

Targeted Search in Matching Markets*

Anton Cheremukhin, Paulina Restrepo-Echavarria,
Antonella Tutino

February 14, 2017

Abstract

We endogenize the degree of randomness in the matching process by proposing a model where agents have to pay a search cost to locate potential matches more accurately. The model features a tension between an agent's desire to find a more productive match and to maximize the odds of finding a match. This tension drives a wedge between the shape of sorting patterns and the shape of the underlying match payoff function. We show the empirical relevance of the latter prediction by applying the model to the U.S. marriage market.

JEL: E24, J64, C78, D83.

Keywords: Matching, sorting, assignment, search.

* Anton Cheremukhin: Federal Reserve Bank of Dallas, 2200 N Pearl St, Dallas TX 75201, chertosha@gmail.com, 214-922-6785. Paulina Restrepo-Echavarria: Federal Reserve Bank of St. Louis, One Federal Reserve Bank Plaza, Broadway and Locust St., St. Louis MO 63166, paulinares@me.com. Antonella Tutino: Federal Reserve Bank of Dallas, 2200 N Pearl St, Dallas TX 75201, tutino.antonella@gmail.com, 214-922-6804. Authors thank Fernando Alvarez, Gady Barlevy, Hector Chade, Jan Eeckhout, Mike Golosov, Christian Hellwig, Oleg Itskhoki, Boyan Jovanovic, Ricardo Lagos, Jim Peck, Robert Ulbricht, Pierre-Olivier Weill, and David Wiczer for discussions and suggestions which helped improve the paper. We also thank seminar participants at multiple conferences. All errors are our own. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas, St. Louis or the Federal Reserve System. First draft: January 6, 2012.

1 Introduction

The search for a partner—whether in business or in personal life—includes both productive and strategic considerations: People seek productive partnerships that maximize their pay-offs, but also search strategically by considering the likelihood a potential partner will also be interested in them. As a result, we observe that matching is not perfectly assortative nor purely random. For instance, the empirical marriage literature documents an abundance of matches between inferior and superior partners.¹ Understanding why this is the case is important for understanding the roles played by preferences and frictions in the matching process, and for evaluating losses from mismatch.

One reason for apparent mismatch may be that the econometrician does not observe all the match-relevant characteristics.² However, even studies with exhaustive information about the qualities of potential matches and their choice patterns leave a large part of the variation in the data unexplained.³

Another important reason is search frictions. Matches between inferior and superior types may form if one of the partners cannot afford to "wait for the best" and decides to "settle for the rest".⁴ Who marries whom is also influenced by who meets whom.⁵ Although, circumstances are a major driver of the meeting process, whom to meet is also a choice. To our knowledge, the existing literature is silent on how to model this choice. Instead, it is common to postulate an exogenous random meeting process.

¹See survey by Chiaporri and Salanie (2015) and references therein.

²See Choo and Siow (2006) and Galichon and Salanie (2012).

³See Hitsch et. al. (2010).

⁴See Eeckhout (1999), Shimer and Smith (2000), Adachi (2003).

⁵See e.g. Belot and Francesconi (2013).

In this paper we propose a parsimonious way to model the choice of whom to meet, that captures all of the salient features of the meeting process. One reason that people cannot start their search by simply contacting the best partners is that they are uncertain who those are. However, people can choose to put some effort into finding out about the options and locating the best partners. Empirical studies of discrete choice suggest that people respond to incentives by putting more effort and locating best options with a larger probability when stakes are higher.⁶ The tendency to choose options with a higher payoff more often has also been extensively documented for marriage market choices.⁷ To accommodate this empirical pattern, we build on the growing literature on stochastic discrete choice in our modeling strategy. We assume that agents choose who to contact in a probabilistic way. The strategies chosen by agents are their distributions of interest. In this way, agents are free to target one or several potential partners by contacting them with higher probability. However, increasing the likelihood of locating the best partners involves cognitive effort. We capture this intuition by associating a proportional cost with a measure of distance between an "uninformed" strategy and the agent's strategy of choice. The measure of distance has the interpretation of cognitive effort associated with locating best partners. We borrow the specification of cost from the burgeoning literature exploring the connection between information frictions and discrete choice.⁸ With this approach, our model naturally captures the intuition that agents choose to contact more often the potential partners that yield a higher expected payoff.

Another reason for contacting a potential partner is confidence that interest is likely to be mutual, while partners that are unlikely to reciprocate

⁶See the large literature starting with Hey and Orme (1994).

⁷See Fisman et. al. (2008) and Hitsch et. al. (2010).

⁸See Cheremukhin et. al. (2015) and Matejka and McKay (2015).

are worth excluding. Thus, people act strategically not only when deciding whether to form a match or wait for a better option, but also when choosing who to contact. Our model captures this strategic motive, as agents' equilibrium strategies respond positively both to the utility they derive from a match, and the likelihood that interest will be mutual. Thus, our model endogenizes the randomness of the meeting process by formalizing the fact that when looking for a match agents have two incentives: first, to maximize their actual payoffs (which we refer to as the productive motive), and second, to maximize the odds of forming a match (which we refer to as the strategic motive).

We first present a bare-bones one-shot model that has the minimal ingredients to showcase the mechanism. We define a matching equilibrium of the model as the pure-strategy Nash equilibrium between agents' probabilistic strategies. This solution concept is less restrictive than the stability requirement coming from cooperative games. Nonetheless, we prove existence of equilibrium, derive the necessary conditions characterizing it, and sufficient conditions for uniqueness. Furthermore, we show that the equilibria that emerge from a positive and finite cost are inefficient relative to the constrained Pareto allocation, although the outcomes are constrained efficient in both limiting cases.

To better understand the interaction between search frictions, we extend the model to a repeated setting along the lines of Adachi (2000) and Eeckhout (1999). The extended model captures two channels through which search frictions operate: agents have both the choice of who to meet and the choice of who to reject upon meeting. The first type of friction stems from the inability of agents to perfectly distinguish among potential matches, and the second type of friction is a consequence of impatience (time cost of

search). First, we find that both frictions are necessary for existence of mismatch: if agents could easily locate their best partner or wait indefinitely to meet them, the outcome is that of the frictionless model. Second, we find that the two types of frictions reinforce each other: if agents can distinguish their best matches better, the equilibrium likelihood of meeting is higher, which increases the continuation value of waiting, just like an increase in patience. Finally, the extended model highlights that the difference between the outcomes for transferable and non-transferable utility environments is a consequence of patience, while in our baseline model with zero patience the two environments are identical.

Our theoretical results are invaluable for building intuition. The meeting rates in the model are determined by the interaction between the productive and strategic motives. The relative strength of the two motives depends on the search cost. When the cost is low, locating the best partner requires little effort, and the likelihood of reciprocation increases. In this case, the reciprocity of interest is the paramount determinant of who meets whom, while the importance of joint productivity may be undermined by the lack of mutual interest. The strategic motive dominates. This case supports multiple equilibrium arrangements of mutual interest, but only a subset of these arrangements approaches the set of stable allocations in the frictionless limit.

On the other hand, when it is costly to locate the best partner, the likelihood of contacting an inferior partner increases. The likelihood of reciprocation also becomes similar for most potential partners. The strategic motive is mitigated and the productive motive dominates. Payoffs become the driving force behind who meets whom. In the unique equilibrium, every agent's strategy is to chase the partner that would yield the highest payoff,

while impact of mutuality of interest is minimal.

This property can be used for empirical identification of agents' preferences. In particular, the literature describes "horizontal" (attraction to agents similar to themselves) and "vertical" (attraction monotone in a commonly agreed-upon ranking) preferences, but finds it hard if not impossible to distinguish between these cases empirically.⁹ Indeed, both cases lead to identical assortative stable matching in the frictionless case. In contrast, the equilibrium matching rates predicted by our model differ markedly for these two cases. When preferences are "horizontal", the strategic and productive motive pull agents in the same direction, as look-alikes both get the highest payoff from each other and their interest is also more likely to be mutual. In this case, a stochastic version of assortative matching is preserved in equilibrium. However, in the case of vertical preferences, there is common agreement on the ranking of agents and everybody tries to chase the top type despite the lower likelihood of mutual interest. The productive and the strategic motive pull in opposite directions. This case gives rise to a novel and surprising equilibrium pattern that reminisces neither positive nor negative assortative matching. We call it the mixing equilibrium.¹⁰ This matching pattern in the data could be used to distinguish empirically between horizontal and vertical preferences.

We discuss how preferences can be identified from the data on matching rates and from the data on contact rates if those are observed. We show that the tension between the productive and the strategic motive drives a wedge between the shape of sorting patterns and the shape of the underlying match payoff function recovered by the model. The correlation between the shape of the equilibrium matching rate and the shape of the underlying

⁹See Hitsch et. al. (2010) and related studies.

¹⁰Our taxonomy of equilibria in this case follows that of Burdett and Coles (1999).

payoff function depends on whether these motives co-move or clash. If preferences are horizontal the resulting correlation between the matching rate and the underlying payoff function is high, while if preferences are vertical the resulting correlation between can be low. This result has important implications for empirical inference as it implies that one can derive the wrong conclusion about the form of preferences, and hence, the amount of losses from mismatch in the marriage market by just looking at the matching rate.

To show the empirical relevance of this result, we take our model to the data. We use the model to rationalize observed matching rates in the U.S. marriage market and estimate the underlying match payoff function between males and females. We use a standard dataset for matching in the U.S. marriage market, and construct matching rates based on income, age, and education separately. We find that the model does a very good job rationalizing the observed marriage matching rates based on income, age, and education, and for these three cases we estimate the underlying payoffs implied by our model.

We find that the preferences with respect to income and age are vertical and that preferences with respect to education have both vertical and horizontal areas. For income, this means that marrying someone with higher income is always better. For age, this means that females have a strong preference for older males independent of their own age, while males are virtually indifferent about the age of their spouse. For education, people with low levels of education and people with high levels of education prefer someone with their same level of education, generating a region horizontal preferences. However, people with a medium level of education tend to prefer highly educated people, generating a region with vertical preferences.

The correlation between the three observed matching rates and the re-

covered payoff functions ranges from 0.4 to 0.7. This means that strategic considerations uncovered by endogenizing randomness can drive a significant wedge between the shape of the observed sorting pattern and the shape of the underlying payoff function. Ignoring these considerations may result in misleading implications about the degree of mismatch present in the market and hence about the size of the losses associated with it.

Our paper effectively blends two approaches to introducing randomness used in the empirical literature. The first approach introduces search frictions by assuming that it takes time to find a match, as in Shimer and Smith (2000). The second approach introduces unobserved characteristics as a tractable way of accounting for the deviations of the data from the stark predictions of the frictionless model, as in Choo and Siow (2006) and Galichon and Salanie (2012). We introduce a search friction into the meeting process between agents by endogenizing their choice of whom to meet. We build on the discrete choice rational inattention literature—i.e., Cheremukhin, Popova, and Tutino (2015) and Matejka and McKay (2015)—that endogenizes the multinomial logit discrete choice model by introducing cognitive constraints that capture limits to processing information. Consequently, the equilibrium matching rates in our model have a multinomial logit form similar to that in Galichon and Salanie (2012). Unlike Galichon and Salanie, the equilibrium of our model predicts strong interaction between distributions of agents’ interest that is driven entirely by conscious strategic choices of agents, rather than by some unobserved characteristics with fixed distributions.

The search and matching literature has models like Menzio (2007) and Lester (2010) that nest directed search and random matching to generate outcomes with an intermediate degree of randomness.¹¹ Our paper contributes

¹¹Also, see Yang’s (2013) model of “targeted” search that assumes random search within perfectly distinguishable market segments.

to the literature by providing a theory of how people meet that produces equilibrium outcomes featuring endogenous randomness in between uniform random matching and the frictionless assignment, without nesting these two frameworks. Nevertheless, both complete randomness and frictionless assignment are special cases of our model.

Finally, the paper contributes to the literature on directed search and coordination frictions, as in Eeckhout and Kircher (2010) and Shimer (2005).¹² [Talk about queues and congestion here. Also worth mentioning Menzio and Shi (2009) on block recursive equilibria.] The directed search paradigm generally predicts socially efficient and assortative equilibrium outcomes. In contrast, our targeted search model does not appear to possess a market mechanism that can implement the constrained efficient allocation, nor does it guarantee assortativeness.

The paper proceeds as follows: Section 2 describes the model and derives the theoretical predictions. In Section 3 we discuss how the model can help identify preferences, and then take the model to the U.S. marriage market data. Section 4 states some final remarks.

2 The Model

In this section, we present a model that endogenizes the degree of randomness in matching. We build on the frictionless matching environment of Becker (1973), males and females are heterogeneous in their type and all are searching for a match. Both males and females know the distribution and their preferences over types on the other side of the market, but there is noise—agents cannot locate potential partners with certainty—and they can pay a search cost to help locate them more accurately.

¹²See Chade, Eeckhout and Smith (2016) for a neat summary of this literature.

Agents face a search cost and choose search effort by choosing an optimal probability distribution over types. This distribution reflects the likelihood of locating a particular agent. A more targeted search, or a probability distribution that is more concentrated on a particular group of agents implies a higher cost as it means that the agent is exerting more search effort to locate a potential match of a certain type. As agents optimally choose how targeted their search strategy is by choosing a probability distribution, the degree of randomness in matching is endogenous. When they play a more targeted strategy, the degree of randomness goes down.

The probability distribution needs to satisfy two properties: 1) By the nature of the choice between a finite number of options, the distribution must be discrete; and 2) for strategic motives to play a role, agents should be able to vary each element of the distribution and consider small deviations of each element in response to changes in the properties of the options. Hence, this probability distribution cannot be confined to a specific family of distributions.

The choice of functions in economics that satisfy these requirements is very limited. We use the Kullback-Leibler divergence (relative entropy) as the measure of search effort. This specification accommodates both full choice of a distribution and a discrete choice problem. In addition, it turns out that, in our specific case of a choice among discrete options, this specification enhances tractability and leads to closed-form solutions. Specifically, the solution has the form of a multinomial logit that is well understood and already widely used in empirical studies of discrete choice environments.

After choosing their optimal probability distribution over types, both males and females make a draw from their distribution. If the draw is reciprocated, a match is formed if it is mutually beneficial and the output from

the match is split between the two parties. As the search involves balancing the costs and benefits of prospective matches, some participants will not find partners immediately.

2.1 The environment

There are F females indexed by $x \in \{1, \dots, F\}$ and M males indexed by $y \in \{1, \dots, M\}$. Both males and females are heterogeneous in types and are actively searching for a match. A match between female x and male y generates a payoff Φ_{xy} .¹³ If a male and a female match, the payoff is split between them. We normalize the outside option of both to zero. We denote the share of the payoff appropriated by the female ε_{xy} and the share of the payoff appropriated by the male η_{xy} such that $\Phi_{xy} = \varepsilon_{xy} + \eta_{xy}$.¹⁴ The payoff and the split generated by any potential (x, y) match are known ex-ante to female x and male y .

Each female chooses a discrete probability distribution, $p_x(y)$, which reflects the probability with which female x will target male y (seek him out). Each female x rationally chooses her strategy (i.e. the probability of targeting a male y) while facing a trade-off between a higher payoff and a higher cost of searching. Likewise we denote the strategy of a male $q_y(x)$. It represents the probability of a male y targeting a female x . Each agent can vary and choose each element of their distribution. Placing a higher mass on any particular potential match, implies that the agent choosing the distribution has exerted more search effort, will target a potential partner more accurately and hence, will have a higher probability of matching with them.

A female's total cost of searching is given by $c_x(\kappa_x(p_x(y)))$. This cost is

¹³Note that we do not place any restrictions on the payoff function.

¹⁴Given that our aim is to understand the implications of having endogenous randomness, we focus on the simplest model and assume fixed payoff shares.

a function of the search effort, κ_x , and hence of the probability distribution, $p_x(y)$, chosen by female x . Likewise, we denote a male's cost of searching by $c_y(\kappa_y(q_y(x)))$, where the cost is a function of the search effort, κ_y , and hence of the probability distribution, $q_y(x)$, chosen by male y .

Figure 1 illustrates the strategies of males and females. Consider a female $x = 1$. The solid arrows show how she assigns a probability of targeting each male $p_1(y)$. Similarly, dashed arrows show the probability that a male $y = 1$ assigns to targeting a female $q_1(x)$. Once these are selected, each male and female will make one draw from their respective distribution to determine which individual they will actually target. A match is formed between male y and female x if and only if: 1) according to the female's draw of y from $p_x(y)$, female x targets male y ; 2) according to the male's draw of x from $q_y(x)$, male y also targets female x ; and 3) their payoffs are non-negative.




Figure 1. Strategies of males and females

Since negative payoffs lead to de facto zero payoffs due to the absence of a match, we can assume that all payoffs are non-negative:

$$\Phi_{xy} \geq 0, \quad \varepsilon_{xy} \geq 0, \quad \eta_{xy} \geq 0.$$

Each female x chooses a strategy $p_x(y)$ to maximize her expected net payoff:

$$Y_x = \max_{p_x(y)} \sum_{y=1}^M \varepsilon_{xy} q_y(x) p_x(y) - c_x(\kappa_x(p_x(y))).$$

The female gets her expected return from a match with male y minus the cost of searching. The probability of a match between female x and male y is given by the product of the distributions $q_y(x) p_x(y)$. Note that in equilibrium the matching rate that female x faces from male y equals male y 's strategy $q_y(x)$. As matching rates are equilibrium objects, they are assumed to be common knowledge to participating parties.

The cost function is given by $c_x(\kappa_x) = \theta_x \kappa_x$, where θ_x is the search cost. Here, we are using the linear case for simplicity, but all of our proofs will hold for more general cost functions. As mentioned earlier, κ_x reflects search effort and needs to accommodate the full choice of a discrete distribution and attach a cost to it. One function that satisfies these requirements is the following:¹⁵

$$\kappa_x = \sum_{y=1}^M p_x(y) \ln \frac{p_x(y)}{1/M}, \quad (1)$$

where $p_x(y)$ must satisfy $\sum_{y=1}^M p_x(y) = 1$ and $p_x(y) \geq 0$ for all y .

Note that, κ_x , is increasing in the distance between a uniform distribution $\{1/M\}$ over males and the strategy, $p_x(y)$. If an agent does not want to exert any search effort, she can choose a uniform distribution $p_x(y) = \frac{1}{M}$ over types, the cost of search is zero, and her search is random. As she chooses

¹⁵In the model of information frictions used in the rational inattention literature, κ_x , would represent the relative entropy between a uniform prior $\{1/M\}$ over males and the posterior strategy, $p_x(y)$. This definition is a special case of Shannon's channel capacity where information structure is the only choice variable (See Thomas and Cover (1991), Chapter 2).

a more targeted strategy the distance between the uniform distribution ($\frac{1}{M}$) and her search strategy $p_x(y)$ is greater, increasing κ_x and the overall cost of searching, and her search will be less random. By increasing search effort agents bring down uncertainty about the location of a prospective match, which allows them to target their better matches more accurately.

Similarly, male y chooses his strategy $q_y(x)$ to maximize his expected payoff:

$$Y_y = \max_{q_y(x)} \sum_{x=1}^F \eta_{xy} p_x(y) q_y(x) - c_y(\kappa_y(q_y(x))),$$

where

$$\kappa_y = \sum_{x=1}^F q_y(x) \ln \frac{q_y(x)}{1/F}, \quad (2)$$

and $q_y(x)$ must satisfy $\sum_{x=1}^F q_y(x) = 1$ and $q_y(x) \geq 0$ for all x .

2.2 Matching Equilibrium

Definition 1. A matching equilibrium is a set of strategies of females, $\{p_x(y)\}_{x=1}^F$, and males, $\{q_y(x)\}_{y=1}^M$, that simultaneously solve problems of males and females.

The equilibrium of the matching model can be interpreted as a pure-strategy Nash equilibrium of a strategic form game. In what follows we shall apply the results for concave n-person games from Rosen (1965). The game consists of the set of players, the set of actions and the player's payoffs. The set of players is given by $\mathcal{I} = \{x \in \{1, \dots, F\}, y \in \{1, \dots, M\}\}$.

The set of actions $s \in S$ is given by the cartesian product of the sets of strategies of females $p_x(y) \in S_x$ and males $q_y(x) \in S_y$, where $S_x =$

$$\left\{ p_x(y) \in R^M, p_x(y) \geq 0, \sum_{y=1}^M p_x(y) \leq 1 \right\} \text{ and } S_y = \left\{ q_y(x) \in R^F, q_y(x) \geq 0, \sum_{x=1}^F q_y(x) \leq 1 \right\}.$$

The payoffs $u_i(s) = \{Y_x(s), Y_y(s)\}$ are defined as follows:

$$Y_x(p_x(y), q_y(x)) = \sum_{y=1}^M \varepsilon_{xy} q_y(x) p_x(y) - c_x \left(\sum_{y=1}^M p_x(y) \ln \frac{p_x(y)}{1/M} \right),$$

$$Y_y(q_y(x), p_x(y)) = \sum_{x=1}^F \eta_{xy} p_x(y) q_y(x) - c_y \left(\sum_{x=1}^F q_y(x) \ln \frac{q_y(x)}{1/F} \right).$$

Theorem 1. *A matching equilibrium exists.*

Proof. Note that the strategy set of each player is a unit simplex, and therefore a non-empty, convex and compact set. For a pure-strategy Nash equilibrium to exist, each payoff function $u_i(s)$ needs to be continuous in the strategies s , and $u(s_i, s_{-i})$ needs to be quasiconcave in s_i . Indeed, under the assumption that cost functions are continuous, non-decreasing and (weakly) convex, the payoff functions are continuous and concave in the own strategies of players. For the case of a linear cost function these restrictions are trivially satisfied. □

To show uniqueness, we need to introduce some additional notation. Note that for each player $i \in \mathcal{I}$ the strategy set can be represented as $S_i = \{s_i \in R^{m_i}, h_i(s_i) \geq 0\}$ where h_i is a concave function. Indeed in our case, the functions $h_x(p_x(y)) = \left[-p_x(1), \dots, -p_x(M), 1 - \sum_{y=1}^M p_x(y) \right]$ and $h_y(q_y(x)) = \left[-q_y(1), \dots, -q_y(F), 1 - \sum_{x=1}^F q_y(x) \right]$ are concave. Following Rosen, define the gradient $\nabla u(s) = [\nabla_1 u_1(s), \dots, \nabla_m u_m(s)]^T$ and Hessian: $U(s) = \nabla_i \nabla_j u_i(s)$. Then, if the constraints h_i are concave and the symmetrized Hessian $U(s) + U^T(s)$ is negative definite for all $s \in S$ then the

payoff functions are diagonally strictly concave for $s \in S$. We can then use the result that if h_i are concave functions, if there exist interior points $\tilde{s}_i \in S_i$ such that $h_i(\tilde{s}_i) > 0$, and if the payoff functions are diagonally strictly concave for all $s \in S$ then the game has a unique pure strategy Nash equilibrium.

Theorem 2. *The matching equilibrium is unique if*

- a) *cost functions are non-decreasing and convex;*
- b) $\left. \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \right|_{p_x^*(y)} = \theta_x > \varepsilon_{xy} p_x^*(y);$
- c) $\left. \frac{\partial c_y(\kappa_y)}{\partial \kappa_y} \right|_{q_y^*(x)} = \theta_y > \eta_{xy} q_y^*(x).$

Proof. If the cost functions $c(\kappa)$ are (weakly) increasing and (weakly) convex in κ , then the payoffs of all males and females are continuous and also concave in their strategies. Assuming that the cost functions are twice continuously differentiable functions, the Hessian of this game is the matrix of all second derivatives. The diagonal elements are all non-positive, consistent with concavity of the payoffs:

$$\frac{\partial^2 Y_x}{\partial p_x \partial p_x} = -\frac{\partial c_x}{\partial \kappa_x}(\kappa_x) \frac{1}{p_x(y)} - \frac{\partial^2 c_x}{\partial \kappa_x \partial \kappa_x}(\kappa_x) \left(1 + \ln \frac{p_x(y)}{1/M}\right)^2 \leq 0,$$

$$\frac{\partial^2 Y_y}{\partial q_y \partial q_y} = -\frac{\partial c_y}{\partial \kappa_y}(\kappa_y) \frac{1}{q_y(x)} - \frac{\partial^2 c_y}{\partial \kappa_y \partial \kappa_y}(\kappa_y) \left(1 + \ln \frac{q_y(x)}{1/F}\right)^2 \leq 0.$$

The off-diagonal elements are all non-negative:

$$\begin{aligned} \frac{\partial^2 Y_x}{\partial p_x \partial q_y} &= \varepsilon_{xy} \geq 0, \\ \frac{\partial^2 Y_y}{\partial q_y \partial p_x} &= \eta_{xy} \geq 0, \end{aligned}$$

The remaining cross-derivatives are all zero. Note also that the Hessian is itself symmetric so there is no need to symmetrize it. To guarantee that the

Hessian is negative definite, we require the following “Diagonal dominance” conditions:

$$\left| \frac{\partial^2 Y_x}{\partial p_x \partial p_x} \right| > \left| \frac{\partial^2 Y_x}{\partial p_x \partial q_y} \right|,$$

$$\left| \frac{\partial^2 Y_y}{\partial q_y \partial q_y} \right| > \left| \frac{\partial^2 Y_y}{\partial q_y \partial p_x} \right|.$$

Diagonal dominance conditions postulate that diagonal elements of the Hessian are larger in absolute value than any off-diagonal elements, which in turn guarantees that the Hessian of the game is negative definite. It is clear that when the cost functions are linear, these conditions simplify to $\theta_x \frac{1}{p_x(y)} > \varepsilon_{xy}$ and $\theta_y \frac{1}{q_y(x)} > \eta_{xy}$. While Rosen’s version requires that these conditions hold globally for all $s \in S$ which would imply $\theta_x > \varepsilon_{xy}$ and $\theta_y > \eta_{xy}$, these conditions could be relaxed to require diagonal dominance to be satisfied only along the equilibrium path. For this we note that since the constraints are given by unit simplexes (for which index equals one, and every KKT point is complementary and non-degenerate) we can invoke the generalized Poincare-Hopf index theorem of Simsek, Ozdaglar, and Acemoglu (2007) which in this case implies that the equilibrium is unique if the Hessian is negative-definite at the equilibrium point. Thus, the diagonal dominance conditions need to hold only along the equilibrium path, i.e. if conditions (b) and (c) are satisfied. $\square$

Note that the assumptions we make to prove uniqueness are by no means restrictive. The assumption that cost functions are non-decreasing and convex is a natural one. The additional “diagonal dominance” conditions in our case can be interpreted as implying that the search cost of reducing noise should be sufficiently high for the equilibrium to be unique. If these conditions do not hold, then there can be multiple equilibria. This is a well-known

outcome of the assignment model, which is a special case of our model under zero search costs. In a frictionless environment, the multiplicity of equilibria is eliminated by requiring that the matching be “stable”, i.e., that there is no profitable pairwise deviation. In our framework, ensuring “stability” would require that all males know the location of all females to be able to check all pairwise deviations. Since locating agents is very costly in our model, the equilibrium outcome generically does not satisfy “stability.”

The result of Theorem 2 is intuitive. Recall that there are two motives for female x to target male y . The productive and the strategic motive. The payoff of a female depends on the product of the portion she appropriates from the output of the match and the probability of reciprocation. While her portion of the split does not depend on equilibrium strategies, the strategic motive does. When the search cost, θ_x , is very low, females (and males) are able to place a high probability of targeting one counterparty and exclude all others. As a result, when θ_x is extremely low, the strategic motive dominates. It does not matter what portion of the payoff female x will get from a match with male y if the male chooses not to consider female x . When the strategic motive dominates, multiplicity of equilibria is a natural outcome. In the extreme, any pairing of agents is an equilibrium since no one has an incentive to deviate from any mutual reciprocation.

As θ_x and θ_y increase, probability distributions become less precise, as it is increasingly costly to target a particular counterparty. That is, the search costs dampen the strategic motive and the productive motive plays a bigger role. At some threshold level of θ each agent will be exactly indifferent between following the strategic motive and seeking a better match. This level of costs is precisely characterized by the “diagonal dominance” conditions of Theorem 2. Agents require the strategic motive, characterized by the off-

diagonal element of the Hessian of the game, to be lower than the productive motive, captured by the diagonal element. Above the threshold the unique equilibrium has the property that each agent places a higher probability on the counterparty that promises a higher payoff, i.e., the productive motive dominates.

Under the assumptions on the cost functions made earlier we can also obtain a characterization result. The derivatives of the constrained payoff functions with respect to own strategies are:

$$\begin{aligned}\frac{\partial Y_x}{\partial p_x} &= \varepsilon_{xy} q_y(x) - \frac{\partial c_x}{\partial \kappa_x}(\kappa_x) \left(1 + \ln \frac{p_x(y)}{1/M}\right) - \lambda_x, \\ \frac{\partial Y_y}{\partial q_y} &= \eta_{xy} p_x(y) - \frac{\partial c_y}{\partial \kappa_y}(\kappa_y) \left(1 + \ln \frac{q_y(x)}{1/F}\right) - \lambda_y.\end{aligned}$$

When cost functions are non-decreasing and convex, it is easy to verify that first-order conditions are necessary and sufficient conditions for equilibrium. Rearranging the first-order conditions for males and females, we obtain

$$\begin{aligned}p_x^*(y) &= \exp\left(\frac{\varepsilon_{xy} q_y^*(x)}{\theta_x}\right) / \sum_{y'=1}^M \exp\left(\frac{\varepsilon_{xy'} q_{y'}^*(x)}{\theta_x}\right), \\ q_y^*(x) &= \exp\left(\frac{\eta_{xy} p_x^*(y)}{\theta_y}\right) / \sum_{x'=1}^F \exp\left(\frac{\eta_{x'y} p_{x'}^*(y)}{\theta_y}\right).\end{aligned}\tag{3}$$

These necessary and sufficient conditions for equilibrium cast the optimal strategy of female x and male y in the form of a best response to optimal strategies of males and females, respectively.

Equilibrium conditions (3) have an intuitive interpretation. They predict that the higher the female's private gain from matching with a male, the higher the probability of targeting that male. Similarly, the higher the probability that a male targets a particular female, the higher the probability that

that female targets that male. Overall, females target males that promise higher expected private gains by placing higher probabilities on those males. Males are naturally sorted in each female's strategy by the probability of the female targeting each male. The strategies of males have the same properties due to the symmetry of the problem.

In equilibrium, a male's strategy is the best response to the strategies of females, and a female's strategy is the best response to the strategies of males. Theorem 2 predicts that an increase in θ reduces the complementarities between search strategies of females and males. Once θ is sufficiently high, the intersection of best responses leads to a unique equilibrium. Note that, by the nature of the index theorem used in the proof of uniqueness, it is enough to check diagonal dominance conditions locally in the neighborhood of the equilibrium. There is no requirement for them to hold globally. This suggests a simple way of finding equilibria of our model in most interesting cases. We first need to find one solution to the first-order conditions (3) and then check that diagonal dominance conditions are satisfied.

Now, consider the properties of equilibria for two limiting cases. First, as the search costs go to zero, targeting strategies become more and more precise. In the limit, in every equilibrium each female places a unit probability on a particular male, and that male responds with a unit probability of considering that female. Each equilibrium of this kind implements a matching of the classical assignment problem (not all of them are stable).

Second, in the opposite case when search costs go to infinity, optimal strategies of males and females approach a uniform distribution. This unique equilibrium implements the standard uniform random matching assumption extensively used in the literature. Thus, the assignment model and the random matching model are special cases of our targeted search model, when θ

is either very low or very high.

2.3 Efficiency

To evaluate the efficiency of the equilibrium, we compare the solution of the decentralized problem to a social planner's solution. We assume that the social planner maximizes the total payoff, which is a utilitarian welfare function. To achieve a social optimum, the planner can choose the strategies of males and females. If there were no search costs, the planner would always choose to match each male with the female that produces the highest output. The socially optimal strategies of males would be infinitely precise.

To study the constrained efficient allocation we impose on the social planner the same costs of search that we place on males and females. Thus, the social planner maximizes the following welfare function:

$$W = \max_{p_x(y), q_y(x)} \sum_{x=1}^F \sum_{y=1}^M \Phi_{xy} p_x(y) q_y(x) - \sum_{x=1}^F c_x(\kappa_x(p_x(y))) - \sum_{y=1}^M c_y(\kappa_y(q_y(x)))$$

subject to (1-2) and to the constraints that $p_x(y)$ and $q_y(x)$ are well-defined probability distributions.

Under the assumption of increasing and convex cost functions, the social welfare function is concave in the strategies of males and females. Hence, first-order conditions are necessary and sufficient conditions for a maximum. Rearranging and substituting out Lagrange multipliers, we arrive at the following characterization of the social planner's allocation:

$$\begin{aligned} p_x^o(y) &= \exp\left(\frac{\Phi_{xy} q_y^o(x)}{\theta_x}\right) / \sum_{f'=1}^F \exp\left(\frac{\Phi_{xy'} q_{y'}^o(x)}{\theta_x}\right), \\ q_y^o(x) &= \exp\left(\frac{\Phi_{xy} p_x^o(y)}{\theta_y}\right) / \sum_{m'=1}^M \exp\left(\frac{\Phi_{x'y} p_{x'}^o(y)}{\theta_y}\right). \end{aligned} \quad (4)$$

The structure of the social planner's solution is very similar to the structure of the decentralized equilibrium given by (3). From a female's perspective, the only difference between the two strategies is that the probability of targeting a male depends on the social gain from a match rather than on her private gain. Notice that the same difference holds from the perspective of a male. Thus, it is socially optimal for both females and males to consider the total payoff, while in the decentralized equilibrium they consider only their private payoffs.

This result is reminiscent of goods with positive externalities where the producer undersupplies the good if she is not fully compensated by the marginal social benefits that an additional unit of the good would provide to society. In our model, additional search effort exerted by an individual male or female has a positive externality on the whole matching market.

For instance, when a male chooses to increase his search effort, he can better locate his preferable matches. As a consequence, the females he targets will benefit (through an increase in the personal matching rate) and the females he does not target will also be better off as his more targeted strategy will help them exclude him from their search (through a decrease in the personal matching rate). Nevertheless, in this environment agents can not appropriate all the social benefits (the output of a match) they provide to society when increasing their search effort. They only get a fraction of the payoff. This failure of the market to fully compensate both females and males with their social marginal products leads to under-supply of search effort by both sides in the decentralized equilibrium.

Because the social gain is always the sum of private gains, there is no feasible way of splitting the payoff such that it implements the social optimum. When θ is finite and positive, a socially optimal equilibrium has to satisfy

the following conditions simultaneously:

$$\varepsilon_{xy} = \Phi_{xy}, \quad \eta_{xy} = \Phi_{xy}.$$

In the presence of heterogeneity, these optimality conditions can hold in equilibrium only if the total payoff is zero, as private gains have to add up to the total payoff, $\varepsilon_{xy} + \eta_{xy} = \Phi_{xy}$. Therefore, we have just proven the following theorem:

Theorem 3. *The matching equilibrium is socially inefficient for any split of the payoff if all of the following hold:*

- 1) *cost functions are increasing and convex;*
- 2) $\Phi_{xy} > 0$ *for some* (x, y) ;
- 3) $\Phi_{xy} \neq \Phi_{xy'}$ *for some* x, y *and* y' ;
- 4) $\Phi_{xy} \neq \Phi_{x'y}$ *for some* y, x *and* x' ;
- 5) $0 < \left. \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \right|_{p_x^*} = \theta_x < \infty$;
- 6) $0 < \left. \frac{\partial c_y(\kappa_y)}{\partial \kappa_y} \right|_{q_y^*} = \theta_y < \infty$.

Proof. See Appendix A. □

The first two conditions are self-explanatory; the case when all potential matches yield zero payoffs is a trivial case of no gains from matching. Conditions 5 and 6 state that marginal costs of reducing noise have to be finite and positive in the neighborhood of the equilibrium. When θ is zero, the best equilibrium of the assignment model is socially optimal. When θ is very high, the random matching outcome is the best possible outcome. For all intermediate values of marginal costs, the decentralized equilibrium is socially inefficient.

Conditions 3 and 4 together require heterogeneity to be two-sided. If heterogeneity is one-sided, i.e. condition 3 or condition 4 is violated, then

the allocation of intentions towards the homogeneous side of the market will be uniform. In this case, search becomes one-sided and equilibrium allocations are efficient contingent on the actively searching side appropriating 100 percent of the payoff.¹⁶

One notable property of the equilibrium is that, by considering only fractions of the total payoff when choosing their strategies, males and females place lower probabilities on pursuing their best matches. This implies that in equilibrium, probability distributions of males and females are more dispersed and the number of matches is lower than is socially optimal.

Another way of thinking about the inefficient quantity of matches is to consider the reduction in strategic complementarities. To illustrate these complementarities consider the case of a female who chooses her strategy under the assumption that all males implement socially optimal strategies. Because a female considers only her private gains from matching with a male, the female's optimal response would be to target less accurately the best males than is socially optimal. In a second step, taking as given these strategies of females, males will be discouraged to reduce noise further, not only by the fact that they appropriate fractions of the total gains from a match, but also by the fact that females do not choose strategies that are as targeted as is socially optimal. These complementary disincentives will lower the probabilities of males pursuing their best match. Iterating in this way on strategies of males and females, at each step the probability of targeting the best matches decreases. As a result, agents will target their *better* matches instead of the best possible matches.

The inefficiency that arises in the two-sided model can in principle be corrected by a central planner. This can be done by promising both males

¹⁶See Appendix B for a version of the model with one-sided heterogeneity.

and females that they will receive the entire payoff of each match and then by collecting lump-sum taxes from both sides of the market to cover the costs of the program. Nevertheless, to do so, the planner himself would need to acquire extensive knowledge about the distribution of the payoffs, which is costly. We leave this point for future research.

2.4 Implications for Sorting

To better understand the effect of the productive and strategic motives, it is useful to consider simple examples of payoffs to understand the relative importance of these motives for equilibrium matching rates. Let us consider a matching market where there are just two males and two females, with types labeled, high (H) and low (L). Let us also consider two specific cases of the form of the payoff function.

Case one: The high type female is better off with a high type male, and the low type female is better off with a low type male. The same property is true for males. We shall generally refer to a payoff function where for each type the best option on the other side is different - as the case of *relative* advantage. Case two: Both females prefer the high type male, and both males prefer the high type female. We shall generally refer to a payoff for which everyone's best option is the same type as the case of *absolute* advantage.

In the case of relative advantage, the strategic and the productive motives are aligned. The productive motive points all agents in different directions, and the strategic motive ensures that the same agent that implies a higher payoff is also the one that is more likely to reciprocate (because agents have no incentive to compete for the same match). However, in the case of absolute advantage, the productive motive points all agents in the same direction,

while the strategic motive tends to coordinate agents on paying attention to those whom their competitors are less likely to consider, to maximize the odds of finding a match. Thus, there is a conflict between the two motives as they pull intentions in different directions.

If the payoff function exhibits relative advantage, and the search costs are low enough, our model can have two different equilibrium patterns. The first pattern is when the high type is more likely to target the high type and the low type to target the low type (HH, LL). This is the case of positive assortative matching (PAM). The second pattern is when the high type is more likely to target a low type, because the low type is more likely to reciprocate (HL, LH). This is the case of negative assortative matching (NAM). However, if search costs are high, only the PAM equilibrium survives because the productive motive dominates.

If the payoff function exhibits *absolute* advantage, and search costs are low enough, in addition to the PAM and NAM equilibria that we described above, there is a third equilibrium pattern, which we call a *mixing* equilibrium. In the mixing equilibrium, both females target the high type male, and both males target the high type female. Moreover, for high enough search costs, the unique equilibrium has the mixing pattern, while the PAM and NAM equilibria disappear. These patterns are illustrated in Figure 1.

This last result is in stark contrast with the literature on optimal assignment, which predicts a PAM equilibrium as the only stable outcome. The prediction of the assignment model is driven by the strategic motive. If search costs are low, the high types look at only each other, so it makes no sense for the low types to target the high types as, despite a higher potential payoff, the chance their interest will be reciprocated is zero. However, when search costs are high enough, the strategic motive is dampened to the extent that

the productive motive starts to play a dominant role. The productive motive instructs people to place a higher probability on the type that promises a higher payoff. Hence, the unique mixing equilibrium.




Figure 1: Three types of equilibria and their sorting patterns

Note: We show by an arrow the direction in which each agent places the highest probability.

This basic intuition has important implications for empirical inference. If the productive and strategic motives are perfectly aligned, as they are in the relative advantage case, then the shape of the equilibrium matching pattern looks very similar to the shape of the payoff function. Indeed, agents will always place the highest probability on the types that give the highest payoff and we shall see a larger number of matches between those pairs of types. The presence of a conflict between these motives, as in the absolute advantage case, drives a wedge between the shape of the payoff and the shape of the matching rates. On the one hand, you should still see more matches between pairs of types that are more productive. On the other hand, there is a large number of competing agents that would be able to compensate for the lower payoff by a higher probability of reciprocation. The main consequence of this result is that when the payoff function is such that the two motives are in conflict, the pattern of who marries whom may differ substantially from the pattern of who would be better off with whom.

To quantify this difference, we run a set of Monte Carlo simulations and compute the correlation between the equilibrium matching rate and the underlying payoff function. For the Monte Carlo simulations, we assume three males and three females and draw each element of the 3-by-3 payoff matrix from a uniform distribution. We make 25,000 draws. We then find all equilibria and corresponding matching rates for each draw of the payoff function. We discard all the payoffs that produce multiple equilibria. For the draws that have a unique equilibrium, we compute the correlation between the matrix of equilibrium matching rates and the payoff matrix. In Figure 2 we show the probability density functions of correlations for three classes of payoff functions: those exhibiting absolute advantage, relative advantage, or no clear advantage pattern.




Figure 2: Correlation of Matching rate and Payoff

We find that, indeed, in the case of absolute advantage the correlation is significantly lower than that in the case of relative advantage. The intermediate shapes of payoff generate intermediate values of the correlation. Thus, when our model is the true data-generating process, the conflict between the productive and strategic motive could drive a substantial wedge between the shape of the underlying payoff function and the shape of the matching rate. Consequently, the empirical researcher could easily arrive at wrong conclusions about the shape of the underlying payoff by simply looking at the shape of the matching rates. As we shall discuss at the end of the empirical section, this is indeed what workhorse models of the marriage market do.

To put this result in context, we note that both random matching models a la Shimer and Smith (2000), and directed search models a la Eeckhout and Kircher (2010) can produce a substantial wedge between the shape of the payoff function and the shape of the matching rates. In the case of random matching the distribution is uniform, while in the case of directed search matching is fully assortative. Matching patterns in both of these cases are accommodated by our model under extreme (very high or very low) values of search costs. Our model also spans the continuum of matching patterns in between these two extremes.

To show that the wedge between the matching pattern and the payoff function is indeed present in the data and empirically relevant, in the empirical section, we explore three prominent examples of matching patterns in the marriage market. We show that, when viewed through the lens of our model, they exhibit absolute advantage or a combination of absolute and relative advantage. Also, we observe a substantial wedge between the shape of the underlying payoff function and the matching rate.

2.5 Invertibility

Our model builds on the interaction of strategic motives of agents and is hence more complicated computationally than leading examples in the literature. This fact has both bonuses and drawbacks. We find that in our model the mapping between the payoff and the matching rate is not invertible. By that we mean that there exist matching rate patterns that cannot in principle be generated by our model. Also, we cannot exclude the possibility that some matching rate pattern could be generated by more than one payoff function (although we could not find an example of this in practice).

This implies that our model is testable. That is, we could observe data on matching rates that would be at odds with the predictions of our model. To illustrate this point we perform a Monte-Carlo exercise by drawing the elements of the 2-by-2 payoff matrix from a uniform distribution, computing the equilibrium and the corresponding matching rates. We normalize the total expected number of matches to one and plot all the possible vectors of equilibrium matching rates on a (3-dimensional) simplex. Figure 3 illustrates our findings.




Figure 3: Subset of model-generated matching rates in a 3D simplex

We find that large white spaces remain in the simplex, implying that many shapes of matching rates cannot be obtained as an equilibrium outcome of our model. The intuition for this result is simple. If, for example, both types of males search actively, an equilibrium cannot allocate the majority of prospective matches to just one of the males and generate almost no matches for the other male. This is why we argue that our theory of targeted search is testable. It implies certain restrictions on equilibrium outcomes which may

or may not be rejected by the data. This result is in contrast with workhorse models in the marriage literature, such as Choo and Siow (2006) and Galichon and Salanie (2012), which can rationalize any observed matching rates.

Given the non-invertibility of the mapping between the payoff and the matching rates, how do we test the model and estimate the payoff function from matching rate data? Our empirical methodology proceeds in three steps. First, we make a few identifying assumptions. In particular, we assume that search costs are identical across agents on both sides of the marriage market, $\theta_x = \theta_y = \theta$. This assumption will allow us to identify the ratio of the payoff to search cost, Φ_{xy}/θ , for each pair of types. In addition, we assume that each payoff is split equally between males and females, i.e. $\epsilon_{xy} = \eta_{xy} = \Phi_{xy}/2$.

Second, for any shape of the payoff function, Φ_{xy} , we find all equilibria (if there are more than one) and compute all corresponding equilibrium matching rates implied by the model. Third, we search for a shape of the payoff that maximizes the likelihood function of the data given the predicted matching rates. Whenever a proposed payoff function produces multiple equilibria, we select the one that fits the observed matching rate best, i.e. has the highest likelihood. Maximization of the likelihood function efficiently minimizes the properly weighted sum of distances between the data and the model's prediction and should lead to consistent estimates. Maximum likelihood estimation of discrete games with multiple equilibria have been reasonably well studied in the literature, e.g. Aguirregabiria and Mira (2007). Here we do not employ any computational tricks since the 3-by-3 case can be computed by brute force in reasonable time. The results of such estimation can be treated as an upper bound on the explanatory power of the model. In the empirical section, we apply this method to three prominent examples of sorting in the

marriage market and find that the model fits the data very well.

3 Empirical Application

To take the model to the data, we use a standard dataset for matching rates in the U.S. marriage market. The data on unmarried males and females and newly married couples comes from IPUMS for the year 2001.¹⁷ For computational transparency we attribute both males and females to three equally sized bins, which we refer to as low (L), medium (M), and high (H) types. We consider three dimensions along which males and females evaluate each other in the marriage market: income, age, and education. In each case we choose the cutoffs between bins in such a way as to split the whole U.S. population of each gender to equally sized bins.

In the case of age, we restrict our attention only to adults between the ages 21 and 33. To make them as close as possible to equal size, the bins correspond to ages 21-23, 24-27, 28-33. We discard all younger and older people from the analysis because there is a disproportionate amount of unmarried people in these other age categories who only rarely marry. One reason for this may be that a large fraction of them are not searching for a spouse and are thus not participating in the marriage market. To avoid misspecification due to our inability to observe search effort, we exclude them from our analysis. In the case of education, the natural breakdown into three bins is to have people who never attended college, those who are currently in college, and those who have graduated from college. Income is a continuous characteristic, so the three bins correspond to people with low, medium, and high incomes.

For each of the three cases, we estimate the shape of the payoff function

¹⁷We thank Gayle and Shephard (2015) for kindly sharing the IPUMS data with us.

using the maximum likelihood methodology described earlier. We assume that all currently unmarried males and females are searching, and the number of matches is proxied by the number of couples that were married in the past 12 months, as indicated by answers to the questionnaire. The dataset contains roughly 93,599 unmarried males, 82,673 unmarried females, and 23,572 newly married couples above the age of 21.

The matching rate for the case of income is presented in the left panel of Figure 4. The estimate of the underlying payoff is shown in the right panel of the same Figure. A notable property of the payoff is that it exhibits strong absolute advantage. That is, marrying a spouse with a higher income is always better. We find that the matching rate and the payoff have a correlation of 0.72.




Figure 4: Sorting by income

The matching rate for the sorting by age is presented in the left panel of Figure 5. Looking at the shape of the matching rate, we would expect to see the pattern of relative advantage here, with slightly older males looking for slightly older females. However, the shape of the payoff that best explains

this sorting pattern is very close to absolute advantage. Females have a strong preference for older males independent of their own age. Meanwhile, males are virtually indifferent to the age of their spouse. The highest payoff is produced by males at age 30 marrying females at age 23. The correlation between the matching rate and the payoff is a staggeringly low 0.42.




Figure 5: Sorting by age

The matching rate for sorting by education is presented in the left panel of Figure 6. In this case the payoff exhibits a combination of absolute and relative advantage. Low educated people and high educated people prefer someone with their same level of education, generating a region of relative advantage. However, people with a medium level of education tend to prefer highly educated people, generating a region with absolute advantage. The matching rate and the payoff function have a correlation of 0.52.




Figure 6: Sorting by education

A widely used workhorse model in the marriage literature is the model of Choo and Siow (2006). They estimate a static transferable utility model that generates a nonparametric marriage matching function. This model postulates that, in equilibrium, each pair of cohorts of males and females reaches an implicit agreement on the matching rate among themselves; matching (or staying single) is a voluntary decision. In their model, the payoff is recovered as a simple algebraic function of the matching rates and the number of people searching. The first notable property of this mapping is that it is one-to-one, i.e., for any payoff there is a unique matching rate and for any matching rate one can invert the relationship to compute the payoff.

The second notable property is that the matching rate depends only on the characteristics of the agents directly involved in the match, but not on the characteristics of other agents present in the marriage market. This is because the strategic motive is absent from their model, so the shape of the matching rate mimics closely the shape of the payoff function. An important consequence of these two properties is that any set of matching rates observed in the data can be rationalized by some form of payoff function. Thus, the model of Choo and Siow does not place any constraints on the

data and cannot be tested. This also implies that the distance between the assumptions and implications is minimal: the correlation between the matching rates across pairs of types and the implied values of the payoff are close to one.

We illustrate this feature in Figure 7 where we use the 3-by-3 Monte Carlo simulation from Section 4.1. We plot the correlation between the true underlying payoff and the equilibrium matching rate obtained from our model on the horizontal axis and the correlation between the same matching rate and the corresponding payoff function recovered by the model of Choo and Siow on the vertical axis. We find that in many cases, the shape of the true payoff and of the matching rate descends to 0.4, while the model of Choo and Siow would imply that they have a similar shape with a correlation above 0.75. We color the payoffs with the three patterns of advantage in three different colors. We find that while the correlation depends significantly on the pattern of advantage in our model, in Choo and Siow's model it does not.

The Figure also compares our empirical findings with the Monte Carlo simulation. We find that the three prominent empirical examples that we have considered indeed belong to the range of correlation values commonly generated by payoffs that exhibit absolute advantage.




Figure 7: Monte Carlo results and Data

This result emphasizes the importance of considering the effect of strategic motives on the sorting patterns in empirical research. If a researcher looks at the data through the lens of a model with exogenous randomness, that model by construction ignores any strategic considerations that may affect agents' search strategies. As we have shown, strategic considerations can drive a significant wedge between the shape of the productive complementarities and the shape of the observed sorting pattern. Ignoring endogenous randomness may thus lead to vastly misleading conclusions regarding the amount of mismatch present in a market and the size of the losses associated with it.

4 Final Remarks

In this paper we endogenize the degree of randomness in the matching process by proposing a model where agents must pay a cost for searching to better locate potential matches. If they increase search effort, they increase the probability of targeting a better match. The model features a productive motive that drives agents to target the person that renders a higher payoff and a strategic motive that drives agents to target the person with whom their interest is more likely to be reciprocated. We believe that ignoring these considerations may result in misleading implications about the degree of mismatch present in the market and hence about the size of the losses associated with it.

With endogenous randomness as the driving force of matching patterns, our model is well suited to study a host of real-life matching markets where people typically have limited time and ability to reduce noise. Roth and Sotomayor (1990) and Sönmez and Ünver (2010) provide examples of such markets. Moreover, for many markets, equilibrium outcomes are neither pure random matching nor optimal assignment, as documented in the empirical literature. Our model can be a useful tool for analyzing these markets.

Furthermore, our model describes markets where the degree of centralization is fairly low. This structure encompasses a number of markets ranging from labor markets to education and health care. In many two-sided market models a platform acts both as a coordination device and as a mechanism to transfer utility. Our model can be used to study the optimal degree of centralization and the social efficiency of pricing schemes in these markets. We view the study of the optimal design of centralization in two-sided search environments as an exciting area of future research and a practical application of our theory with far-reaching consequences.

References

- [1] **Aguirregabiria, V and P. Mira, (2007).** "Sequential Estimation of Dynamic Discrete Games." *Econometrica*, 75(1), pp. 1-53, 01.
- [2] **Becker, G. S., (1973).** "A Theory of Marriage: Part I." *Journal of Political Economy*, 81(4), pp. 813-846.
- [3] **Burdett, K. and M. G. Coles, (1999).** "Long-Term Partnership Formation: Marriage and Employment." *Economic Journal, Royal Economic Society*, 109(456), pp. F307-F334.
- [4] **Chade, H., J. Eeckhout and L. Smith (2016).** "Sorting Through Search and Matching Models in Economics." *Journal of Economic Literature*, Forthcoming.
- [5] **Cheremukhin, A., A. Popova and A. Tutino (2015).** "A Theory of Discrete Choice with Information Costs." *Journal of Economic Behavior and Organization*, vol. 113, pp. 34-50.
- [6] **Choo, E. and A. Siow, (2006).** "Who Marries Whom and Why." *Journal of Political Economy*, 114(1), pp. 175-201.
- [7] **Diamond, P. A., (1982).** "Wage Determination and Efficiency in Search Equilibrium." *Review of Economic Studies*, 49(2), pp. 217-227.
- [8] **Eeckhout, J. and P. Kircher, (2010).** "Sorting and Decentralized Price Competition." *Econometrica*, 78(2), pp. 539-574.

- [9] **Galichon, A. and B. Salanié, (2012).** "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models." Working Paper hal-01053710
- [10] **Gayle, G.-L. and A. Shephard, (2015).** "Optimal Taxation, Marriage, Home Production, and Family Labor Supply." October 2015, University of Pennsylvania Working Paper.
- [11] **Lester, B. (2011).** "Information and Prices with Capacity Constraints," *American Economic Review*, 101(4), pp. 1591-1600.
- [12] **Matejka F. and A. McKay, (2015).** "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review*, 105(1), pp. 272-298.
- [13] **Menzio, G. (2007).** "A Theory of Partially Directed Search," *Journal of Political Economy*, 115(5), pp. 748-769.
- [14] **Roth, A. E. and M. Sotomayor, (1990).** "Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis." *Econometric Society Monograph Series*, Cambridge University Press, 1990.
- [15] **Shimer, R. and L. Smith, (2000).** "Assortative Matching and Search." *Econometrica*, 68(2), pp. 343-370.
- [16] **Shimer, R. (2005).** "The Assignment of Workers to Jobs in an Economy with Coordination Frictions." *Journal of Political Economy*, 113, pp. 996-1025.
- [17] **Simsek, A. Ozdaglar, A. and D. Acemoglu. (2007).** "Generalized Poincare-Hopf Theorem for Compact Nonsmooth Regions," *Mathematics of Operations Research*, 32(1), pp. 193-214.

- [18] **Sönmez, T. and M. U. Ünver, (2010).** "Matching, Allocation, and Exchange of Discrete Resources," in J. Benhabib, A. Bisin, and M. Jackson (eds.), *Handbook of Social Economics*, Elsevier.
- [19] **Thomas, C. and J. Cover, (1991).** "Elements of Information Theory". John Wiley & Sons, Inc.
- [20] **Vives, X. (1990).** "Nash Equilibrium with Strategic Complementarities." *Journal of Mathematical Economics*, 19, pp. 305-321.

NOT FOR PUBLICATION

Appendix A: Proof of Theorem 3

The proof proceeds in 3 steps.

Step 1. Under the assumption of increasing convex cost functions, both individual payoff functions and the social welfare function are concave in the strategies of males and females. Hence, first-order conditions are necessary and sufficient conditions for a maximum.

Step 2. We denote by CEFOC the first-order conditions of the decentralized equilibrium and by POFOC the first-order conditions of the social planner. In formulae:

$$\text{POFOC}_{q_y(x)}: \quad \Phi_{xy} \tilde{p}_x(y) - \frac{\partial c_y(\tilde{\kappa}_y)}{\partial \tilde{\kappa}_y} \Big|_{\tilde{q}_y(x)} \frac{1}{\ln 2} \left(\ln \frac{\tilde{q}_y(x)}{1/F} + 1 \right) - \tilde{\lambda}_y = 0$$

$$\text{POFOC}_{p_x(y)}: \quad \Phi_{xy} \tilde{q}_y(x) - \frac{\partial c_x(\tilde{\kappa}_x)}{\partial \tilde{\kappa}_x} \Big|_{\tilde{p}_x(y)} \frac{1}{\ln 2} \left(\ln \frac{\tilde{p}_x(y)}{1/M} + 1 \right) - \tilde{\lambda}_x = 0$$

$$\text{CEFOC}_{q_y(x)}: \quad \eta_{xy} p_x(y) - \frac{\partial c_y(\kappa_y)}{\partial \kappa_y} \Big|_{q_y(x)} \frac{1}{\ln 2} \left(\ln \frac{q_y(x)}{1/F} + 1 \right) - \lambda_y = 0$$

$$\text{CEFOC}_{p_x(y)}: \quad \varepsilon_{xy} q_y(x) - \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \Big|_{p_x(y)} \frac{1}{\ln 2} \left(\ln \frac{p_x(y)}{1/M} + 1 \right) - \lambda_x = 0$$

For the equilibrium to be socially efficient we need to have the following:

$$\tilde{p}_x(y) = p_x(y) \quad \text{for all } x, y$$

$$\tilde{q}_y(x) = q_y(x) \quad \text{for all } x, y$$

Step 3. By contradiction, imagine that the two conditions above hold.

Then, by construction,

$$\left. \frac{\partial c_y(\tilde{\kappa}_y)}{\partial \tilde{\kappa}_y} \right|_{\tilde{q}_y(x)} = \left. \frac{\partial c_y(\kappa_y)}{\partial \kappa_y} \right|_{q_y(x)} = a_y$$

and

$$\left. \frac{\partial c_x(\tilde{\kappa}_x)}{\partial \tilde{\kappa}_x} \right|_{\tilde{p}_x(y)} = \left. \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \right|_{p_x(y)} = a_x.$$

Denote them a_y and a_x respectively.

It then follows that:

$$\begin{aligned} \Phi_{xy}\tilde{p}_x(y) - \tilde{\lambda}_y &= \left. \frac{\partial c_y(\tilde{\kappa}_y)}{\partial \tilde{\kappa}_y} \right|_{\tilde{q}_y(x)} \frac{1}{\ln 2} \left(\ln \frac{\tilde{q}_y(x)}{1/M} + 1 \right) \\ &= \left. \frac{\partial c_y(\kappa_y)}{\partial \kappa_y} \right|_{q_y(x)} \frac{1}{\ln 2} \left(\ln \frac{q_y(x)}{1/M} + 1 \right) \\ &= \eta_{xy}p_x(y) - \lambda_y \end{aligned}$$

i.e. $\Phi_{xy}\tilde{p}_x(y) - \tilde{\lambda}_y = \eta_{xy}p_x(y) - \lambda_y$ for all x and y . We can use the first-order conditions of the females to derive the formulas for λ and $\tilde{\lambda}$:

$$\begin{aligned} (i) \quad M \exp \left(1 + \frac{\tilde{\lambda}_y}{a_y / \ln 2} \right) &= \sum_{x=1}^M \exp \left(\frac{\Phi_{xy}p_x(y)}{a_y / \ln 2} \right) \\ (ii) \quad M \exp \left(1 + \frac{\lambda_y}{a_y / \ln 2} \right) &= \sum_{x=1}^M \exp \left(\frac{\varepsilon_{xy}(x)p_x(y)}{a_y / \ln 2} \right) \\ (iii) \quad (\Phi_{xy} - \varepsilon_{xy})p_x(y) &= \tilde{\lambda}_y - \lambda_y \quad \text{for all } x \end{aligned}$$

Jointly (i) (ii) and (iii) imply:

$$\frac{\sum_{x'=1}^M \exp \left(\frac{\Phi_{x'y}p_{x'}(y)}{a_y / \ln 2} \right)}{\sum_{m'=1}^M \exp \left(\frac{\varepsilon_{x'y}p_{m'}(y)}{a_y / \ln 2} \right)} = \frac{\exp \left(\frac{\Phi_{xy}p_x(y)}{a_y / \ln 2} \right)}{\exp \left(\frac{\varepsilon_{xy}p_x(y)}{a_y / \ln 2} \right)} \quad \text{for all } x$$

Hence,

$$\frac{\exp(\Phi_{xy}p_x(y))}{\exp(\varepsilon_{xy}p_x(y))} = \frac{\exp(\Phi_{x'y}p_{x'}(y))}{\exp(\varepsilon_{x'y}p_{x'}(y))} \quad \text{for all } x \text{ and } x'.$$

Therefore, either:

$$\text{a) } \Phi_{xy} = \varepsilon_{xy} \text{ for all } x \text{ or}$$

$$\text{b) } \Phi_{x'y} = \Phi_{x''y} \text{ and } \varepsilon_{x'y} = \varepsilon_{x''y} \text{ for all } x' \text{ and } x'';$$

Similarly from males' first-order conditions it follows that either :

$$\text{c) } \Phi_{xy} = \eta_{xy} \text{ for all } y \text{ or}$$

$$\text{d) } \Phi_{xy'} = \Phi_{xy''} \text{ and } \eta_{x'y} = \eta_{x''y} \text{ for all } y' \text{ and } y''$$

Cases b) and d) have been ruled out by the assumptions of the theorem. Cases a) and c) jointly imply that $\varepsilon_{xy} = \eta_{xy} = \Phi_{xy} = \varepsilon_{xy} + \eta_{xy}$ which leads to a contradiction $\varepsilon_{xy} = \eta_{xy} = \Phi_{xy} = 0$.

Appendix B: One-sided model

Here we consider a one-sided model where females are searching for males who are heterogeneous in type and females face a search cost. We assume that there is no heterogeneity on the female side of the market. As such the probability that a male reciprocates the intentions of a female is given by q_y . The strategy of a female, denoted $p_x(y)$, represents the probability of female x locating male y . It is also the female's probability distribution. We assume that each female can rationally choose her strategy facing a trade-off between a higher payoff and a higher cost of searching.

Like before a female's cost of searching is given by $c_x(\kappa_x)$. Once the optimal distribution $p_x(y)$ is chosen, each female draws from it to determine which male to contact.¹⁸

Female x chooses a strategy $p_x(y)$ to maximize her expected income flow:

$$Y_x = \max_{p_x(y)} \sum_{y=1}^M \varepsilon_{xy} p_x(y) q_y - c_x(\kappa_x)$$

We normalize the outside option of females to zero. A female receives her expected share of the payoff in a match with male y conditional on matching with that male. She also incurs a cost that depends on search effort:

$$\kappa_x = \sum_{y=1}^M p_x(y) \ln \frac{p_x(y)}{1/M} \quad (5)$$

where the female's strategy must satisfy $\sum_{y=1}^M p_x(y) = 1$ and $p_x(y) \geq 0$ for all y .

Definition 2. *A matching equilibrium of the one-sided matching model is a set of strategies of females, $\{p_x(y)\}_{x=1}^N$, which solve their optimization problems.*

Theorem 4. *If the cost functions are non-decreasing and convex, the one-sided matching model has a unique equilibrium.*

Proof. The payoffs of all females are continuous in their strategies. They are also concave in these strategies when cost functions are (weakly) increasing and convex. Hence, each problem has a unique solution. $\square$

¹⁸As in the model of Section 2, we assume that each female pursues only one male.

When in addition the cost functions are differentiable, it is easy to verify that first-order conditions are necessary and sufficient conditions for equilibrium.¹⁹ Rearranging the first order conditions for the female, we obtain:

$$p_x^*(y) = \exp \left(\frac{\varepsilon_{xy} q_y}{\frac{1}{\ln 2} \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \Big|_{p_x^*}} \right) / \sum_{y'=1}^M \exp \left(\frac{\varepsilon_{xy'} q_{y'}}{\frac{1}{\ln 2} \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \Big|_{p_x^*}} \right). \quad (6)$$

This is an implicit relationship as p_x^* appears on both sides of the expression. If cost functions are linear functions of κ_x , then the derivatives on the right hand side are independent of p_x^* , and the relationship becomes explicit.

The equilibrium condition (6) has an intuitive interpretation. It predicts that the higher is the female's expected gain from matching with a male, the higher is the probability placed on locating that male. Thus, males are naturally sorted in each female's strategy by probabilities of contacting those males.

Efficiency To study the constrained efficient allocation we impose upon the social planner the same constraints that we place on females. Thus, the social planner maximizes the following welfare function:

$$W = \sum_{x=1}^F \sum_{y=1}^M \Phi_{xy} p_x(y) q_y - \sum_{x=1}^F c_x(\kappa_x)$$

¹⁹Taking derivatives of the Lagrangian function corresponding to the problem of female x , we obtain for all y :

$$\varepsilon_{xy} q_y - \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \Big|_{p_x^*} \frac{1}{\ln 2} \left(\ln \frac{p_x^*(y)}{1/M} + 1 \right) = \lambda_x$$

We can invert this first-order condition to characterize the optimal strategy:

$$p_x^*(y) = \frac{1}{M} \exp \left(\frac{\varepsilon_{xy} q_y - \lambda_x}{\frac{1}{\ln 2} \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \Big|_{p_x^*}} - 1 \right).$$

subject to (5) and to the constraint that the $p_x(y)$'s are well-defined probability distributions. Under the assumption of increasing convex cost functions, the social welfare function is concave in the strategies of females. Hence, first-order conditions are sufficient conditions for a maximum. Rearranging and substituting out Lagrange multipliers, we arrive at the following characterization of the social planner's allocation:

$$p_x^o(y) = \exp \left(\frac{\Phi_{xy} q_y}{\frac{1}{\ln 2} \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \Big|_{p_x^o}} \right) / \sum_{y'=1}^M \exp \left(\frac{\Phi_{xy'} q_{y'}}{\frac{1}{\ln 2} \frac{\partial c_x(\kappa_x)}{\partial \kappa_x} \Big|_{p_x^o}} \right). \quad (7)$$

The first observation to make is that the structure of the social planner's solution is very similar to the structure of the decentralized equilibrium. Second, from the female's perspective, the only difference between the centralized and decentralized equilibrium strategies is that the probability of locating a male depends on the social gain from a match rather than on the private gain. Thus, it is socially optimal to consider the whole expected payoff when determining the socially optimal strategies, while in the decentralized equilibrium females only consider their private gains.

To decentralize the socially optimal outcome the planner needs to give all of the payoff to the females, $\varepsilon_{xy} = \Phi_{xy}$, effectively assigning them a share of 1. Note that, if the planner could choose the probability that a male reciprocates a female, q_y , he would also set it to 1.

The only special cases, when the outcome is always efficient are the limiting cases discussed earlier. When search costs are absent, the equilibrium of the model is socially optimal. When costs are very high, the random matching outcome is the best possible outcome. For all intermediate values of costs, the decentralized equilibrium is socially efficient contingent on the female having all the bargaining power.

Appendix C: Repeated two-sided model

Here we extend the two-sided matching model to a repeated setting. Following Adachi (2003) we assume that at the moment that male y and female x meet, each of them has an additional decision to make. Each agent may choose to form a match and receive the corresponding share of the surplus, or refuse to form a match and wait for a better potential partner in future periods if their continuation value is higher than the utility from matching with the proposed partner. The continuation value is assumed to be simply the expected utility of matching in the future discounted at the rate ρ , which is the patience parameter. In the Adachi model the case $\rho = 1$ represented a frictionless case which implied that agents could wait for their preferred match indefinitely at no time cost to them. Notice, that our one-shot model represents the opposite case of $\rho = 0$.

We denote v_x the continuation value of female x and w_y the continuation value of male y . Each agent chooses her strategy and pays the cost of search before the game starts, and then makes a sequence of draws from the chosen distribution. The time-0 problems of the agents are like before:

$$Y_x = \sum_{y=1}^M EU_x(y) q_y(x) p_x(y) - \theta_x \left(\sum_{y=1}^M p_x(y) \ln \frac{p_x(y)}{1/M} \right) + \lambda_x \left(1 - \sum_{y=1}^M p_x(y) \right),$$

$$Y_y = \sum_{x=1}^F EU_y(x) p_x(y) q_y(x) - \theta_y \left(\sum_{x=1}^F q_y(x) \ln \frac{q_y(x)}{1/F} \right) + \lambda_y \left(1 - \sum_{x=1}^F q_y(x) \right).$$

Where the continuation values are defined as the solutions to the Bellman programs:

$$v_x = \rho \sum_{y=1}^M EU_x(y) q_y(x) p_x(y) + \rho \left(1 - \sum_{y=1}^M q_y(x) p_x(y) \right) v_x,$$

$$w_y = \rho \sum_{x=1}^F EU_y(x) p_x(y) q_y(x) + \rho \left(1 - \sum_{x=1}^F p_x(y) q_y(x) \right) w_y.$$

and the expected returns are either equal to match utilities if both partners agree to a match, or to continuation values if they do not reach an agreement:

$$EU_x(y) = v_x + (\eta_{xy} - v_x) I(\eta_{xy} \geq v_x) I(\varepsilon_{xy} \geq w_y),$$

$$EU_y(x) = w_y + (\varepsilon_{xy} - w_y) I(\eta_{xy} \geq v_x) I(\varepsilon_{xy} \geq w_y).$$

An equilibrium of the model is a set of strategies $\{p_x(y)\}_{x=1}^F, \{q_y(x)\}_{y=1}^M$, reservation values $\{v_x\}_{x=1}^F, \{w_y\}_{y=1}^M$, and expected utilities $\{EU_x(y)\}_{x=1}^F, \{EU_y(x)\}_{y=1}^M$ that jointly solve the problems of the agents and satisfy the remaining system of equations above. Since the maximization problems are well-defined, the first-order conditions are still necessary conditions and must be satisfied in equilibrium. However, because the remaining functions are continuous, but not everywhere differentiable, the model may have multiple equilibria for many different combinations of parameters and it is hard to establish definitive results regarding uniqueness.

So far, this model explicitly postulates non-transferable utility, but it can easily be extended to the case of transferable utility. Specifically, the TU case allows for redistributing the surplus in the cases when joint surplus of the match exceeds the sum of continuation values of the agents. Therefore, the last two equations are replaced in the TU case by:

$$EU_x(y) = v_x + (\eta'_{xy} - v_x) I(\eta_{xy} + \varepsilon_{xy} \geq v_x + w_y),$$

$$EU_y(x) = w_y + (\varepsilon'_{xy} - w_y) I(\eta_{xy} + \varepsilon_{xy} \geq v_x + w_y).$$

where the utilities adjusted for the payments are defined as:

$$\eta'_{xy} = v_x + \frac{\eta_{xy}}{\Phi_{xy}} (\eta_{xy} + \varepsilon_{xy} - v_x - w_y),$$

$$\varepsilon'_{xy} = w_y + \frac{\varepsilon_{xy}}{\Phi_{xy}} (\eta_{xy} + \varepsilon_{xy} - v_x - w_y).$$

Note that in the one-shot model in the main text the TU case and the NTU case are identical as the continuation values are zero.

The repeated game with patience is instructive as it highlights two independent sources of search frictions: the costs of waiting and the costs of distinguishing among agents. According to Smith et al. (1999), search costs are divided into external and internal costs. External costs include the monetary costs of searching and contacting partners as well as the opportunity costs of the time spent on searching. These costs are captured by the parameter ρ in the repeated model. Internal costs include the mental effort associated with the search process, sorting the incoming information, and integrating it with what the agent already knows. Modeling the internal costs is the novel feature of our model. Internal costs are captured by parameter θ which describes the agents' ability to evaluate available information, depending on intelligence, prior knowledge, education and training.

When the agents are unable to distinguish partners until they meet, i.e. when the cost parameter θ approaches infinity, we obtain the Adachi (2003) model. In that case, if the patience parameter ρ approaches 1, the model replicates the frictionless matching outcome as agents are able to wait as long as necessary to meet their best match. Similarly, when agents cannot wait and match everybody that they meet, i.e. the patience parameter ρ is

set to zero, we obtain our baseline one-shot model. In that case, if agents are nonetheless able to perfectly distinguish among potential partners, i.e. the parameter θ approaches zero, the model, with a refinement permitting only stable matchings, also reproduces the frictionless matching outcome. This discussion highlights that both internal and external costs of search are necessary to obtain outcomes where superior agents are matched with inferior agents in equilibrium: the agents need to be reasonably impatient and be unable to perfectly distinguish among potential partners. In the Figures below we illustrate, using the surplus example from Adachi (2003), the regions of the parameter space in which the equilibrium is unique and where the equilibrium involves matches between all pairs of types.




Figure 8: Number and types of equilibria depending on parameters

Appendix C: Notes for self

From Fudenberg and Tirole, chapter 12.3.

Suppose that each player i 's strategy set S_i is a subset of R^{m_i} . Then the union set S is a subset of R^m . Define meet and join of any x and y in S as:

meet: $x \wedge y = [\min(x_1, y_1), \dots, \min(x_m, y_m)]$

join: $x \vee y = [\max(x_1, y_1), \dots, \max(x_m, y_m)]$

Definition: S is a sublattice of R^m if $s_1 \in S$ and $s_2 \in S$ imply $s_1 \wedge s_2 \in S$ and $s_1 \vee s_2 \in S$.

Topkis result: if $S_i = R^{m_i}$ and if u_i is twice continuously differentiable in s_i , then u_i is supermodular in s_i if and only if for any two components s_{ik} and s_{il} of s_i the cross derivative $\partial^2 u_i / \partial s_{ik} \partial s_{il} \geq 0$.

Definition: A supermodular game is such that for each i , S_i is a sublattice of R^{m_i} , u_i has increasing differences in (s_i, s_{-i}) and u_i is supermodular in s_i .

Remark: Supermodularity in s implies both supermodularity in s_i and increasing differences in (s_i, s_{-i}) .

Hence, if u_i is twice continuously differentiable, u_i is supermodular if and only if $\partial^2 u_i / \partial s_k \partial s_l \geq 0$ for any two components of s .

Topkis 1979: If for each i , S_i is compact and u_i is upper semi-continuous in s_i for each s_{-i} , and if the game is supermodular, the set of pure strategy Nash equilibria is nonempty and possesses greatest and least equilibrium points.

This all does not work, since the simplex in R^{m_i} (with the usual vector ordering) $\{p_y \in R^{m_i}, \sum_y p_y = 1, p_y \geq 0\}$ is not a lattice.

If our game was on a lattice, we would get nice comparative statics, but we don't.

From Asu Ozdaglar:

Given a strategic form game $[I, S_i, u_i]$, assume that the strategy sets S_i are

nonempty, convex, and compact sets, $u_i(s)$ is continuous in s , and $u_i(s_i, s_{-i})$ is quasiconcave in s_i . Then the game has a pure strategy Nash equilibrium.

We will next establish conditions that guarantee that a strategic form game has a unique pure strategy Nash equilibrium, following the classical paper [Rosen 65].

Assume that for player $i \in I$ the strategy set is given by $S_i = \{s_i \in R^{m_i}, h_i(s_i) \geq 0\}$ where h_i is a concave function. Therefore, S_i is a convex set.

Therefore, the set of strategy profile S being a cartesian product of convex sets, is also a convex set.

A strategy s^* is a pure-strategy Nash equilibrium if and only if for all $i \in I$, s_i^* is the solution of:

$$\max_{s_i} u_i(s_i, s_{-i}) \text{ subject to } h_i(s_i) \geq 0$$

$$\text{Define gradient: } \nabla u(s) = [\nabla_1 u_1(s), \dots, \nabla_m u_m(s)]^T$$

$$\text{Define Jacobian/Hessian: } U(s) = \nabla_i \nabla_j u_i(s)$$

Definition: Payoff functions are diagonally strictly concave for $s \in S$ if for every $s', s'' \in S$ it holds that: $(s' - s'')^T \nabla u(s') + (s'' - s')^T \nabla u(s'') > 0$

Theorem: For the strategy sets given above assume h_i to be concave functions and that there exist $\tilde{s}_i \in S_i$ such that $h_i(\tilde{s}_i) > 0$. If the payoff functions are diagonally strictly concave for $s \in S$ then the game has a unique pure strategy Nash equilibrium.

Proposition: If the constraints h_i are concave and the symmetric Jacobian $U(s) + U^T(s)$ is negative definite for all $s \in S$ then the payoff functions are diagonally strictly concave for $s \in S$.