search equilibrium, agents form subintervals and are only matches to agents inside their own class, a result that closely mirrors that of Burdett and Coles (1997). The two main differences between our paper and theirs is my inclusion of the heterogenous trait ψ and the fact that we assume that the monopolist utilizes a two-part tariff pricing structure. We are primarily interested in the signs of the per-interaction prices–be it positive, the non-distortionary (or, as we see later, the 'no externalities') level of 0, or even negative. Bloch and Ryder (2000), on the other hand, study how two separate pricing structures (uniform and commissions on the matching surplus) affect equilibrium behavior in agents. Another similar paper is Damiano and Li (2007), where the monopolist again faces two sides of the market with each side having characteristics distributed on a compact interval. The monopolist is able to choose a sorting and pricing structure in order to maximize revenue. They then show that the revenue-maximizing sorting is efficient.

We would like to make two final notes regarding the existing literature. To relate the horizontal (heterogenous) matching component used in our paper to the traditional horizontal-differentiation model presented by Salop (1979), note that the horizontal component utilized in our paper can be rationalized by preferences on the Salop circle. In particular, for any distribution function F of agents on the circle, we can choose the cost of matching away from one's ideal type so that the idiosyncratic matching shock in our model is distributed according to F.

The second note relates to the following proposition:

Proposition 1 (Browning, Chiappori, and Weiss (2014))**.** *Suppose utility is transferable. A stable assignment must maximize total surplus over all possible assignments.*

Proposition 1 then tells us that in the case where utility is transferable, a *stable assignment*, which corresponds to the steady-state equilibrium in our model, must maximize total social surplus. In other words, in our environment with perfectly observable types and no bargaining, the monopolist charges prices to essentially reallocate surplus amongst agents *as though* utility is transferable. By maximizing total surplus, the monopolist in our model shifts the allocation of utility amongst agents from being non-transferable to utility being transferable.

3 Model

3.1 Environment

We study an infinite-horizon search model of matching, where each side of a two-sided market search for partners on the other side. Each side consists of two types of agents, characterized by a quality θ: High ("Studs") and Low ("Duds") (abbreviated H and L). Types are perfectly observable to the monopolist.[5] Time is continuous and both agents and the firm discount at the

[5]This is an assumption that is technically difficult to relax. For an exposition/explanations as to why this is the case, and some basic pricing results, please refer to this paper's online appendix at http://www.melatinungsari.com/research.html.

same rate r. We will be focusing on a steady-state equilibrium, which is defined by equating the total inflow into the platform with the total outflow out of the platform. Using these steady-state conditions on inflows and outflows allows us to solve for the prorportion of Studs on the platform, α, that makes the mass N of agents on the platform constant.

Once agents enter the market, they receive a stream of draws–that is, potential partners. This stream is characterized by a Poisson process with arrival rate λ. Throughout, we will make one of two assumptions about the rate of draws agents face: linear returns to matching (LRM) or constant returns to matching (CRM).[6]

Assumption 1 (Assumption 1A (LRM))**.** *Agents receive a rate of draws λ proportional to the mass of agents on the platform, normalized to N.*

Assumption 2 (Assumption 1B (CRM))**.** *Agents receive a constant rate of draws λ , normalized to* 1*.*

Linear returns to matching means that the frequency of draws is proportional to the mass of agents on the platform and that thick markets make search faster. Linear returns may be more realistic on online matching platforms like dating and job search sites where the choice set can easily and effectively be filtered down to a manageable size, effectively allowing users to search through potential matches faster when there are more agents on the platform. Constant returns to matching may be more appropriate for traditional forms of search where finding potential matches is time consuming and these frictions put an upper bound on the number of draws an agent can consider, regardless of the size of the market.

For each draw, agents incur a search cost of s, and meet a Stud with probability α and Dud with probability $1 - \alpha$. This probability α is endogenous and will be determined in equilibrium by using steady-state conditions.

Once the agents meet, they learn each others' types, receive an idiosyncratic matching shock ψ (fit) and decide whether or not to match. If both agents decide to match, they do so forever. If one or both do not want to match, they return to search. The idiosyncratic shock ψ is drawn from a distribution on $[0, m]$ with cumulative distribution F and density function f. Note that the two sides of the market have identical ψ distributions. A intuitive interpretation for ψ is given below.

The matching utility that an agent i receives when matching to agent j is given by

$$u_i(\psi_{ij}, \theta_j) = \theta_j + \psi_{ij} \tag{1}$$

We'll condense this to u_i where appropriate. There are two characteristics of the matching utility that we would like to elaborate on. First of all, note that agents receive the other agent's type as utility and not their own. This also means that all agents in the model prefer the Studs to

[6]Note that LRM is sometimes referred to as a quadratic search technology, owing to the quadratic nature of the total number of draws in the market as a function of the number of agents, and CRM is sometimes referred to as a linear search technology based on the linear rate of total draws in the market as a function of the total number of agents.

Duds. Secondly, note that the shock ψ is the same for both partners, that is the shocks ψ that i and j receive are *perfectly correlated*. We can interpret ψ as either a more general heterogenous preference or a horizontal preference. For example, the shock ψ may measure both agents' mutual 'fit' with one another. We argue that symmetric or highly correlated preferences are common to many heterogenous traits, and especially to horizontal traits. For example, suppose all agents prefer partners who are geographically closer. Then two agents who are close to one another will *both* find each other more attractive than agents who live far away. In any case, we will relax this assumption of perfect correlation of ψs in simulations done in the next section, which is devoted to an in-depth explanation of the intermatch and intramatch externalities.

3.2 Agent's Problem

For all but a measure zero proportion of the time, agents are engaged in search and make no decisions. In the instant when an agent receives a draw, they simply decide whether to match or to continue searching.

Definition 1. *Define a strategy for type i as*

$$S_i : \{\theta_L, \theta_H\} \times [0, m] \;\to \{accept, reject\}$$

The interpretation for the strategy is that, for a given own-type, it takes the ψ-draw and type of the opposite agent as arguments and returns a decision on whether or not to form a match. Then $\overrightarrow{S}$ is the strategy profile for the continuum of agents on both sides. Define $o_i(S) \equiv Pr[match|draw, i, S]$. We'll often truncate this to o_i or o_{θ_i}. Define $U_{\mathrm{i}}(s)$ as the expected lifetime utility for i while searching, which we'll also truncate to U_{θ_i} or U_i as appropriate. Define M_i as the set of potential partners that will generate matchs (mutual acceptance) for i. set of $\psi's$ accepted by i and a draw of type θ. The dynamic program for this environment gives us the following condition for a small time period dt:

$$U_{\mathrm{i}} = \frac{U_{\mathrm{i}}(1 - \lambda o_i dt) + \lambda o_i E_{M_i}[u_i]dt - s\lambda dt}{1 + r\,dt}$$

Taking the limit as $dt \to 0$ and rearrainging, we have

$$U_{\mathrm{i}} = \frac{o_i E_{M_i}[u_i] - s}{r/\lambda + o_i} \tag{2}$$

Then, at any time t where agent i has received a draw, the continuation value of rejecting that draw is precisely $C_i \equiv U_{\mathrm{i}}$.

Optimality requires that an agent accept any draw whose utility exceeds U_i. Note that (**??**) is independent of the current draw, while u_i is strictly increasing in ψ_{ij} and θ_j. This implies that i's optimal strategy takes the form of a *threshold strategy.*

Definition 2. *Agents optimally accept a draw when the $\psi_{ij} + \theta_j$, and choose to continue searching otherwise. Formally,*

$$S_i^*(\theta_j, \psi_{ij}) = \begin{cases} \textit{reject} & \textit{if } \psi \geq \psi_i^*(\theta_j) \\ \textit{accept} & \textit{if } \psi < \psi_i^*(\theta_j) \end{cases} \tag{3}$$

There are then two thresholds for each agent of quality θ, $\psi_\theta^* = \{\psi_\theta^*(H), \psi_\theta^*(L)\}$. The optimal thresholds satisfy the following equation:

$$\theta_H + \psi_\theta^*(H) = \theta_L + \psi_\theta^*(L) = C_\theta \tag{4}$$

Now, we characterize the agent's optimal thresholds and expected continuation utilities. The characterization of the optimal thresholds and expected continuation utilities will allow us to rank the relative positions of the four thresholds. In particular, as we will soon see, the Studs will be *pickier* when it comes to matching. An agent i is said to be *pickier* in the matching process than agent j if $C_i > C_j$.

Proposition 2. *For any $\theta \in \{\theta_L, \theta_H\}$, $\psi_\theta^*(L) > \psi_\theta^*(H)$.*

This proposition informs us that any fixed type of an agent is pickier about dating duds than studs. This is a direct consequence of studs being seen as universally more attractive. In the next proposition, we study the relationship between the continuation values of the two types of agents.

Proposition 3. *In a steady-state equilibrium, $C_H > C_L$.*

Corollary 1. *For all $\theta \in \{\theta_L, \theta_H\}$, $\psi_H^*(\theta) > \psi_L^*(\theta)$*

From the above propositions, we can deduce that H always has the upper hand in determining whether or not a match forms when two agents meet. This means that we can reduce the set of thresholds to characterize. In particular, a Dud's threshold for matching with a Stud ($\psi_L^*(H)$) is a non-binding constraint since the decision whether or not to marry depends solely on $\psi_H^*(L)$. Thus, we now only have to characterize three thresholds: $\psi_H^*(H), \psi_H^*(L), \psi_L^*(L)$.

With these continuation values in hand, we can express (**??**) explicitly for Studs and Duds:

$$C_H = \frac{(\alpha \int_{\psi_H^*(H)}^{a} (\theta_H + \psi)\, f(\psi)\, d\psi + (1-\alpha) \int_{\psi_H^*(L)}^{a} (\theta_L + \psi)\, f(\psi)\, d\psi) o_H - s}{r/\lambda + o_H} \tag{5}$$

where

$$C_L = \frac{(\alpha \int_{\psi_H^*(L)}^{a} (\theta_H + \psi)\, f(\psi)\, d\psi + (1-\alpha) \int_{\psi_L^*(L)}^{a} (\theta_L + \psi)\, f(\psi)\, d\psi) o_L - s}{r/\lambda + o_L} \tag{6}$$

where $o_H = \alpha\,(1 - F\,(\psi_H^*(H))) + (1-\alpha)\,(1 - F\,(\psi_H^*(L)))$ and $o_L = \alpha\,(1 - F\,(\psi_H^*(L))) + (1 - \alpha)\,(1 - F\,(\psi_L^*(L)))$.

An important note to make is that the following proposition holds $\alpha$and λ fixed. In equilibrium, however, α and λ are not fixed and are determined endogenously. This is because of the previously mentioned fixed point problem: α and λ affects the thresholds, which then in turn affect α and λ. However, it is useful to have the comparative statics by holding them constant, since the agents themselves take them as constant when they optimize.

Proposition 4. *For distributions with continuous support, in a steady-state equilibrium the following is true:*

1. *All thresholds strictly decrease with s.*
2. *All thresholds strictly decrease with r.*
3. $\frac{\partial C_H}{\partial \alpha} > 0$
4. $\frac{\partial \psi_H^*(H)}{\partial \alpha} > 0$, $\frac{\partial \psi_H^*(L)}{\partial \alpha} > 0$

The first result is that thresholds strictly decrease with search costs. This means that higher search costs cause agents to become less picky since the cost of searching on the platform has now increased. A lower r, which corresponds to more patient agents, causes agents to become less picky. All else equal, agents are now content with staying on the platform and waiting for longer to obtain a better match and random draw ψ. The last two results of the previous proposition relate to the behavior of the studs when the proportion of the studs, α, changes. In particular, the results show that studs are both better off and pickier when there are more studs on the platform. This is not a surprising result since studs are seen as the superior type by everyone on the platform. We are not able to prove similar results relating to how the optimal strategy of the duds changes with respect to changes in the proportion of studs on the platform. This is because this result would depend strongly on whether or not studs are willing to match to duds. In particular, one can imagine a situation where studs are significantly better than duds (i.e. $\theta_H - \theta_L$ is very large), in which case studs would never match to duds. This situation would cause studs to be much worse off in the presence of many duds. In particular, the value of the platform as a means to obtain a match decreases to the duds–he is not able to increase his odds of obtaining a match by joining a platform in which the vast majority of others on the platform do not wish to match to him.

3.3 Dynamics

We will be focusing on steady-state equilibrium in this paper. To do so, we introduce the dynamics in the model. At time t, there is a time-invariant exogenous inflow rate, normalized to 1, of duds joining the platform. The exogenous inflow rate of Studs is $i_H \in R$. There is also a mass of agents of each type already on the platform, denoted N_t. For each quality, agents leave based on the outflow rate for that quality, the mass of agents of that quality $N_t\alpha_t$ and $N_t(1-\alpha_t)$, respectively, and the rate of draws. We can then state the transitional equation for N_t, which is the following

$$N_{t+dt} = N_t = N_t + \, i_H \, dt + 1 \, dt - N_t \lambda_t \alpha_t \, o_{tH} \, dt - N_t \lambda_t (1-\alpha_t) \, o_{Lt} \, dt$$

Rearranging, taking the limit, and assuming a constant N_t, we have

$$N = \frac{i_H + 1}{\lambda \left(\alpha o_H + (1-\alpha) o_L \right)} \tag{7}$$

Now, we are ready to define a steady-state. The steady-state will determine the values of α and λ.

Definition 3. *A steady state consists of inflows, outflows, and a constant distributions of agents $\{N, \alpha\}$ on the platform. It is characterized by the following equations:*

$$i_H = o_H \alpha N \tag{8}$$

$$1 = o_L(1-\alpha)N \tag{9}$$

$$N = \frac{i_H + 1}{\lambda\left(\alpha o_h + (1-\alpha)o_L\right)} \tag{10}$$

The proportion of H types on the platform, α, can then be solved for from the steady-state equations.

$$\alpha = \frac{i_H o_L}{i_H o_L + o_H} \tag{11}$$

There are some intuitive properties of α that are worth highlighting. First of all, note that α becomes smaller with a higher o_L. As duds types leave the platform faster, the proportion of studs types on the platform will increase. Similarly, the quicker studs leave the platform, the lower α becomes since o_H has increased. This analysis can also be conducted on the proportion of agents that are let into the platform: as i_H increases, α increases; and as more duds enter, the lower the proportion of studs on the platform becomes.

Given this expression for α, we can complete our analysis of λ and N. Plugging 11 into 7, we have

$$N = \frac{\frac{i}{o_H} + \frac{1}{o_L}}{\lambda}.$$

For CRM we have $\lambda = 1$ and

$$N = \frac{i}{o_H} + \frac{1}{o_L} \tag{12}$$

For LRM we have$\lambda = N$ and

$$N = \sqrt{\frac{i}{o_H} + \frac{1}{o_L}} \tag{13}$$

4 Multidimensional Matching Externalities

Disclaimer: We are in the midst of converting this paper to continuous time. Some propositions and proofs may still be in discrete time.

We will now discuss the externalities studied in this paper in detail. First, we will consider the intermatch externality, where self interested agents ignore the fact that their matching choices change the distribution of agents available to others, and that this effect is less likely to be costly with heterogeneous preferences than with homogeneous preferences. For expositional ease, we will first consider a frictionless model in the vein of Gale and Shapley (1962). Before we do so, we would like to introduce two important concepts relating to matching markets, which are the *matching function* or *assignment* and *stability*.

Definition 4 (Matching μ)**.** *A matching μ is a mapping from each agent to their match.*

Definition 5 (Stability - Transferable and Non-Transferable Utility)**.** *In the case where utility is non-transferable:*

- *A matching μ is stable if there is no a and b such that $b \succ_a \mu(a)$ and $a \succ_b \mu(b)$. (a, b) is called a blocking pair.*

In the case where utility is transferable:

- *A matching μ is stable if $\exists$ a feasible allocation rule $v : A \cup B \to \mathbb{R}$ giving the payoff for each matched agent such that there is no a and b such that $u(a, b) > v(a) + v(b)$. (a, b) is called a blocking pair.*

As mentioned before, the intermatch and intramatch externalities are not specific to search models, but generally present in non-transferable matching models even without search costs. Consider a matching market identical to the one described above, but without search costs and with a finite and equal set of agents on each side. Agents can observe the type of every possible match and costlessly propose to potential suitors, so there will be a stable matching where every agent matches to the most preferred partner that will accept her (him). Clearly, one's matching decision will affect others, since it changes the set of available partners for all other agents. In particular (and here, we introduce some new notation), the man m that a woman w receives must be withheld from some other woman w', who counterfactually would have received m as a match if w had matched differently.

However, self interested agents do not care about what happens in other matches. Shapley and Shubik (1971) showed that, with perfectly transferable utility, transfers can allow such an agent w' to force others to internalize the costs they impose on her in a manner analogous to the Coase Theorem. However, with non-transferable utility there is no mechanism available to internalize externalities. Note also that agents in a multidimensional matching environment will generically have to make tradeoffs between the various traits they care about. Unless every agent can match to the most desirable possible partner along all traits simultaneously, they will have to choose between matches that are better along one dimension and matches that are better along another.

Thus, if agents have preferences over a vertical trait θ, where all agree on the preference ordering; and a heterogeneous trait ψ, where every agent's preference ordering is independent and identically distributed, we can expect agents without access to transfers to make an inefficient trade-off between the traits. This is because, along the vertical trait, a pursuing a good match for oneself must mean a bad match for another—there are only so many high type agents, and this intermatch ensures that taking one out of the market worsens the outcome for at least one other agent along this trait. With the heterogeneous trait, by contrast, pursuing a good match for oneself has no effect on the distribution of remaining agents in expectation. Thus, good θ matches impose a negative externality on some other agent, while good ψ matches do not, and self interested agents will match too aggressively on θ.

With modular preferences over θ, or, more generally, vertical traits, this effect is especially stark. In a model where one's utility from a match is their match's type (or more generally where match utility is the sum of the agent's types) the total surplus is invariant to the matching assignment—it doesn't matter who matches to whom, simple algebra shows that the total surplus in the market is the sum of every agents type, with different assignments simply changing the order of summation. In a model with a modular vertical trait and another trait, as in this paper, total surplus must then depend only on sorting along the other trait, since the surplus accruing from the vertical trait is invariant to assignment.

Formally, consider a frictionless two-sided matching market with

- n agents on each side.
- k modular vertical traits $\theta = (\theta_1, ..., \theta_k)$, where the vector for vertical characteristics of a a male agent (m) i is θ_i^m, and his jth vertical characteristics is $\theta_{i,j}^m$.
- l other traits $\psi = (\psi_1, ..., \psi_l)$.
- match surplus for man i and i's match $\mu(i)$ given by

$$u(\theta_i^m, \theta_{\mu(i)}^w, \psi_i^m, \psi_{\mu(i)}^w) = \sum_{j=1}^{k}(\theta_{i,j}^m + \theta_{\mu(i),j}^w) + f(\psi_i^m, \psi_{\mu(i)}^w)$$

Define TSS_θ as the total social surplus from matching on the vertical type θ, TSS_ψ as the total social surplus from matching on heterogenous traits ψ, and TSS as the overall total surplus (i.e. $\text{TSS} = \text{TSS}_\theta + \text{TSS}_\psi$). Then, we have

$$\text{TSS}_\theta \equiv \sum_{i=1}^{n}(\sum_{j=1}^{k}(\theta_{i,j}^m + \theta_{\mu(i),j}^w))$$

$$\text{TSS}_\psi \equiv \sum_{i=1}^{n} f(\psi_i^m, \psi_{\mu(i)}^w)$$

$$\text{TSS} = \text{TSS}_\theta + \text{TSS}_\psi$$

Proposition 5. *In the above environment, a non-transferable utility matching must exhibit weakly lower TSS_ψ than the TSS_θ-maximizing assignment, which is also the transferable utility stable matching.*

Proof. Appendix. □

This only shows that the non-transferable matching must have a weakly lower TSS due to ψ traits, but generally, in a non-transferable utility framework, agents will prefer to get better θ draws, even if the overall θ endowment is unchanged, and will thus make inefficient tradeoffs against ψ matching, so the inequality will typically be strict. It will also often be that the TSS_θ will be higher for non-transferability than for the first best assignment, but it is possible that non-transferable utility-induced inefficiencies can cause both to be lower in non-transferability utility frameworks.

This externality can be translated into the search environment as well, but the mechanism of action is a bit more involved. Without frictions, we can talk of specific agents preventing specific other agents from getting a desired match, making the externality extremely clear. In a search model with a continuum of agents, by contrast, we can only talk about measurable masses of agents, distributions of potential draws, and expectations over matching outcomes. However, in a steady-state search model, inflow distributions must equal outflow distributions, and thus the distribution of matches that agents receive each period must equal the inflow distribution. Thus, if a nonzero mass of women match to high type men, that equal mass of high quality men is unavailable to other women who match in that period, and they must receive lower quality matches on average than if they were able to match to this group. Again, a vertical trait induces a clear externality, while a heterogeneous trait generally does not. The mechanism by which this assignment happens is search, however, so counterfactually different assignments must be implemented through differing distributions of agents and cutoff strategies—that is, this externality is mediated by the familiar thick market and congestion externalities. For example, if high types only accept one another, this will change the distribution of agents on the platform, and that will change the distribution of draws agents face and the set of agents who will accept them.

The other externality we study is the intramatch externality. This sort of externality appears in many economic environments, notably with transaction costs in Coasian models. Fundamentally, it arises from a wedge between private match utility u_s and match surplus u. If u_s is not proportional to u, agents value a given match differently from a social planner, and may make socially inefficient acceptance and rejection decisions. This externality can appear (along with the intermatch externality) in a one-dimensional model due to a tradeoff between time (discounting or search costs) and match quality, and is in fact present in models like Burdett and Coles (1997), though it has generally not been discussed explicitly in the literature. A simple example is the case with utility being your match's type and non-transferability. Total surplus may be quite high when a high type matches to a low type, and it may be socially optimal for the match to proceed so as to avoid more time costs, but the high type only receives the low type utility, ignoring the large benefit she provides to her partner, and may choose to reject. Because there are no transfers, the low type cannot offer some on their large benefit in order to induce a match. In multidimensional models there is an additional tradeoff between traits, and thus this externality can appear even in environments without time costs.

In our model, there is a wedge between agents' private matching utilities (u) and the resulting match surplus (u_s) for vertical traits, due to the assumption that utility obtained by an agent is his match's type, but not with the heterogeneous trait ψ, which is symmetric across the sides of the market—that is, your ψ draw is the same as your match's ψ draw, so u and u_s due to ψ are proportional. This means that both externalities work in the same direction—people match too aggressively on the vertical trait. Ideally, we would break the model out into cases with each externality in order to decompose the effects, and that is an avenue for future work, but we will argue to this coincidence of externalities fits with the stylized facts of multidimensional matching markets.[7]

Specifically, many traits over which agents have heterogeneous preferences will tend to exhibit symmetric payoffs and thus no or small wedges. As mentioned before, horizontal traits are a common type of heterogeneous preference, and they have symmetric payoffs by construction—if two agents m and w prefer matches closer to one another, then if m gets a high payoff from w along this dimension they must be close, which means that w also gets a high payoff from m. This can also apply to traits like race, religion, values, or even traits that factor into decision making like patience and risk aversion where agents often prefer a match similar to themselves.[8] Even with heterogeneous traits that are not horizontal, symmetric payoffs are quite plausible—traits like mutual chemistry are likely to be approximately symmetric. With vertical traits, by contrast, symmetry is possible but requires strong functional form assumptions that are orthogonal to the homogeneous preference ordering assumption. Thus, a model with a wedge only for the vertical trait captures these arguments in a simple and stylized way.

We will now explore several simulations of a frictionless non-transferable utility market analogous to the baseline model of this paper (without a strategic platform), as well as several permutations of this model. The results of the simulation are in Table 1. We also estimate the corresponding transferable utility assignment, again using the Shapley and Shubik (1971) result that transferable utility gives the TSS-maximizing assignment. This allows us to see the first best outcome and evaluate the externalities caused by non-transferable utility. We include these simulations both to show that these externalities are extremely general and do not depend on the assumption of large search frictions, as well as to explore the effect of alternate assumptions on match surplus function on the externalities in this environment.

Table 1: Frictionless Matching Simulations

[7] We will also study frictionless analogues with different assumptions on which traits induce wedges shortly.

[8] You may prefer a match with similar patience and risk aversion since you?ll be making many joint decisions and you?d prefer concordance.

Correlated ψ, 50 agents on each side, 40 iterations					
	β	R^2	TSS_θ	TSS_ψ	$TSS = TSS_\theta + TSS_\psi$
Non-transferable Utility, $a = 0$	0.79	0.61	25.00	45.81	70.81
Non-transferable Utility, $a = 1$	0.82	0.67	31.51	45.60	77.11
Transferable Utility, $a = 0$	-0.01	0.01	25.00	48.40	73.40
Transferable Utility, $a = 1$	0.78	0.59	30.93	47.52	78.46
Uncorrelated ψ, 50 agents on each side, 40 iterations					
Non-transferable Utility, $a = 0$	0.82	0.66	25.00	40.33	65.33
Non-transferable Utility, $a = 1$	0.87	0.75	31.85	40.06	71.91
Transferable Utility, $a = 0$	0.03	0.02	25.00	44.55	69.55
Transferable Utility, $a = 1$	0.83	0.70	31.81	42.69	74.50

These Monte Carlo simulations study markets with 50 agents on each side, and utilize 40 iterations each. A man m and a woman w obtain the following utility when deciding to match:

$$u_m = \frac{(1+a)\theta_m^a \theta_w + \psi_m}{2}, \; u_w = \frac{(1+a)\theta_w^a \theta_m + \psi_w}{2}$$

Note that a is a parameter determining the supermodularity of the utility function. When a is zero, the utility function is analogous to that of the baseline model, while $a = 1$ gives a vertical payoff component comprised of the product of one's own type and one's match's type. This means that higher types generate more utility, and thus make matching high types to high types and low types to low types generates more surplus than matching high and low types, all else equal. We'll treat the case where $\psi_m = \psi_w$, as in the baseline model, and also treat the case when they are independently drawn, dropping the symmetry assumption. Note that, with supermodular utility, the vertical trait in fact exhibits symmetric payoffs, so we will be able to model all four permutations of wedge and no-wedge over the two traits. ψ and θ are drawn from $U[0, 1]$, and the $(1 + a)$ term is a normalizing factor to account for the fact that, with the supermodular specification, utilities will generally be lower since types are drawn from $[0, 1]$ with an average of 0.5, so when $a = 0$, $\theta_m^a \theta_w$ is 0.5 on average, while when $a = 1$, the match's type is now multiplied by another draw from $[0, 1]$, yielding 0.25 in expectation with random assignment. Defining an agent's own type as (θ, ψ), match's type as (θ', ψ'), and a match μ as $\mu(\theta, \psi) = (\theta', \psi')$, where $\mu_1(\theta, psi) = \theta'$ and $\mu_2(\theta, \psi) = \psi'$. With vertical traits, we are interested in how aggressive the matching is along this trait, since we have predicted that it will be too high with non-transferable utility. Since high types are mutually desirable and low types the opposite, we can expect sorting along the vertical trait to induce assortative matching as per Becker (1973), with high types using their attractiveness to ensure a high type match and low types being stuck with other low types. Thus, more effort to match along the vertical trait should induce more correlation between own θ and match's θ, and in a sufficiently large market with only vertical sorting, $E(\mu_1(\theta, \psi)) = \theta$. Thus we will specify a simple regression

$$\mu_1 = \beta\theta + \epsilon$$

and report the β and R^2 vales. Additionally, we will report TSS, as well as TSS_θ and TSS_ψ, with the expectation that TSS_ψ and TSS should be higher for transferable utility and TSS_θ should be weakly higher for non-transferable utility.

Comparing non-transferable utility and transferable utility when $a = 0$ and ψs are correlated, the baseline environment for the paper, we see that TSS_θ is the same for both, as guaranteed by the above proposition—with modular utility, the matching assignment doesn't matter. However, in non-transferable utility agents sort quite a bit on the vertical trait—high types get better matches, but those better matches come at a cost to low types. The coefficient is relatively close to 1 and one's own vertical type explains 60% of the variation in match's vertical type. This comes at a clear cost to TSS_ψ, which is much lower for the non-transferable utility case, and thus TSS is also much lower, demonstrating the cost of these externalities.

Comparing non-transferable utility and transferable utility when $a = 1$ and ψs are correlated, we see sorting on the vertical trait for both the first best and the non-transferable utility assignments since supermodularity means that some assortation is desirable, but there is more sorting in theta with non-transferable utility and worse assignments in ψ in terms of TSS_ψ. Now there is no wedge for either trait, so these externalities come from intermatch, and intermatch is also weaker since the supermodularity means the costs imposed on low types are scaled by their own type and are thus lower than the benefits accruing to high types. However, there is still a modest loss in TSS due to intermatch.

Comparing non-transferable utility and transferable utility when $a = 0$ and ψs are uncorrelated, we again see sorting on the vertical trait only for non-transferable utility assignments. Now there is a wedge for both traits, so the intramatch externalities should largely cancel. We see a large shortfall in TSS_ψ due to intermatch with non-transferable utility, demonstrating the significance of this externality, especially with modular or close to modular utility.

Finally, comparing non-transferable utility and transferable utility when $a = 1$ and ψs are uncorrelated, we again see sorting on the vertical trait for both assignments. Now there is a wedge only for ψ, so the externalities have opposite effects. The intermatch externality clearly dominates here, since there is still more sorting along the vertical trait in the non-transferable utility case and less along ψ.

5 Resolving Multidimensional Matching Externalities

Disclaimer: We are in the midst of converting this paper to continuous time. Some propositions and proofs may still be in discrete time.

Having established the ubiquity of these externalities in settings without transfers, we now ask what instruments a platform could utilize to counteract them. We look both at pricing strategies, like two-part tariffs, and changes to the structure of the market itself, such as partitioning the platform along quality. In each subsection, we derive some basic results about the efficiency of the instrument, then provide some qualitative discussion of its implementability. Note in each case that expected utility is constant across types and own type does not directly enter into the utility function, so it is without loss of generality for us to ignore misreporting issues –incentive compatibility constraints will be satisfied for every agent in all of the following specifications.

Lemma 1. *Total surplus has a well defined maximum in the space of cutoffs* $\Psi \equiv \{\psi \in R^4\}$.

Proof. Appendix. □

Lemma 2. *Suppose that* i *is the inflow rate of Studs onto the platform. Then,*

$$iEU_h + EU_l = i\,\frac{o_h h + \alpha \int_{\psi_{hh}}^{\infty} x\, f(x)\, dx + (1-\alpha)\int_{\psi_{hl}}^{\infty} x\, f(x)\, dx - s}{\frac{r}{\beta} + o_h} +$$

$$\frac{o_l l + \alpha \int_{\psi_{hl}}^{\infty} x\, f(x)\, dx + (1-\alpha)\int_{\psi_{ll}}^{\infty} x\, f(x)\, dx - s}{\frac{r}{\beta} + o_l} = iEV_h + EV_l$$

Proof. Appendix. □

Lemma 3. *Any* ψ *generates an identical set of accepted matches to some* ψ' *of the form* $\{\psi_{hh}, \psi'_{hl}, \psi'_{hl}, \psi_{ll}\}$.

Proof. Appendix. □

Lemma 4. *First best* ψ^* *with censored search is invariant to search costs.*

Proof. Appendix. □

5.1 Pricing

We'll first consider the exclusive use of price instruments, starting with the standard two-part tariffs and moving on to a more complex pricing strategy.

For the two-part tariffs, we first present the non-distortionary pricing model in which there are no externalities. In particular, agents are identical. In the second subsection, we study the cases where there are two types of vertically-differentiated agents. In the third subsection, we discuss some impediments to separating the types into distinct platforms, and focus on two cases for the idiosyncratic matching shock ψ: Case I, where there are small idiosyncratic shocks, and Case II, where ψ is distributed along a uniform distribution on $[0, m]$. All cases considered in this section are for when agents' types are perfectly observable by the monopolist.

The assumption of perfect observability of types does not match what actually happens in real-life; this being said, it allows for tractability. This assumption can be relaxed and we refer the reader to the online appendixes for this paper for further research on the case where types are unobservable, and cases where we consider ψ to be distributed according to a smooth distribution and also a two-point distribution.[9]

[9] This may be found at http://www.melatinungsari.com/research.html.

5.1.1 Two-Part Tariffs in a Model With No Externalities

Consider the baseline model where there is only one vertical type and perfectly correlated matching shocks ψ. This model with no externalities provide a baseline for the more complicated model later on, where externalities are introduced through the α term and the tradeoff between choosing to match more aggressively on the homogenous versus heterogeneous dimension. To solve for the steady-state equilibrium, note that the monopolist acts a social planner who maximizes total surplus in the economy. In particular, the monopolist will set prices to maximize the agent's expected discounted utility of being on the platform. Now, note that the only way prices affect the agent's problem in the case with no externalities is through the agent's threshold ψ^*. For ease of exposition, since utility is additively separable, write the expected discounted utility of agents as

$$EU(\psi^*) \equiv U - f - p\, D(\psi *)$$

where U is the benefit from matching, D is the expected number of dates the agent goes on before leaving the platform, p is the per-interaction price charged by the platform, and f is the fixed price. Recognizing that prices affect the threshold ψ^*, we take first-order conditions with respect to ψ^* to obtain

$$\frac{\partial EU(\psi^*)}{\partial \psi *} = 0 \implies \frac{\partial U}{\partial \psi^*} = p \frac{\partial D(\psi *)}{\partial \psi *} \implies p = 0$$

Now, note that the monopolist profit function is given by

$$\Pi = 2\{if + \sum_{t=0}^{\infty} \delta^{t+1} ip(1-o)^t\} \tag{14}$$

The profit function linearly increases in i, which implies that the monopolist optimally sets $i = 1$.[10] This fully characterizes the solution to the monopolist's optimization problem for the case where there are no externalities.

Proposition 6. *In the model with no externalities and a differentiable F, the monopolist optimally sets per-interaction fees to be zero, lets all agents in, and sets the fixed fee to extract all remaining surplus.*

Proof. Appendix. □

This result provides a benchmark for the case where there are no distortions caused by externalities imposed by different types of agents on each other. In particular, the firm in acting as a social planner does not have to correct the matching behavior of agents on the platform by

[10] The profit function takes into account the cohort that entered the platform in steady state and follows them throughout the entirety of their lives. This is in contrast to a similar (but incorrect way) of calculating profit, which is by calculating the sum of profits that is able to be extracted from the mass of agents *in* steady-state, regardless of the when they entered the platform.

charging a non-zero per-interaction price since agent's optimality conditions are already aligned with those of society.

5.1.2 Two-Part Tariffs in a Model With Externalities

Consider now a model where there are two types of agents, namely the studs (H types) and the duds (L types). We will analyze the case in which the firm is not able to separate the two types of agents into distinct platforms. The firm may not be able to separate the two types of agents due to a variety of reasons. For example, an online dating firm may specialize and operate only a single website, focusing on aggressively advertising and managing the website. The firm may also face high fixed and operating costs to operate more than one website, or may be faced with legal constraints that prevent them from segregating the two types.

Another explanation as to why there could be two different platform configurations may simply be because of the matching technology employed by the firm. To further illustrate this idea, consider the real-life online dating platform eHarmony. This dating platform operates with a different matching technology than other online dating platforms such as Match. The reason that eHarmony is different than Match because eHarmony makes all of their users undergo a 'personality quiz' in which the users answer a long list of questions designed to extract information about their characteristics and essentially, type. eHarmony then able to isolate a high type from a low one, and *limit the user's choice set* by only offering them access to agents who are similar to themselves in types. Abstracting away from these details, we can model this as a two-platform model where the firm fully observes the types of all agents and splits the population into two separate platforms that cater to specific types.

This is in contrast to the matching technology employed by Match, where the website does not try and extract more information about the types of the agents on the platform. In particular, the firm allows agents to match freely with anyone while on the platform. The question now is whether or not the monopolist would in fact, in the case where they are able to, operate two platforms instead of one. As proven in a later section, the answer to this question is yes, under some unrestrictive assumptions on the distribution function $F(.)$ of the idiosyncratic match shock ψ. This result is intuitive but not trivial. In particular, the analysis is complicated by the fact that the two platform configurations could have different outflow rates associated with the agents leaving the platforms, and the fact that the proportion of studs on the platform, α, enters the analysis for the one platform type. Our result proves that under this model setup and assuming perfectly observable types, the monopolist does strictly better by splitting the types into separate platforms following the strategy of eHarmony than operating one platform like Match and letting all types freely search and match.

In solving for equilibrium, we focus on four cases for the idiosyncratic matching shock ψ. The first case is the basic case with no idiosyncratic shocks. We then extend this to the case where the idiosyncratic matching shocks are sufficiently 'small'; that is when the shocks ψ are small relative to the size of the types θ_H and θ_L. In this case, we are able to explicitly solve for the steady-state equilibrium. The second case that we study is when ψ is uniformly distributed on the interval $[0, m]$. In this case, we obtain some partial equilibrium results and also run numerical simulations to solve for a full steady-state equilibrium.[11] Before proceeding to these cases, we will first outline the monopolist's problem in all entirety.

Monopolist's Problem

We assume that the monopolist *fully commits* to a constant path of prices throughout time, charging $\{p^* = (p^*_H, p^*_L), f^* = (f^*_H, f^*_L)\}$ in each period.

Definition 6 (Monopolist Profit)**.** *The firm's expected profit in steady-state is*

$$\Pi_{ss} \equiv 2\left\{i_H f_H + i_L f_L + \sum_{t=0}^{\infty} \delta^{t+1}(p_H i_H (1-o_H)^t + p_L i_L (1-o_L)^t)\right\}$$

This function can be thought of as the profit obtained by following the group of agents of type H and L that entered in the steady-state until the end of their lives.[12]

By evaluating the summation, we simplify the expression for the steady state profit to obtain

$$\Pi_{\text{ss}} = 2\{f_H i_H + f_L i_L + \frac{\delta i_H p_H}{1-\delta(1-o_H(p_H))} + \frac{\delta i_L p_L}{1-\delta(1-o_L(p_H,p_L))}\} \tag{15}$$

Given that types are perfectly observable, the monopolist is only concerned about the *individual rationality* (participation) constraints. The optimization problem for the monopolist is defined as follows.

$$\max_{p_H, p_L, f_H, f_L} \Pi_{\text{ss}}(\alpha(p_H,p_L), p_H, p_L)$$
$$\text{s.t.} \quad (IR_H) \quad C_H(\alpha(p_H,p_L), p_H) - f_H \geq 0$$
$$(IR_L) \quad C_L(\alpha(p_H,p_L), p_H, p_L) - f_L \geq 0$$

subject to the steady-state conditions in Definition 3.

The pricing problem is complicated by the fact that changes in prices affect not only one group of agents, but also the interactions between the group and the other through the steady-state α. However, since types are observable, the firm will extract all surplus from the agents, as proven

[11] This provides a preview of the results for when F is distributed according to a two-point mass distribution; that is, ψ is either $\gamma > 0$ with probability $1-q$ or 0 with probability q. Here, we are not able to provide any analytical results, but we provide numerical simulations for solving for equilibrium. The final case is when ψ is distributed according to any smooth distribution. We are able to obtain analytical results for this case, where the results are slightly more restrictive than that of when ψ is uniformly distributed.

[12] As mentioned in a previous footnote, this is in contrast to another possible (but incorrect) way of expressing the profit, which would be to take into account *all agents on the platform* in steady-state, regardless of the time period in which they joined the platform:

$$\Pi = 2\sum_{t=0}^{\infty} \delta^t \sum_{j \in \{L,H\}} \{f_j\, i_{j,t} + \mu_{j,t}\, p_j\}$$

below. This implies that the pricing problem is equivalent to the planning problem in which the firm maximizes the total surplus of the agents on the platform, only to charge prices to extract it all from them.

Proposition 7. *In any steady state equilibrium with observable types, firm sets fixed prices $f = \{f_H, f_L\}$ and per-interaction prices $p = (p_H, p_L)$ so that IR_H and IR_L bind.*

Proof. Appendix. □

Case I: No to Small Idiosyncratic Matching Shocks
Suppose, to begin, that there are no idiosyncratic shocks in this model. An agent of type i obtains the following utility when she chooses to match with agent j:

$$u_i(j) = \frac{\theta_j}{1-\delta}$$

An agent obtains the type of her match forever, which is discounted by the discount rate of δ. Agents are risk-neutral and share the same discount rate with the firm. We can think of there being two steady-state equilibrium: a *pooling* equilibrium where H and L types match not only amongst themselves but also with each other, and a *separating* equilibrium with H types only match with other H types, and L types only with other L types.

Since types are perfectly observable, the profit-maximizing monopolist will extract all surplus from the agents by using a combination of per-interaction prices $p = (p_H, p_L)$ and fixed fees $f = (f_H, f_L)$. All agents have the same outside option of obtaining zero utility when not on the platform. The individual rationality constraints for each type will bind, giving

$$C_H - f_H = 0$$

$$C_L - f_L = 0$$

Since types are perfectly observable, we note that solving the monopolist problem is equivalent to solving the social planner's problem in which total social surplus is maximized. An agent's strategy under this setup is defined as

$$s_i^* = \begin{cases} \text{match with } H \\ \text{match with } L \text{ and } H \end{cases} \tag{16}$$

In equilibrium, since H determines whether or not a match forms, L's only strategy will be whether or not she will match with another L.

Under this setup, the total social surplus will be a function of all the types of agents on the platform. For concreteness, suppose that there are only 4 agents in the market, with types H, H, L, L. There are only two cases—the first in which $H - H$ match and $L - L$ match, and the other where two $H - L$ matches form. In the first case, the total amount of surplus that may be extracted is $2(\frac{\theta_H}{1-\delta}) + 2(\frac{\theta_L}{1-\delta})$. In the second case, the total amount is surplus that may be extracted is $2(\frac{\theta_H}{1-\delta} + \frac{\theta_L}{1-\delta})$. It is clear then that the total surplus that may be extracted by the monopolist is the same for both cases; this implies that the monopolist does not care about who matches with whom, but only about the total surplus that may be extracted. This result obtains from the fact that the matching utility is additively separable (modular). Another important fact to note is that an agent's contribution to the total social surplus is his or her own type. This is due to the fact that since an agent obtains his or her partner's type in a marriage, each agent then contributes his or her own type to the surplus.

We can now define the total social surplus on the platform. This is given by

$$\text{Total Social Surplus (TSS)} = \mu_H \sum_{t=0}^{\infty} \frac{\theta_H}{1-\delta}\delta^t(1-o_H)^t i_H + \mu_L \sum_{t=0}^{\infty} \frac{\theta_L}{1-\delta}\delta^t(1-o_L)^t i_L \tag{17}$$

Note that prices are implicit in the expression of total social surplus. I can ignore explicit expressions of prices since for any given expected discounted lifetime utility of being on the platform can be *fully* extracted by the monopolist given a combination of prices. The prices only enter in the expression of total social surplus in the outflows o_H and o_L, where higher prices would cause agents to be less picky and so, leave the platform more quickly.

Now, we focus on the simplifying the expression for TSS by looking at optimal inflows. Since all agents agree on the ranking of the types, H types are considered desirable by all other agents on the platform. Thus, i_H is optimally set to be 1. We need only characterize what the optimal i_L is. In the following proposition, we prove a strong and interesting result about this model, namely that given the setup and *any* range of parameters, the monopolist will price per-interaction prices to extract all surplus from the agents and make them leave *in the first period.* Also, given any range of parameters, the monopolist will choose to let all the L types in, regardless of the value of θ_L.

Proposition 8. *Consider the model with no idiosyncratic matching shocks. In the steady-state equilibrium, the profit-maximizing firm will*

1. *Set* $i_H = i_L = 1$.
2. *Charge* p_H, p_L *sufficiently high to induce all agents to leave in the first period that they are on the platform.*

Proof. Appendix. □

There is an intuitive interpretation for this result in terms of the discount rate δ. Consider three cases for the discount rate: $\delta = 1, \delta = 0$, and $\delta \in (0, 1)$. In the case where $\delta = 1$, any pricing strategy that the firm sets will give the same TSS, since agents essentially weight each period the same. In the case where $\delta = 0$, the model reduces to that of a one-period model, where it is obviously optimal for the profit-maximizing monopolist to extract all surplus by charging the highest possible per-interaction prices. For any $\delta \in (0, 1)$, note that TSS is strictly decreasing in the the number of periods that the agents stay on the platform. This being, the pricing strategy that forces everyone to leave immediately gives a higher TSS than any other pricing strategy.

Thus, the optimal pricing strategy for the firm is to set per-interaction prices sufficiently high to induce all agents to leave in the first period.

Now, we would like to make strides towards understanding the model where the idiosyncratic shock is sufficiently small. We have already shown what the steady-state equilibrium looks like in the context of no idiosyncratic shocks. What happens when these shocks are small, i.e. when the *support* of the distribution of ψ is small? It turns out that the result for when the shocks are small relative to the size of the types (θ_H, θ_L) is the same as in the case where there are no shocks for a range of parameters that occurs with positive probability. A sufficient but not necessary condition for the shock ψ being small is $\frac{\theta_L}{m} \geq \frac{\delta}{1-\delta}$.

Proposition 9. *Suppose $\psi \in [0, m]$. Then for parameterizations such that $\frac{\theta_L}{m} \geq \frac{\delta}{1-\delta}$, the profit-maximizing monopolist*

1. *Sets $i_H = i_L = 1$.*

2. *Charge p_H, p_L sufficiently high to induce all agents to leave in the first period that they are on the platform.*

Proof. Appendix. □

We have now shown that for the set of parameters $\{\theta_L, \delta, m\}$ that satisfy $\frac{\theta_L}{m} \geq \frac{\delta}{1-\delta}$, the monopolist firm will set per-interaction prices sufficiently high to force all agents to match immediately in the first period. In particular, for relatively small values of δ, the condition on the parameter space is satisfied.

The intuition for these results is as follows. The monopolist only cares about the total amount of social surplus in the economy. The total social surplus declines as time goes on and agents are still on the platform by a factor of δ in each period. This being, the monopolist will want to force agents to match immediately while on the platform *unless* there is a non-zero probability of agents obtaining a sufficiently high draw ψ, which will improve their utility and so increase the amount of surplus that can be extracted from them. In the case where the idiosyncratic matching shock ψ is small relative to the size of the types θ_H and θ_L, the incentives for the monopolist to let agents stay on the platform for long are corresponding small. In particular, with both agents and the monopolist valuing the future very little (implying low values of δ), there is no incentive for either side to stay longer on the platform when the benefit of waiting on the platform for a better ψ is small.

One natural question to ask now is whether the firm charges different per-interaction and fixed prices to the studs and duds under the parameter range where the pricing strategy of charging sufficiently high prices to induce agents to match immediately is implemented. What is observed is that since the monopolist firm lets all agents in to the platform and outflow rates for both studs and duds are exactly 1, the probability of meeting a stud or dud is exactly (denoted α) is also exactly $\frac{1}{2}$. Since the chance of meeting a stud is the same for both types of agents, and per-interaction prices are set to induce both agents' acceptance thresholds to be 0 (agents marry their first match), agents expected lifetime utility of being on the platform is the same regardless of type. Thus, we obtain the following result that the profit-maximizing monopolist charges the same fixed prices to both H and L types, despite the fact that types are observable.

Proposition 10. *Suppose* $\frac{\theta_L}{m} \geq \frac{\delta}{1-\delta}$ *and types are perfectly observable. In the steady-state equilibrium, the profit-maximizing monopolist sets sufficiently high per-interaction prices* p_H^*, p_L^* *so that all agents match and exit in the first period and* $f_H^* = f_L^*$.

Proof. Appendix. □

One fact to note is that although the fixed prices are the same for both types of agents, the optimal per-interaction prices may or may not be the same. In particular, the only restriction on per-interaction prices is that they have to be sufficiently higher than a certain threshold. After they reach the threshold, whether or not the per-interaction prices for studs and duds are the same is inconsequential. However, since agents do not actually pay the per-interaction prices in the first period that they are searching on the platform if they marry in the first period, this optimal contract essentially does not differentiate between types since it only consists of fixed prices. The per-interaction prices for both types are just set to provide strong incentives (i.e. force you) to immediately marry the first person you go on a date with.

Case II: The Idiosyncratic Matching Shock ψ is Uniformly Distributed on $[0, m]$

Now, we study the case when the random matching shock ψ is draw from a uniform distribution on $[0, m]$. In particular, we would like to elaborate on the assumptions that the following proposition makes. The facts to note are the following:

1. In studying the case where the shocks are uniformly distributed, we fix the pricing strategy for the H types. This implies that the analysis done is a partial equilibrium analysis. Fixing the pricing strategy for the H types is necessary to keep the H type optimality conditions the same. This is because by fixing the pricing strategy, we are able simplify the problem greatly due to the complexities that arise from the fixed point analysis with the proportion of studs on the platform, α.

2. We assume that $i_L \in (0, 1)$; that is, that the monopolist is letting in a nonzero proportion of the L types. In practice, to check whether or not the solution we obtain is optimal, we may compare the profit obtained from the equilibrium in the proposition with the one where the firm lets no duds in (i.e. $i_L = 0$).

3. We also assume that for a fixed p_L that $\psi_L^*(L) \in (0, m)$. This is to ensure that prices have a *bite*; that is, changing prices will actually affect how the L type optimizes. For example, if the threshold were at 0, L types are already choosing to always match with other L types and so prices have no effect on the threshold. The duds would ideally like to lower their threshold with an increased p_L but are unable to do so. In the case where $\psi_L^*(L) = m$, L types never match with other L types anyway, and so prices have no effect on the thresholds.

Proposition 11. *Suppose $0 < i_L < 1$, fix i_H, p_H and p_L, suppose F is the uniform distribution with support $[0, m]$, and $\psi_L^*(L) \in (0, m)$. The firm can strictly increase profits by increasing both i_L and p_L by $\epsilon > 0$.*

Proof. Appendix. □

Corollary 2. *Suppose the conditions in the previous proposition hold. In equilibrium, if the firm chooses to let any L types onto the platform, it will let all of them in (i.e. set $i_L = 1$).*

Proof. Appendix. □

Note that this corollary *does not* state whether it is always optimal to let any L types in (i.e. $i_L > 0$) or whether $p_L > 0$. In particular, it may be that if $\theta_H - \theta_L$ is very large (or rather if the L type is very 'weak'), then the optimal choice for i_L may be to set $i_L = 0$ to eliminate the negative externality that L imposes on H.

Despite the small number of primitives in this model, analytical results are difficult to obtain. This is primarily due to two factors, the first being that the the effects of perturbations in the model depend on the exact specification of the distribution function of ψ, $F(.)$. The second reason is because of the fixed point nature of the equilibrium. In particular, even though agents take the proportion of high types on the platform, α, as given when optimizing, in equilibrium their thresholds will determine α. In particular, since prices affect thresholds, prices will then also affect α through both the thresholds of the agents and the outflow rates for both types. In assuming a particular form for F, we are able to conduct simulations for a range of parameters in the model that illustrate the main points of the model without getting into the complications of striving for analytical results. Next, by fixing the values for θ_L, δ, and m, we run simulations that solve for the steady-state equilibrium for varying values of θ_H. As we shall see, the results from the simulation agree with the analytical results that we have presented here.

Implementation:

5.1.3 Match-Dependent Pricing

We now consider a more complex pricing strategy: match-dependent pricing. By this we mean that, in addition to quality dependent fixed fees and per-interaction prices, the platform may charge a distinct price upon matching for each potential pairing. In practice, the platform need only set two such prices–one for matching to a high quality partner when you are high quality, and one for matching to a high quality partner when you are low quality[13]. With this flexibility, the platform can essentially transform the NTU environment into a TU environment by providing the transfers that agents can't make themselves. Shimer and Smith (2001) show that, in a decentralized search-and-matching setting with heterogeneous agents, TU is not enough by itself to achieve first best. However, they show that the addition of per-interaction prices is sufficient to achieve first best in their model, and so it is with ours.

[13]Generalizing to n quality levels, at most $(n-1)n$ such prices will be necessary.

The intuition is as follows: the platform wants to implement the three optimal cutoffs ψ_{hh}, ψ_{hl}, and ψ_{ll}. The difference between an agent's cutoffs when matching to Studs and Duds is exactly the difference in match quality, since the agent's optimization is based on their continuation value, which is independent of the current draw. By adding an additional cost to matching to high types that gap will be increased (or decreased) by precisely the additional cost. Then, any $\psi_{hh} - \psi_{hl}$ and $\psi_{hl} - \psi_{ll}$ can be induced via match dependent pricing. Given those cutoff differences, appropriate per-interaction prices can be chosen to implement the desired lower cutoffs–ψ_{hl} and ψ_{ll}–giving us the complete cutoff scheme. Fixed fees soak up residual surplus in the monopoly case and ensure IR constraints are satisfied in the planner's case. Define ϕ_q as a price charged to agents of quality q upon matching to a high quality partner. Then the matching utility for agent i is $u(i,j) = \psi_{ij} + \theta_j - \phi_{q_i} I_{\theta_j = h}$.

Proposition 12. *There exist prices (f, p, ϕ) maximizing TS.*

Proof. Appendix. □

Implementation: This sort of aggressive price discrimination may be difficult to implement in practice as it would involve not just differing prices for different users, but match dependent prices that would make it more costly to match to certain types of partners, discriminating against them along two dimensions. By construction, this price discrimination must be transparent–it is exactly the price differentials between different types of matches that induce optimal behavior. This makes it even more likely to irk users. This sort of pricing is also likely to be illegal in many jurisdictions. Finally, in a more general case with more than two quality grades, the number of prices to be determined increases quadratically, and, rather than setting cutoffs directly, the platform must simultaneously determine this menu of prices to induce the appropriate cutoffs. Thus, this instrument is technically challenging to implement as well.

5.2 Splitting Platforms

The strategies we've considered thus far have only utilized prices, but a platform may also use the structure of the market itself as an instrument to improve efficiency. An obvious and extremely simple approach is to split the platforms along the quality dimension, creating one market for Studs, and another for Duds. We find that, with CRM–that is, when platforms have constant returns to scale–splitting the platforms can always improve total surplus. The intuition is simple: agents match too much on the quality trait, and therefore reject good fit, low quality matches they should accept ("type I errors") and accept bad fit, high quality matches they should reject ("type II errors"). Segregating agents by quality avoids this temptation, and with CRM the rate of draws is the same as on a mixed platform, making split platforms more efficient. This mirrors a similar result from Damiano and Li (2007), but for an entirely different reason. In Damiano and Li (2007), forcing assortation via platform partition is optimal because agents have supermodular utility, so matching high types with high types is more efficient than mixed matching–without partition, agents don't match closely enough on quality. In our case, however, partition prevents agents from matching too aggressively on quality to the detriment of fit. However, we show that this result depends strongly on the CRM assumption. If there are increasing returns to scale in platform size (LRM), it's always possible to generate more surplus on a single, mixed-quality platform because of the increased market thickness. Define TS_m^* as the maximum total surplus for mixed platforms, and TS_s^* as the maximum for split platforms.

Proposition 13. *Given CRM, there exists cutoffs ψ such that $TS_s(\psi) \geq TS_m^*$.*

Proof. Appendix. ☐

We've shown that separate platforms can generate more total surplus with appropriate cutoffs, but we haven't shown that the decentralized outcome induces them. Clearly, any set of instruments that can achieve first best can achieve these cutoffs, but the following result shows that even two-part tariffs can do so:

Proposition 14. *There exist prices (f, p) inducing $TS_s^* > TS_m^*$.*

Proof. Appendix. ☐

Clearly, when CRM holds and there are no direct benefits to having a larger platform, a planner or monopolist should operate separate platforms for Studs and Duds. However, this does not hold under LRM. When the rate of draws increases in platform size, it will generally be preferable to aggregate users onto a single platform, assuming their matching behavior can be appropriately influenced. Specifically, with linear returns to matching, the rate of draws is proportional to the mass of agents on the platform–a mixed platform can have the same rate of Stud draws as a Studs-only platform while still having the same rate of Dud draws as a Duds-only platform. Therefore, if we assume search costs are low (s=0)[14], a mixed platform has no disadvantages relative to split platforms–even if Studs and Duds refuse to match to one another, they'll do just as well, and find matches just as fast, as if they were on separate platforms. The intuition of the proof comes from this insight: we can exactly replicate the total surplus of split platforms with a mixed platform where there are no cross-quality matches. From there, if we raise the Stud-Stud and Dud-Dud cutoffs, eliminating some lower fit matches, and add a proportional mass of Stud-Dud matches at the top of the fit distribution, we improve expected fit while holding expected time to match constant, strictly improving total surplus.

Proposition 15. *Given LRM and $s = 0$, $\exists$ cutoffs ψ such that $TS_m(\psi) > TS_s^*$.*

Proof. Appendix. ☐

Implementation: This is a comparatively easy strategy to implement, though with a continuum of quality the optimal number of platforms may be quite large (infinite with CRM, though that assumption will inevitably break down at sufficient granularity.)

[14]Gaining economies of scale from a single, large platform relies on giving users more draws and allowing them to be pickier. Search costs undermine that strategy, and the optimality of a single platform would depend on a comparison of the relative costs of search and the benefits of improved matching on fit.

5.3 Censored Search

We now consider a second non-price instrument, censoring the draws an agent can observe. Curating agent choice sets is ubiquitous on online matching platforms – there are typically more partners available than can be presented at one time. Platforms will often limit the observable set of partners to those it considers a good match. At the very least, a platform will sort the set of potential matches, showing some earlier than others. Regardless, this can be framed as censoring, at least temporarily, some partners. Aside from lowering search costs and acting as an matching expert, our analysis thus far illuminates a third reason a platform may pursue this strategy: eliminating matching externalities. However, this instrument's reach is limited–it can't directly induce agents to accept matches they don't want. Censored search can eliminate "type II errors", but can't directly eliminate "type I errors". The cost of this limitation becomes obvious when we consider the externalities we've discussed thus far: generally, the platform wants Studs to be less picky when matching to Duds in order to accept good fit matches. An obvious approach using censored search would be to censor the lower fit Stud-Stud matches, forcing Studs to be less selective with Duds. However, direct inspection of the Stud's continuation value – which is also their expected utility from the platform – and corresponding indifference condition for a reservation value show that the cutoff for accepting a Dud is own continuation value plus a constant–to decrease a Stud's cutoff for Duds by x, the platform must lower their expected utility by x. In special cases where there is a mass point of Duds just below the Stud's cutoff, slightly lowering the expected utility for Studs to get Duds better matches can be optimal, but this is extremely sensitive to the fit distribution and the parameters of the model. Note that every draw that is censored generates no search costs, and incurs no per-interaction price – these are potential draws that are never actually realized.

We'll now illustrate how, under the right circumstances, a platform can utilize censored search to improve total surplus. The benefit of lowering the Studs' lower cutoff is proportional to the mass of Duds who are now accepted by Studs that would have previously rejected them, so a Bernoulli fit distribution will best illustrate the potential use of the instrument. Formally, $\psi_{ij} \sim Bern(p)$. We'll first establish a few properties of the decentralized equilibrium with binary fit, then find parameter ranges where censoring bad fit Stud-Stud matches is surplus improving. For concision, we'll denote a match by $\theta_i\theta_j\psi_{ij}$. To avoid the multiplicity of cases and equilbria that can occur in this setting, we'll focus on a case where $l = 0$, and $i = 1$. We'll assume $s = 0$ to simplify the interpretation of our results – if s is positive, the platform can benefit from the more obvious advantage of censored search: eliminating search costs by censoring draws that will be rejected anyway. With $s = 0$, that channel is shut down and the platform can only utilize censored search to change users' matching outcomes. We'll also assume that indifferent matches are always accepted.

Lemma 5. *In the decentralized equilibrium, $hh1$ and $ll1$ must be accepted, while $hl0$ and $ll0$ are always rejected.*

Proof. Appendix. □

Proposition 16. *There exist parameters such that censoring hh1 draws improves total surplus.*

Proof. Appendix. □

Proposition 17. *For sufficiently high* $\frac{h}{r}$*, censoring* $hh1$ *draws cannot improve total surplus.*

Proof. Appendix. □

5.3.1 Adding per-interaction costs

It seems that, by itself, censored search is of ambiguous value. There are cases where it can improve total surplus, but they are very sensitive to the parameters of the model. However, with the addition of per-interaction prices, censored search is a much more powerful tool In particular, positive per-interaction costs act as additional search costs, forcing agents to be less selective. This eliminates "type I errors" at the cost of generating more "type II errors". Since censored search can eliminate the latter but not the former, this gives censored search more traction. With high enough per-interaction costs, all "type I errors" are eliminated, and the platform can achieve first best.

Proposition 18. *There exist prices* $(\mathbf{f}, p\mathbf{1})$ *and censoring cutoffs* ψ_c *maximizing total surplus.*

Proof. Appendix. □

Implementation: Unlike the pure pricing instruments, censored search only requires that the platform be able to identify "bad" matches–they don't need to devise a menu of prices to indirectly induce ideal matching. Even when adding a per-interaction cost, the platform need only ensure that it is sufficiently large. This suggests that censored search can both be extremely powerful and easily implementable. Even if extremely high per-interaction prices are infeasible in practice, any increase in per-interaction costs gives censored search more traction, illustrating the complementarity of these two instruments.

6 Conclusion

This paper presents a multidimensional matching model with search where economic agents are differentiated both in terms of *quality* and *fit*. In presenting a more complex modeling of preferences, we highlight a previously unstudied source externalities – the tradeoff between quality and fit. We study how a strategic platform can leverage a variety of instruments to improve surplus on the platform, and find the combination of censored search and two-part tariffs particularly effective in our setting. Our model is broad, covering constant and increasing returns in the matching technology as well as both of the common specifications for search frictions – search costs and time discounting. We also include extensions to frictionless matching. However, our model is extremely simple, and thus suggests a variety of directions for further research. One promising direction is competition – our model assumes a monopoly, but many matching platforms operate in a competitive setting. Modular utility simplifies our analysis greatly, but supermodularity may be more realistic for many applications, and may change the welfare analysis for some instruments. Finally, we assume a setting where the platform can utilize fixed fees to perfectly price discriminate. This, along with modular utility, ensures that incentive compatability constraints will always be satisfied. Relaxing these assumptions would allow a rich analysis of the constraints incentive compatibility may place on a strategic matching platform in this environment.

Appendix

Proposition 5:

Proof. Total surplus is then

$$\sum_{i=1}^{n} u_i = \sum_{i=1}^{n}(\sum_{j=1}^{k}(\theta^m_{i,j} + \theta^w_{\mu(i),j}) + f(\psi^m_i, \psi^w_{\mu(i)})) = \sum_{i=1}^{n}(\sum_{j=1}^{k}(\theta^m_{i,j} + \theta^w_{i,j}) + f(\psi^m_i, \psi^w_i))$$

for any μ. For assignments μ and η, let TSS_μ and TSS_η and denote the overall total social surplus generated from the assignments (i.e., the total social surplus from both the vertical and heterogenous dimensions; for example, $\text{TSS}_\mu = \text{TSS}_{\mu,\theta} + \text{TSS}_{\mu,\psi}$). Then,

$$\text{TSS}_\mu - \text{TSS}_\eta = \text{TSS}_{\mu,\theta} - \text{TSS}_{\eta,\theta} + \text{TSS}_{\mu,\psi} - \text{TSS}_{\eta,\psi} = \text{TSS}_{\mu,\psi} - \text{TSS}_{\eta,\psi}$$

and, since transferable utility stable matchings maximize TSS by Shapley and Shubik (1971), the transferable stable matching μ^* must satisfy

$$\mu^* = \operatorname*{argmax}_\mu \text{TSS} + \operatorname*{argmax}_\mu \text{TSS}_{\mu,\psi}$$

□

Lemma 2:

Proof. For each match, undiscounted contribution to total surplus is $m = 2\psi_{ij} + \theta_i + \theta_j$, where $u_i = \psi_{ij} + \theta_j$ and $u_j = \psi_{ij} + \theta_i$. Defining $v_i = \psi_{ij} + \theta_i$ and $v_j = \psi_{ij} + \theta_j$, we have $m = v_i + v_j$. Then, replacing u_i with $v_i$and u_j with v_j and computing the continuation value as before, we have $\frac{\alpha\int_{\psi_{hh}}^{\infty}(h+x)f(x)dx+(1-\alpha)\int_{\psi_{hl}}^{\infty}(h+x)f(x)dx-s}{r/\beta+o_h} + \frac{\alpha\int_{\psi_{hl}}^{\infty}(l+x)f(x)dx+(1-\alpha)\int_{\psi_{ll}}^{\infty}(l+x)f(x)dx-s}{r/\beta+o_l}$. The result follows immediately. □

Lemma 3:

Proof. Let $\psi^{'}_{hl} = max\{\psi_{hl}, \psi_{lh}\}$. A Stud-Stud or Dud-Dud match is accepted under ψ' if is accepted under ψ, as the cutoffs are identical. A match between a Stud and a Dud is accepted under ψ iff $\psi_{ij} \geq max\{\psi_{hl}, \psi_{lh}\}$. Then, given that a match is accepted under ψ' iff $\psi_{ij} \geq \psi^{'}_{hl} = max\{\psi_{hl}, \psi_{lh}\}$, the result is proved. □

Lemma 4:

Proof. Suppose a draw will be rejected by one of the agents. Then the platform can improve utility by censoring the draw. Then every uncensored draw will be accepted, and search cost per agent will be s. s is constant in ψ, so ψ^* is invariant to s.

□

Proposition 8:

Proof. Consider the strategy where p_H, p_L is significantly high to induce all agents to leave the platform immediately. In this strategy, the corresponding outflow rates for both types will be exactly 1. Given that $o_H = o_L = 1$, we then note that the only way in which changing i_L affects TSS is by affecting how quickly the agents match, which is expressed through the outflow rates. Given that the outflow rates are held constant at 1, under this pricing strategy, the TSS becomes an increasing linear function in i_L, implying that the firm optimally sets $i_L = 1$. Now, all that remains to be shown is the optimality of charging prices that make all agents leave immediately.

By evaluating the summation in (17), I can write the expression as

$$\text{TSS} = \mu_H\left(\frac{i_H\theta_H}{(1-\delta)(1-\delta(1-o_H))}\right) + \mu_L\left(\frac{i_L\theta_L}{(1-\delta)(1-\delta(1-o_H))}\right) \tag{18}$$

By taking derivatives, I note that

$$\frac{\partial \text{TSS}}{\partial o_H}, \frac{\partial \text{TSS}}{\partial o_L} \geq 0$$

which implies that the optimal pricing strategy is then to set $o_L = o_H = 1$. Under this pricing strategy, we have shown earlier that the corresponding optimal inflows are $i_L = i_H = 1$.

□

Proposition 9:

Proof. Consider again the expression for total social surplus. Since the matching utility is additively separable, we may write the total social surplus as

$$\text{Total Social Surplus (TSS)} = \sum_{j\in\{L,H\}} \mu_j \sum_{t=0}^{\infty} \Big(\frac{\theta_j}{1-\delta} + \frac{E(\psi \mid j \text{ matches})}{1-\delta}\Big)\delta^t(1-o_j)^t i_j \tag{19}$$

where μ_j is the mass of type θ_j on the platform. By the same reasoning as in the previous proposition, given the pricing strategy where the firm sets per-interaction prices sufficiently high to force everyone to match in the first period, it is clear that $i_H = i_L = 1$. Now, consider two pricing scenarios:

1. Per-interaction prices are set sufficiently high to force everyone to match in the first period and $o_H = o_L = 1$.

2. Per-interaction prices are set such that $\exists$ some measure λ_H of H agents and λ_L of L agents that wait on the platform and do not immediately match.

By assumption of the bounded support of ψ, the worst draw that any agent can obtain is any period is $\psi = 0$ and the best is $\psi = m$. Thus, we can bound the total social surplus from scenario 1 by the case where agents obtain a shock of $\psi = 0$ in each period:

$$\text{TSS}_1 \geq \mu_H\left(\frac{\theta_H}{1-\delta}\right) + \mu_L\left(\frac{\theta_L}{1-\delta}\right)$$

Note that the proportion $(1-\lambda_H)$ and $(1-\lambda_L)$ of H and L agents, respectively, we can write the total social surplus in scenario 2 as:

$$\text{TSS}_2 = \lambda_H\mu_H \ \text{TSS}_{2,H} + \lambda_L\mu_L \ \text{TSS}_{2,L} + (1-\lambda_H)\mu_H \ \text{TSS}_{1,H} + (1-\lambda_L)\mu_L \ \text{TSS}_{1,L}$$

where $\text{TSS}_{n,j}$ is the total surplus obtained from $j \in \{L, H\}$ in case $n \in \{1, 2\}$.

We now want to show that $\text{TSS}_1 - \text{TSS}_2 \geq 0$ for some parameter range. To do this, we compare the worst case scenario under pricing strategy 1, which is that agents obtain $\psi = 0$ in all subsequent periods, and the best case scenario under pricing strategy 2, which is that agents wait for a period and obtain $\psi = m$ in the following period. Under the best case scenario in the second pricing strategy, agents will leave in the next period since that case is the one where they obtain the most they can while on the platform and the firm then extracts $\delta(\frac{\theta_j+m}{1-\delta})$ from each agent of type $j \in \{L, H\}$.

Taking the difference between the social surpluses gives

$$\begin{aligned}
& \lambda_H\mu_H\left(\tfrac{\theta_H}{1-\delta} - \delta(\tfrac{\theta_H+m}{1-\delta})\right) + \lambda_L\mu_L\left(\tfrac{\theta_L}{1-\delta} - \delta(\tfrac{\theta_L+m}{1-\delta})\right) &\geq& \ 0 \\
\iff \quad & \lambda_H\mu_H\left(\theta_H - \tfrac{\delta m}{1-\delta}\right) + \lambda_L\mu_L\left(\theta_L - \tfrac{\delta m}{1-\delta}\right) &\geq& \ 0 \\
\iff \quad & \theta_L - \tfrac{\delta m}{1-\delta} &\geq& \ 0 \\
\iff \quad & \tfrac{\theta_L}{m} &\geq& \ \tfrac{\delta}{1-\delta}
\end{aligned}$$

□

Proposition 12:

Proof. By Lemmas 1 and 3, there exists $\psi^* = (\psi^*_{hh}, \psi^*_{hl}, \psi^*_{lh}, \psi^*_{ll})$ such that $\psi^*_{lh} = \psi^*_{hl}$ and ψ^* maximizes total surplus. Let $d_h \equiv \psi^*_{hh} - \psi^*_{hl}$ and $d_l \equiv \psi^*_{hl} - \psi^*_{ll}$. Let $d \equiv h - l$. Then set high type match dependent price $\phi_h = d - d_h$ and low type match dependent price $\phi_l = d - d_l$. Then $\psi^*_{hh} - \psi^*_{hl} = \psi_{hh} - \psi_{hl}$ and $\psi^*_{hl} - \psi^*_{ll} = \psi_{hl} - \psi_{ll}$. Now let p_h be such that $\psi^*_{hl} = EU_h[p_h]$ and let p_l be such that $\psi^*_{ll} = EU_l[p_l]$ (solve algebraically). Choose $f_i = EU_i(\psi^*, p, \phi)$. Then $EU_i(\psi^*, p, \phi, f) = 0$.

□

Proposition 13:

Proof. Define ψ^* as a set of cutoffs inducing TS^*_m. Define $\psi_h \equiv F^{-1}(\alpha F(\psi^*_{hh}) + (1-\alpha)(\psi^*_{hl}))$ and $\psi_l \equiv F^{-1}(\alpha F(\psi^*_{hl}) + (1-\alpha)F(\psi^*_{ll}))$.

Then

$$1 - F[\psi_l] = \alpha F(\psi^*_{hl}) + (1-\alpha)(\psi^*_{ll})$$

and

$$1 - F[\psi_h] = \alpha F(\psi^*_{hh}) + (1-\alpha)(\psi^*_{hl})$$

or, equivalently, $o_{ls} = o_{lm}$ and $o_{hs} = o_{hm}$.

Then

$$EV_{lm} - EV_{ls} = \frac{o_l l + \alpha \int_{\psi_{hl}}^{\infty} xf(x)dx + (1-\alpha)\int_{\psi_{ll}}^{\infty} xf(x)dx - s}{r+o_l} - \frac{o_l l + \int_{\psi_l}^{\infty} xf(x)dx + -s}{r+o_l}$$

$$= \frac{\alpha \int_{\psi_{hl}}^{\infty} xf(x)dx + (1-\alpha)\int_{\psi_{ll}}^{\infty} xf(x)dx - \int_{\psi_l}^{\infty} xf(x)dx}{r+o_l}$$

Then it suffices to show that

$$\alpha \int_{\psi_{hl}}^{\infty} xf(x)dx + (1-\alpha)\int_{\psi_{ll}}^{\infty} xf(x)dx \le \int_{\psi_l}^{\infty} xf(x)dx$$

Applying integration by substitution, this is equivalent to

$$\alpha \int_{F[\psi_{hl}]}^{1} F^{-1}[y]dy + (1-\alpha)\int_{F[\psi_{ll}]}^{1} F^{-1}[y]dy \le \int_{F[\psi_l]}^{1} F^{-1}[y]dy$$

Without loss of generality, assume $\psi_{hl} > \psi_{ll}$. Canceling out overlapping regions of integration and doing some algebra, we have

$$(1-\alpha)\int_{F[\psi_{ll}]}^{F[\psi_l]} F^{-1}[y]dy \le \alpha \int_{F[\psi_l]}^{F[\psi_{hl}]} F^{-1}[y]dy$$

Note that, given a well defined CDF F, F^{-1} is an increasing function. Thus,

$$(1-\alpha)\int_{F[\psi_{ll}]}^{F[\psi_l]} F^{-1}[y]dy \le (1-\alpha)\int_{F[\psi_{ll}]}^{F[\psi_l]} \psi_l dy = (1-\alpha)\psi_l(F[\psi_l]-F[\psi_{ll}]) = \alpha(1-\alpha)\psi_l(F(\psi_{hl}^*)-F(\psi_{ll}))$$

and

$$\alpha \int_{F[\psi_l]}^{F[\psi_{hl}]} F^{-1}[y]dy \ge \alpha \int_{F[\psi_l]}^{F[\psi_{hl}]} \psi_l dy = \alpha\psi_l(F[\psi_{hl}] - F[\psi_l]) = \alpha(1-\alpha)\psi_l(F(\psi_{hl}^*) - F(\psi_{ll})).$$

This proves the inequality for low types. An analogous argument holds for high types.

□

Proposition 14:

Proof. By Proposition 13, there exists $TS_s > TS_m^*$, so $TS_s^* > TS_m^*$. Given a (ψ_h^*, ψ_l^*) inducing TS_s^*, let p_h be such that $\psi_h^* = EU_h[p_h]$ and let p_l be such that $\psi_l^* = EU_l[p_l]$. Choose $f_i = EU_i(\psi^*, p, \phi)$. Then (f, p) induces TS_s^*. □

Proposition 15:

Proof. Define ψ^* as a set of cutoffs inducing TS_s^*. Choose ψ_{hl} such that $1-\epsilon < F(\psi_{hl}) < 1$, $\psi_{hl} > \psi_h^*$, and $\psi_{hl} > \psi_l^*$. Then define ψ_{ll} such that

$$1 - F[\psi_l^*] = \alpha/(1-\alpha)(1 - F(\psi_{hl})) + 1 - F(\psi_{ll}). \tag{20}$$

It follows that $\psi_{ll} > \psi_l^*$. We also have that $\alpha \int_{\psi_{hl}}^{\infty} xf(x)dx + (1-\alpha)\int_{\psi_{ll}}^{\infty} xf(x)dx > \int_{\psi_l^*}^{\infty} xf(x)dx$. Similarly, define ψ_{hh} such that

$$1 - F[\psi_h^*] = 1 - F(\psi_{hh}) + (1-\alpha)/\alpha(1 - F(\psi_{hl})). \tag{21}$$

It follows that $\psi_{hh} > \psi_h^*$. We also have that

$$\alpha \int_{\psi_{hh}}^{\infty} xf(x)dx + (1-\alpha)\int_{\psi_{hl}}^{\infty} xf(x)dx > \int_{\psi_h^*}^{\infty} xf(x)dx. \tag{22}$$

From previous equations, we have $o_{lm} = (1-\alpha)o_{ls}$ and $o_{hm} = (1-\alpha)o_{hs}$. We'll now show that expected time discounting is identical for split and mixed platforms: $EV_{lm} = \frac{o_{lm}(l+E_m\psi)}{r/\sqrt{i/o_{hm}+1/o_{lm}}+o_{lm}}$ and $EV_{ls} = \frac{o_{ls}(l+E_s\psi)}{r\sqrt{o_{ls}}+o_{ls}}$. Then time discounting is identical iff

$$1/(\sqrt{i/o_{hm} + 1/o_{lm}}o_{lm}) = 1/\sqrt{o_{ls}}$$

Substituting, we have

$$1/(\sqrt{i/o_{hm} + 1/o_{lm}}o_{lm}) = 1/\sqrt{(1-\alpha)o_{lm}}$$

Equivalently:

$$io_{lm}/o_{hm} + 1 = 1/(1-\alpha)$$

or

$$\alpha = \frac{io_{lm}}{io_{lm} + o_{hm}}$$

But this is always satisfied as it is exactly the definition of α.

Then $EV_{lm} - EV_{ls} = \frac{o_{ls}(l+E_m\psi - l - E_s\psi)}{r\sqrt{o_{ls}}+o_{ls}}$. By a previous equation, this expression must be positive. An analogous argument holds for Studs. Thus $TS_m > TS_s^*$.

□

Lemma 5:

Proof. $hh1$ must be accepted as agents follow cutoff strategies and rejecting hh1 would yield an expected utility of 0, while accepting yields a strictly positive payoff. If $hl1$ is accepted, $ll1$ must also be accepted, as $c_h \geq c_l$. If $hl1$ is rejected, $ll1$ must be accepted by the same argument hh1 must be. $hl0$ and $ll0$ both generate a payoff of zero for at least one agent, while continuation values are strictly positive. □

Proposition 17:

Proof. There are 32 possible strategy profiles in this game–there are 2^4 potential acceptance strategies for Studs and two additional cutoffs for Duds. However, equilbrium requires a cutoff strategy, and 1. and 2. rank the Stud payoff of $hh0$ above $hl1$. Additionally, $hh1$ must be accepted and $hl0$ and $ll0$ must be rejected. Thus, the only possible strategy profiles are $\{hh1, hh0, hl1, ll1\}$, $\{hh1, hh0, ll1\}$, and $\{hh1, ll1\}$. To simplify the comparison between the censored and uncensored search cases, we'd like $\{hh1, hh0, ll1\}$ to be the unique decentralized equilibrium, so we need to find conditions on the parameters to ensure profitable deviations for the other two cases.

For $\{hh1, hh0, hl1, ll1\}$, we have $\alpha = \frac{2p}{p+\sqrt{(8-3p)p}}$, $o_h = \frac{p\left(-p+\sqrt{(8-3p)p}+2\right)}{p+\sqrt{(8-3p)p}}$, and $o_l = p$. This yields $Uh = -\frac{p\left(2h+p+\sqrt{(8-3p)p}\right)}{p^2-p\left(\sqrt{(8-3p)p}+r+2\right)-\sqrt{(8-3p)p}r}$, which we need to be greater than 1 to generate a profitable deviation away from accepting $hl1$.Then $\frac{2p(h+p-1)}{p+\sqrt{(8-3p)p}} > r$.

For $\{hh1, hh0, ll1\}$, we have $\alpha = \frac{1}{4}\left(\sqrt{p}\sqrt{p+8}-p\right)$, $o_h = \frac{1}{4}\left(\sqrt{p}\sqrt{p+8}-p\right)$, and $o_l = \frac{1}{4}p\left(p-\sqrt{p}\sqrt{p+8}+4\right)$. This yields $Uh = \frac{2p(h+p)}{p(r+2)+\sqrt{p(p+8)}r}$, which we need to be greater than 1 to to exclude a profitable deviation to hl1. Then $\frac{2p(h+p-1)}{p+\sqrt{p(p+8)}} > r$. We also need there be to no profitable deviation away from $hh0$: $\frac{2p^2}{hp+h\sqrt{p(p+8)}} < r$.

For $\{hh1, ll1\}$, we have $\alpha = 1/2$, $o_h = p/2$, and $o_l = p/2$. This yields $Uh = \frac{(h+1)p}{p+2r}$. we need to ensure a profitable deviation to $hh0$: $r > \frac{p}{2h}$. Then $\{hh1, hh0, ll1\}$ is the only equilibrium.

Now we need to ensure that, with $hh0$ censored, $\{hh1, hl1, ll1\}$ is the only equilbrium. The set of admissable strategy profiles is now $\{hh0, hl1, ll1\}$, $\{hh1, ll1\}$, and we've already eliminated the latter.

For$\{hh1, hl1, ll1\}$, we have $\alpha = 1/2$, $o_h = p$, and $o_l = p$. This yields $Uh = \frac{(h+2)p}{2(p+r)}$, which we need to be less than 1 to exclude a profitable deviation away from hl1: $\frac{hp}{2} < r$.

TS given censoring is $2(\frac{(h+2)p}{2(p+r)})$

TS without censoring is $\frac{2p(h+p)}{p(r+2)+\sqrt{p(p+8)}r} + \frac{4p}{p(r+4)+\left(\sqrt{p(p+8)}+4\right)r}$

$\Delta TS = 2(\frac{(h+2)p}{2(p+r)}) - (\frac{2p(h+p)}{p(r+2)+\sqrt{p(p+8)}r} + \frac{4p}{p(r+4)+\left(\sqrt{p(p+8)}+4\right)r})$

Therefore, we need $\Delta TS > 0$.

The set of parameter vectors satisfying these constraints is not empty: $h = 1.3$, $p = .2$,$r = .13$ satisfies all of them. Thus, the platform can improve total surplus by censoring search. However, these constraints are quite restrictive. In particular, $hp/2 < r$ requires a linearly increasing discount rate as h increases–for any r, there is an H such that, for $h > H$, censored search cannot induce acceptance of $hl1$. □

Proposition 18:

Proof. By Lemma D, assume s=0 without loss of generality. By lemmas A and C, there exists $\psi^* = (\psi^*_{hh},\psi^*_{hl}, \psi^*_{lh}, \psi^*_{ll})$ such that $\psi^*_{lh} = \psi^*_{hl}$ and ψ^* maximizes total surplus. Now let p be such that $0 \geq max\{max_{\psi\in\Psi}[EU_h[\psi, p]], max_{\psi\in\Psi}[EU_l[\psi, p]]\}$. Then every draw is accepted by every agent, regardless of agent strategy profiles. Now set$\psi_{\mathbf{c}} \equiv \psi^*$. Then agents accept every draw above the ψ^* cutoffs, and only receive those draws. Choose $f_i = EU_i(\psi^*, p)$. Then $EU_i(\psi^*, p, f) = 0$.

□

References

Armstrong, M. (2006). Competition in Two-Sided Markets. *RAND Journal of Economics 37(3)*, 668-691.

Adachi, H. (2003). A Search Model of Two-Sided Matching under Nontransferable Utility. *Journal of Economic Theory 113(2)*, 182-198.

Abdulkadiroğlu, A., P. Pathak, and A. Roth. (2005). The New York City High School Match. *American Economic Review Papers and Proceedings (May 95)*, 364-367.

Abdulkadiroğlu, A., P. Pathak, A. Roth, and T. Sönmez. (2005). The Boston Public School Match. *American Economic Review Papers and Proceedings (May 95)*, 368-371.

Azevedo, E. and J. Leshno. (2012). A Supply and Demand Framework for Two-Sided Matching Markets. *Working Paper*.

Becker, G. S. (1973). A Theory of Marriage: Part 1. *Journal of Political Economy 81(4)*, 813-846.

Bester, H. (1994). Price Commitment in Search Markets. *Journal of Economic Behavior and Organization 25(1)*, 109-120.

Bloch, F. and H. Ryder. (2000). Two-Sided Search, Marriages, and Matchmakers. *International Economic Review 41(1)*, 93-115.

Browning, M., P. Chiappori, and Y. Weiss. (2014). *Economics of the Family*. Cambridge Surveys of Economic Literature: Cambridge University Press.

Bulow, J. and J. Levin. (2006). Matching and Price Competition. *American Economic Review, 96(3)*, 652-668.

Burdett, K. (1978). A Theory of Employee Job Search and Quit Rates. *American Economic Review, 68(1)*, 212-220.

Burdett, K. and M. Coles. (1997). Marriage and Class. *Quarterly Journal of Economics 112*, 141-168.

Burdett, K. and R. Wright. (1998). Two-sided Search with Nontransferable Utility. *Review oF Economic Dynamics 1(1)*, 220-245.

Cabral, L. (2011). Dynamic Price Competition with Network Effects. *Review of Economic Studies 78*, 83-111.

Cornelius, T. (2003). A Search Model of Marriage and Divorce. *Review of Economic Dynamics 6(1)*, 135-155.

Damiano, E. and L. Hao. (2007). Price Discrimination and Efficient Matching. *Economic Theory 30*, 243-263.

Flanders, S. (2014). Mirror Matches: Analyzing Matching Markets with N-Dimensional Preferences. *Working Paper*.

Gale, D. and L.S. Shapley. (1962). College Admissions and the Stability of Marriage. *The American Mathematical Monthly 69(1)*, 9-15.

Hagiu, A. (2006). Pricing and Commitment by Two-Sided Platforms. *RAND Journal of Economics 37(3)*, 720-737.

Halaburda, H., M. Piskorski., and P. Yildirim (2016). Competing by Restricting Choice: The Case of Search Platforms. *Harvard Business School Working Paper 10-098*.

Hoppe, H., B. Moldavanu, and A. Sela. (2011). The Theory of Assortative Matching Based on Costly Signals. *Economic Theory 47*, 75-104.

Hotelling, H. (1929). Stability in Competition. *The Economic Journal 39(153)*, 41-57.

Kojima, F. and P. Pathak. (2009). Incentives and Stability in Large Two-Sided Matching Markets. *American Economic Review 99(3)*, 608-627.

McCall, J. (1970). Economics of Information and Job Search. *The Quarterly Journal of Economics 84(1)*, 113-126.

Mortensen, D. and C. Pissarides. (1994). Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies 61*, 397-415.

Rochet, J. and J. Tirole. (2006). Two-Sided Markets: A Progress Report. *RAND Journal of Economics 37(3)*, 645-667.

Rochet, J. and J. Tirole. (2010). Platform Competition in Two-Sided Markets. *Journal of the European Economic Association 1(4)*, 990-2019.

Rogerson, R., R. Shimer, and R. Wright. (2005). Search-Theoretic Models of the Labor Market: A Survey. *Journal of Economic Literature Vol XLIII*, 959-988.

Roth, A., T. Sönmez, and M.U. Ünver. (2005). Pairwise Kidney Exchange. *Journal of Economic Theory 125*, 151-188.

Roth, A., T. Sönmez, and M.U. Ünver. (2007). Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences. *The American Economic Review 97(3)*, 828-851.

Rysman, M. (2009). The Economics of Two-Sided Markets. *Journal of Economic Perspectives 23(3)*, 125-143.

Salop, S. C. (1979). Monopolistic Competition with Outside Goods. *The Bell Journal of Economics 10(1)*, 141-156.

Shapley, L.S. and M. Shubik. (1971). The Assignment Game I: The Core. *International Journal of Game Theory 1(1)*, 111-130.

Shimer, R. and L. Smith. (2001). Matching, Search, and Heterogeneity. *Advances in Macroeconomics* 1(1).

Weyl, G. (2010). A Price Theory of Multi-Sided Platforms. *American Economic Review 100(4)*, 1642-1672.