Targeted Search in Matching Markets*

Anton Cheremukhin, Paulina Restrepo-Echavarria, Antonella Tutino

February 26, 2016

Abstract

We endogenize the degree of randomness in the matching process by proposing a model where agents have to pay a cost to better distinguish among potential matches. The model features a tension between an agents' 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 productive complementarities. 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, endogenous randomness, 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, Mike Golosov, Christian Hellwig, Oleg Itskhoki, Boyan Jovanovic, Ricardo Lagos, Jim Peck, Robert Ulbricht, and Pierre-Olivier Weill 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

When searching for a partner -whether in business or in life- people look for a productive partnership that can maximize their pay-offs, but they also look strategically for somebody that reciprocates their interests. Sorting through potential matches and pondering these requirements takes time and effort.

Taking this view, the empirical evidence on sorting patterns appears natural, as matches between superior and inferior types should in fact be common in the data if the strategic considerations matter more than the productive ones. Yet, theoretical models of search and matching have difficulties generating these observed patterns. These models typically postulate ad-hoc randomness in the matching process by introducing search frictions or unobserved characteristics. Both approaches treat randomness as exogenous by assuming its distribution.

In this paper we propose a model that endogenizes the degree of randomness in the search and matching process. We build on the idea that when agents search for a match, they know their preferences over types, but distinguishing among them is a costly and time-consuming process. We assume that agents can pay a cost to reduce the noise level associated with the process of distinguishing among potential matches. The trade-off agents face is that they would like to identify better potential matches that are more likely to reciprocate their interests with high probability, but this reduction of the noise level comes at a higher cost (e.g., cognitive effort in sorting through candidates). As a result of this cost-benefit analysis agents optimally choose to identify their best matches only partially and the degree of randomness in the outcome of the matching process is endogenous.

Even though we believe that our theory can potentially apply to different markets, throughout the paper we focus on the marriage market. We build on the frictionless matching environment of Becker (1973) with two-sided heterogeneity, and assume that even though both men and women know the distribution over types of the other side of the market, there is noise—they can not tell what is the type of a potential partner or what are the odds of being reciprocated by them—so they face a cost of distinguishing among potential matches.

In reality as in our model, every individual acquires information by carrying out a screening process to reduce the noise about their potential matches' type. A more detailed screening entails a higher cost but allows the agent to have more certainty about who is the best match that has a higher probability of reciprocating. In the end, men and women optimally choose the probability with which they will target each individual of the opposite sex. We refer to this process of probabilistically targeting their search as choosing an optimal distribution of attention over types.

The more concentrated the distribution of attention around a few types, the higher the probability of being matched with the types that render a higher surplus. However, such a distribution implies a higher cost than a more spread out distribution across types. Once men and women select the optimal distribution of attention, they draw a match from that distribution. If the outcome of the draw is reciprocated and is mutually beneficial, a match is formed. The surplus of the match is split between the two parties through ex-post Nash bargaining. As the search involves balancing costs and benefits of information about prospective matches, agents will optimally choose partially targeted strategies and some participants will remain unmatched.

When choosing the optimal distribution of attention there are two forces at play: i) a productive motive and ii) a strategic motive. The productive motive drives agents to target the person that renders a higher payoff, while the strategic motive drives agents to target the person with whom their interest is more likely to be reciprocated.

These two motives play a crucial role in shaping two theoretical predictions of our model: (1) the uniqueness and inefficiency of the equilibrium; (2) the implications for sorting pattern.

In terms of the uniqueness of the equilibrium, the strategic motive, makes the search strategies of men and women complementary; both will be more interested in someone who is more likely to reciprocate their interests. The strength of these complementarities is key for determining the uniqueness of the equilibrium.[1] In the optimal assignment model, complementarities between search strategies are strong and lead to multiplicity of equilibria,[2] whereas complementarities are absent in random matching models. We show that an increase in the marginal cost of decreasing the noise level makes the search strategies of market participants less complementary and eliminates the multiplicity of equilibria.

Whether the productive or the strategic motive dominates depends on how costly it is to identify the type and odds of reciprocation of a potential partner. When the costs of acquiring information are sufficiently high the productive motive dominates and there is always a positive probability for a low type to be matched with a high type (because the low type knows that the high type cannot correctly identify him as being a low type). In this case the equilibrium that emerges is unique and of the mixing type, i.e. some high

[1]The complementarities in our model are different from search externalities studied by Diamond (1982) where an increase in the number of participants makes it easier for one side of the market and more difficult for the other side to find a match. In our model, the complementarity arises because if one person targets another, the other has an incentive to reciprocate.

[2]Many models of directed search, e.g. Shimer (2005), introduce additional assumptions in order to select one of those equilibria. The requirement of "stability" is another common equilibrium selection strategy.

types match with high types and some high types match with low types.[3] If the costs are below a certain threshold the strategic motive dominates and there are multiple equilibria that can deliver positive assortative matching, negative assortative matching or a mixing equilibrium.[4]

As a result, the equilibria of our model lie *in between* the random matching outcome and the frictionless assignment outcome and encompass both outcomes as limiting cases. In the extreme, if the cost of reducing noise tends to infinity, then agents will not incur on it and the optimal strategy corresponds to the random matching outcome. If the cost of reducing noise is zero, then agents can perfectly identify their best match and the outcome is a frictionless assignment. 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.

The intuition behind the inefficiency is as follows. When an agent decides to pay a higher cost, she is choosing a more targeted strategy. This increases the probability of ending up with a better match and lowers the probability of being paired with a bad match. Because of the complementarities in our model, high types are more likely to assign a higher probability to other high types. Low types would then place a lower probability to this participant in their distribution of attention and target somebody else. Market participants are unable to fully internalize this positive externality. Agents on both sides of the market fail to appropriate all social benefits of their actions, and, as a result, the quantity and the quality of matches are both inefficiently low.

The second theoretical prediction of our model is that the shape of the

[3]The mixing equilibrium is non-assortative when productive complementarities are characterized by absolute advantage. In the case of relative advantage assortativeness is preserved.

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

surplus determines whether the productive and the strategic motives are aligned or lead in different directions. Furthermore, the correlation between the shape of the equilibrium matching rate and the shape of the underlying surplus depends on whether these motives co-move or clash. If the surplus exhibits relative advantage[5] the productive motive points agents in all different directions and the strategic motive assures that the agent that renders the higher payoff is also more likely to reciprocate (because there is less competition for the same agent). Hence, the resulting correlation between the matching rate and the underlying surplus is high. However, if the surplus exhibits absolute advantage,[6] 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 improve the odds of forming a match). Hence, the resulting correlation between the matching rate and the underlying surplus can be low.

The fact that the correlation between the matching rate and the underlying surplus can be low has important implications for empirical inference. It means that one can derive the wrong conclusion about the shape of the underlying productive complementarities by simply looking at the shape of the matching rate.

To show the empirical relevance of this finding we take our model to the data. We use a standard dataset for matching rates in the U.S. marriage market, and we construct the matching rate based on income, age and education separately. For these three cases we estimate the underlying surplus

[5] We say that the surplus exhibits relative advantage when a high type female is better off with a high type male, and a low type female is better off with a low type male and vice-versa. Or more generally when for each type on one side the best option on the other side is different.

[6] We say that the surplus exhibits absolute advantage if there is a single type on each side of the market that is preferred by everyone.

implied by our model.

We find that the surplus based on income and age exhibit absolute advantage and that the surplus based on education exhibits both relative and absolute advantage. This means that marrying someone with higher income is always better. In the case of age it means that women have a strong preference for older males independent of their own age, while men are virtually indifferent about the age of their spouse. For education, low educated people and high educated people prefer someone with their same level of education, displaying 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 correlation between the three observed matching rates and the recovered surpluses, 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 productive complementarities. 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 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 capturing 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 randomness in matching driven entirely by conscious strategic choices of agents, rather than by some unobserved characteristics with fixed distributions.

From a theoretical standpoint, the paper contributes to the search and matching literature by providing a framework that produces equilibrium outcomes between random matching and directed search, as opposed to nesting them. Examples of models that nest directed search and random matching are Menzio (2007) and Lester (2011).

Finally, the paper contributes to the literature on directed search and coordination frictions. The directed search paradigm generally predicts efficient static equilibrium outcomes. In contrast, our targeted search model does not appear to possess a market mechanism that can implement the constrained efficient allocation.

The paper proceeds as follows: Section 2 describes the model and derives the theoretical predictions, in Section 3 we take the model to the U.S. marriage market data and 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 men and women know the distribution over types on the other side of the market, but there is noise–they don't know who is who–and they can pay a cost to help them identify the type of all potential matches to maximize their payoff. We model the search process building on elements of information theory and the rational inattention literature, and assume that each agent can choose how much information to gather about the type of all potential matches. Given that information processing is costly, agents optimally choose a search strategy: a distribution of attention over all possible matches.[7]

Endogenizing the degree of randomness in matching means that agents should be able to decide how accurately they want to target a prospective match, and as such, choose a probability distribution over potential partners. This 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 models in economics that satisfy these requirements is very limited. The model closest to our aim is the model of information frictions used in the rational inattention literature. It uses the specification of Shannon's information as the measure of reduction in noise and attaches a cost to it. This specification both accommodates full choice of a distribution and a discrete choice problem. The measure of cost associated with a change in the distribution that this literature provides has a convenient interpretation as the amount of information processed by the agent. In addition, it turns

[7]Equivalently, we could describe agents as receiving costly signals about potential matches and choosing not only the precision of these signals, but the whole probability distribution.

out that in our specific case of a choice among discrete options, this specification enhances tractability and leads to close 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.

We think about the process of choosing a distribution of attention in the following way. Imagine that every individual carries out a screening process to reduce the noise about their potential matches' type. A more detailed screening entails a higher cost but allows the agent to have more certainty about who is their best match. In the end, men and women choose the total cost they want to pay by deciding how accurately they want to target their best matches (which depends on how much information they want to gather about their type) and as such they optimally choose the probability with which they will target each individual of the opposite sex. We refer to this process as choosing an optimal distribution of attention over types. A more concentrated distribution of attention entails a higher cost because it guarantees a higher probability of being matched with the type that renders a higher payoff. After choosing their optimal distribution of attention over types, both men and women make a draw from their distribution. If the draw is reciprocated, a match is formed if it is mutually beneficial and the surplus from the match is split between the two parties through ex-post Nash bargaining. As the search involves balancing the cost and the precision of information about prospective matches, some participants will not find partners immediately.

2.1 The environment

There are F females indexed by $x \in \{1, ..., F\}$ and M males indexed by $y \in \{1, ..., M\}$. Both males and females are actively searching for a match.

A match between female x and male y generates a surplus Φ_{xy}. If a male and a female match, the surplus is split between them by ex-post unilateral bargaining and we normalize the outside option of both to zero. We denote the fraction of the surplus that the female receives by ε_{xy} and the fraction of the surplus that the male receives by η_{xy}. The surplus and the split generated by any potential (x, y) match are known ex-ante to female x and male y.[8]

Each female chooses a distribution of attention which we denote $p_x(y)$ and it reflects the probability with which the female will target the male (ask him out). Each female 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 processing information about his type. Likewise we denote the strategy of a male $q_y(x)$. It represents the probability of a male y targeting a female x.

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 $q_1(x)$. As mentioned earlier, these probabilities constitute the distribution of attention $p_x(y)$ for females and the distribution of attention $q_y(x)$ for males. 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 from $p_x(y)$, female x targets male y; 2) according to the male's draw from $q_y(x)$, male y also targets female x; and 3) their payoffs are non-negative.

[8] The meeting process can be thought of as having two stages. In the first stage, links between males and females are formed. In the second stage they bargain over the surplus. Furthermore, we are not placing any restrictions on the surplus function.

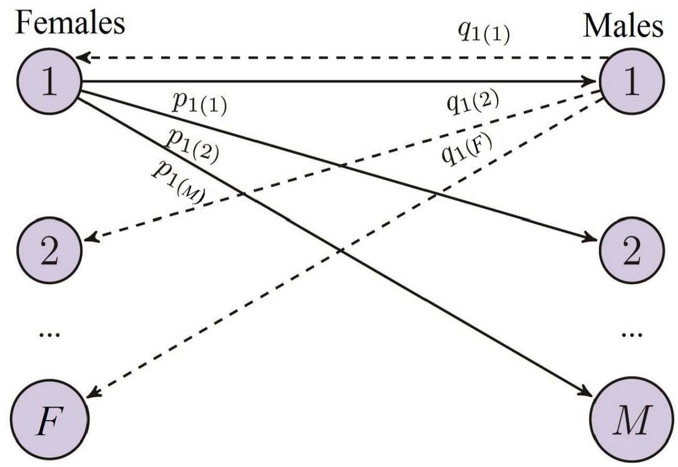

Figure 1. Strategies of males and females

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:

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

The female's cost of searching is denoted by $c_x(\kappa_x)$. This cost is a function of the amount of information processed by female x measured in bits, κ_x. Likewise, we denote a male's cost of searching by $c_y(\kappa_y)$, where the cost is a function of the male's information capacity, denoted κ_y.

We denote $\bar{\mu}_{yx}$ the equilibrium matching rate faced by female x when pursuing male y. It represents the female's perception of the probability that male y targets female x. Similarly, we denote $\bar{\nu}_{xy}$ the equilibrium matching rate faced by male y when considering female x. As matching rates are equilibrium objects, they are assumed to be common knowledge to participating parties, and equal to the distribution of attention of the counter-party in equilibrium.

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

$$Y_x = \max_{p_x(y)} \sum_{y=1}^{M} \varepsilon_{xy} \bar{\mu}_{yx} p_x(y) - c_x(\kappa_x).$$

The female gets her expected return from a match with male y conditional on matching with that male. The probability of a match between female x and male y is the product of the probability that female x targets male y and the probability that male y targets female x.

The search cost, $c_x(\kappa_x)$, is defined as follows:

$$\kappa_x = \sum_{y=1}^{M} p_x(y) \log_2 \frac{p_x(y)}{1/M}, \tag{1}$$

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

Our definition of information, κ_x, represents the relative entropy between a uniform prior $\{1/M\}$ over males and the posterior strategy, $p_x(y)$. This relative entropy can be interpreted as the reduction in the uncertainty of finding a partner that a female can achieve by choosing her distribution of attention. This definition is a special case of Shannon's channel capacity where information structure is the only choice variable.[9] Thus, our assumption is a special case of a uniformly accepted definition of information tailored to our problem.

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

$$Y_y = \max_{q_y(x)} \sum_{x=1}^{F} \eta_{xy} \bar{\nu}_{xy} q_y(x) - c_y(\kappa_y).$$

The male profits from a match with female x conditional on matching with her and pays the cost of search. The search cost on the male's side also

[9] See Thomas and Cover (1991), Chapter 2.

depends on the amount of information processed:

$$\kappa_y = \sum_{x=1}^{F} q_y(x) \log_2 \frac{q_y(x)}{1/F}, \tag{2}$$

where the female's strategy 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, $\{q_y(x)\}_{y=1}^{M}$, and matching rates $\{\bar{\mu}_{yx}\}_{y,x=1}^{M,F}$ and $\{\bar{\nu}_{xy}\}_{x,y=1}^{F,M}$ such that:*

1) strategies solve problems of males and females;

2) matching rates satisfy equilibrium conditions:

$$\bar{\nu}_{xy} = p_x(y), \qquad \bar{\mu}_{yx} = q_y(x). \tag{3}$$

Theorem 1. *A matching equilibrium exists.*

Proof. Note that if we substitute the matching rates (3) into the payoffs of males and females we can express the model as a normal-form game. The equilibrium of the matching model can be interpreted as a Nash equilibrium of this game. The set of distributions mapping compact sets into compact sets is a lattice under the natural ordering. Hence, all the results for lattices described by Vives (1990) apply to it. Since cross-derivatives of objective functions in our case are all non-negative, this game is super-modular. Hence there exists a Nash equilibrium. □

Theorem 2. *The matching equilibrium is unique, if*

a) cost functions are non-decreasing and convex;

b) $\left.\frac{\partial c_x(\kappa_x)}{\partial \kappa_x}\right|_{p_x^*(y)} > \varepsilon_{xy} p_x^*(y)$;

c) $\left.\frac{\partial c_y(\kappa_y)}{\partial \kappa_y}\right|_{q_y^*(x)} > \eta_{xy} q_y^*(x)$.

Proof. The payoffs of all males and females are continuous in their strategies. They are also concave in these strategies when cost functions are (weakly) increasing and convex in information capacities. "Diagonal dominance" conditions (b) and (c) guarantee that the Hessian of the game is negative definite along the equilibrium path. Then, by the generalized Poincare-Hopf index theorem of Simsek, Ozdaglar and Acemoglu (2007), the equilibrium is unique. □

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 marginal cost of information processing 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 marginal information 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 first perfectly identify all females to be able to check all pairwise deviations. Since acquiring that kind of information is very costly in our model, the equilibrium outcome generically does not satisfy "stability".

The result of Theorem 2 is intuitive. There are two motives for female x to target male y. The first motive is that male y may imply a higher payoff compared to other males, we refer to this as the productive motive. The second motive is that male y may have a higher probability to reciprocate, we call this the strategic motive. The payoff of a female depends on the product of the part she appropriates from the surplus and the probability of reciprocation. While her part of the split does not depend on equilibrium strategies, the strategic motive does. When costs of information are very low, females (and males) are able to place a high probability of targeting one counter-party and exclude all others. As a result, when information costs are extremely low, the strategic motive dominates. It does not matter what share of the surplus 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 nobody has an incentive to deviate from any mutual reciprocation.

As information costs increase, distributions of attention become less precise as it is increasingly costly to target a particular counter-party. That is, information processing constraints dampen the strategic motive and the productive motive starts playing a bigger role. At some threshold level of information costs 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. They 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 counter-party which promises a higher payoff, i.e. the productive motive dominates.

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:

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

$$q_y^*(x) = \exp\left(\frac{\eta_{xy} p_x^*(y)}{\frac{1}{\ln 2} \left.\frac{\partial c_y(\kappa_y)}{\partial \kappa_y}\right|_{q_y^*}}\right) \Big/ \sum_{m'=1}^{M} \exp\left(\frac{\eta_{x'y} p_{x'}^*(y)}{\frac{1}{\ln 2} \left.\frac{\partial c_y(\kappa_y)}{\partial \kappa_y}\right|_{q_y^*}}\right). \tag{4}$$

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 (4) 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 a best response to the strategies of females, and a female's strategy is a best response to the strategies of males. Theorem 2 predicts that an increase in information costs reduces the complementarities between search strategies of females and males. Once costs of information are 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 (4) and then check that diagonal dominance conditions are satisfied.

Now, consider the properties of equilibria for two limiting cases. First, as the marginal costs of processing information 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, consider the opposite case when marginal costs go to infinity. In this case, 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 costs of information are either very low or very high.

2.3 Efficiency

In order 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 surplus of the economy, which is a utilitarian welfare function. In order to achieve a social optimum, the planner can choose the strategies of males and females. If no costs of processing information were present, the planner would always choose to match each male with the female that produces the highest surplus. The socially optimal strategies of males would be infinitely precise.

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

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

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

Under the assumption of increasing 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:

$$\begin{aligned}
p_x^o\left(y\right) &= \exp\left(\frac{\Phi_{xy}q_y^o\left(x\right)}{\frac{1}{\ln 2}\left.\frac{\partial c_x(\kappa_x)}{\partial \kappa_x}\right|_{p_x^o}}\right)\Bigg/\sum_{f'=1}^{F}\exp\left(\frac{\Phi_{xy'}q_{y'}^o\left(x\right)}{\frac{1}{\ln 2}\left.\frac{\partial c_x(\kappa_x)}{\partial \kappa_x}\right|_{p_x^o}}\right), \\
q_y^o\left(x\right) &= \exp\left(\frac{\Phi_{xy}p_x^o\left(y\right)}{\frac{1}{\ln 2}\left.\frac{\partial c_y(\kappa_y)}{\partial \kappa_y}\right|_{q_y^o}}\right)\Bigg/\sum_{m'=1}^{M}\exp\left(\frac{\Phi_{x'y}p_{x'}^o\left(y\right)}{\frac{1}{\ln 2}\left.\frac{\partial c_y(\kappa_y)}{\partial \kappa_y}\right|_{q_y^o}}\right).
\end{aligned} \tag{5}$$

The structure of the social planner's solution is very similar to the structure of the decentralized equilibrium given by (4). 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 surplus, while in the decentralized equilibrium they only consider 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 identify his preferable matches. As a consequence, the females he targets will benefit (through an increase in the personal matching rate), and the females that 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 the male can not appropriate all the social benefits (the surplus of a match) he provides to society when increasing his search effort. The male only gets his bargained share of the surplus. The same statement is true for females. 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 surplus such that it implements the social optimum. When information costs are finite and positive, a socially optimal equilibrium has to satisfy the following conditions simultaneously:

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

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

Theorem 3. *The matching equilibrium is socially inefficient for any split of* *the surplus 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) $0 < \left.\frac{\partial c_x(\kappa_x)}{\partial \kappa_x}\right|_{p_x^*} < \infty$*;*

6) $0 < \left.\frac{\partial c_y(\kappa_y)}{\partial \kappa_y}\right|_{q_y^*} < \infty$*.*

Proof. See Appendix A. □

The first two conditions are self-explanatory; the case when all potential matches yield zero surplus is a trivial case of no gains from matching. Conditions 5 and 6 state that marginal costs of information have to be finite and positive in the neighborhood of the equilibrium. When costs of information are zero, the best equilibrium of the assignment model is socially optimal. When costs of information are very high, the random matching outcome is the best possible outcome. For all intermediate values of 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 attention 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 having all the bargaining power.[10]

One notable property of the equilibrium is that, by considering only fractions of the total surplus in choosing their strategies, males and females place lower probabilities on pursuing their best matches. This implies that in equi-

[10]See Appendix B for a version of the model with one-sided heterogeneity.

librium, attention of males and females is 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 application strategies. Because a female only considers her private gains from matching with a male, the female's optimal response would be to pay less attention to (target less accurately) the best males than it is socially optimal. In a second step, taking as given these strategies of females, males will be dis-incentivized not only by the fact that they consider fractions of the total gains from a match, but also by the fact that females pay less attention to them than it is socially optimal. These complementary dis-incentives will lower the probabilities of males pursuing their best match. Iterating in this way on strategies of males and females, at each step we get a reduction in the probability of targeting the best matches. 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 and females that they will get the whole surplus of each match and then collecting lump-sum taxes from both sides of the market to cover the costs of the program. Nevertheless, in order to do so, the planner himself would need to acquire extensive knowledge about the distribution of the surplus, which is costly. We leave this direction of research for future work.

2.4 Implications for Sorting

To better understand the effect of the productive and strategic motives, it is useful to consider simple examples of surpluses 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, and we label their types high (H) and low (L). Let us also consider two specific cases of the form of the surplus.

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 surplus 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 more generally refer to a surplus which has the same type as everybody's best option - 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 makes sure 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 attention in different directions.

If the surplus exhibits relative advantage, and the costs of reducing noise are low enough, our model can have two different equilibrium patterns. The first pattern is when the high type is more likely to pay attention to the high type, and the low type to 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 pay attention to 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 the costs of reducing noise are high, only the PAM equilibrium survives.

If the surplus exhibits *absolute* advantage, and the costs of reducing noise 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 pay more attention to the high type male, and both males pay more attention to the high type female. Moreover, for high enough costs of information, 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 the costs of information are low, the high types only look at each other, so it makes no sense for the low types to pay attention to the high types as, despite a higher potential payoff, the chance for their interest being reciprocated is zero. However, when information 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.

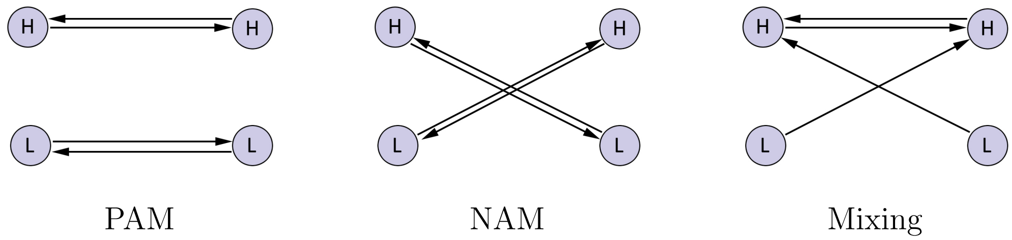

Figure 1: Three types of equilibria and their sorting patterns

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 surplus. 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 surplus 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 surplus 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 surplus. For the Monte Carlo simulations, we assume three males and three females, and draw each element of the 3-by-3 surplus matrix from a uniform distribution. We make 25,000 draws. We then find all equilibria and corresponding matching rates for each draw of the surplus. We discard all the surpluses 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 surplus matrix. In Figure 2 we show the probability density functions of correlations for three classes of surpluses: surpluses exhibiting absolute advantage, relative advantage, or no clear advantage pattern.

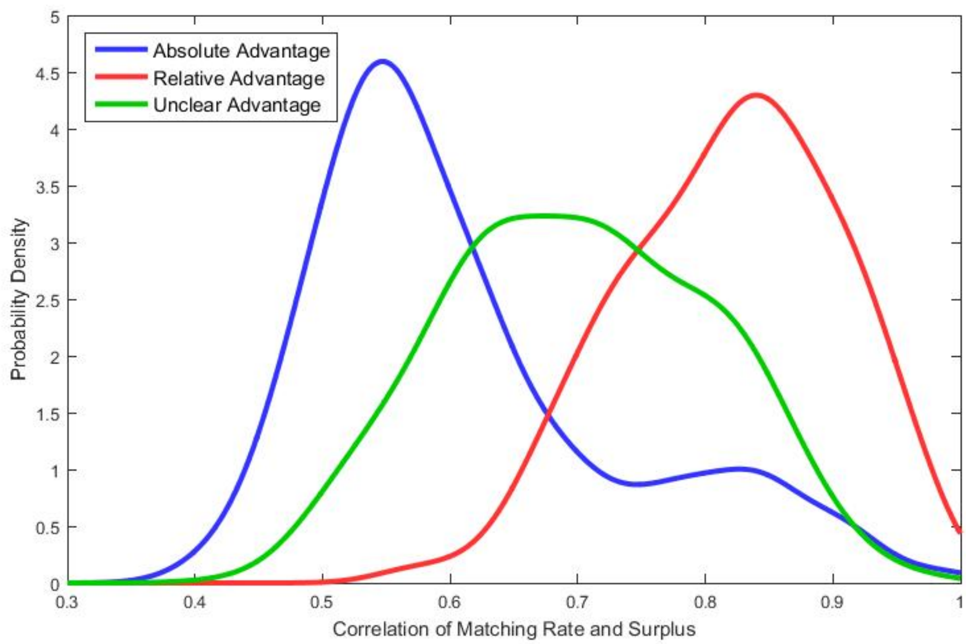

Figure 2: Correlation b/w Matching rate and Surplus

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 surplus 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 productive complementarities and the shape of the sorting pattern. Consequently, the empirical researcher could easily arrive at wrong conclusions about the shape of productive complementarities 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 show that this type of misspecification 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 and show that, when viewed through the lens of our model, they exhibit absolute advantage or a combination of absolute and relative advantage, and there is a substantial wedge between the shape of productive complementarities and the matching rate.

2.5 Invertibility

Our model builds on the interaction of strategic motives of agents and is hence more complicated computationally compared with leading examples in the literature. This fact has both bonuses and drawbacks. We find that in our model the mapping between the surplus and the matching rate is not invertible. In fact, there are matching rate patterns that can not be matched exactly by our model, and, potentially, there are matching rate patterns that could be matched exactly by different surpluses.

This implies that our model is testable. To illustrate this point we perform a Monte-Carlo exercise by drawing the elements of the 2-by-2 surplus 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.

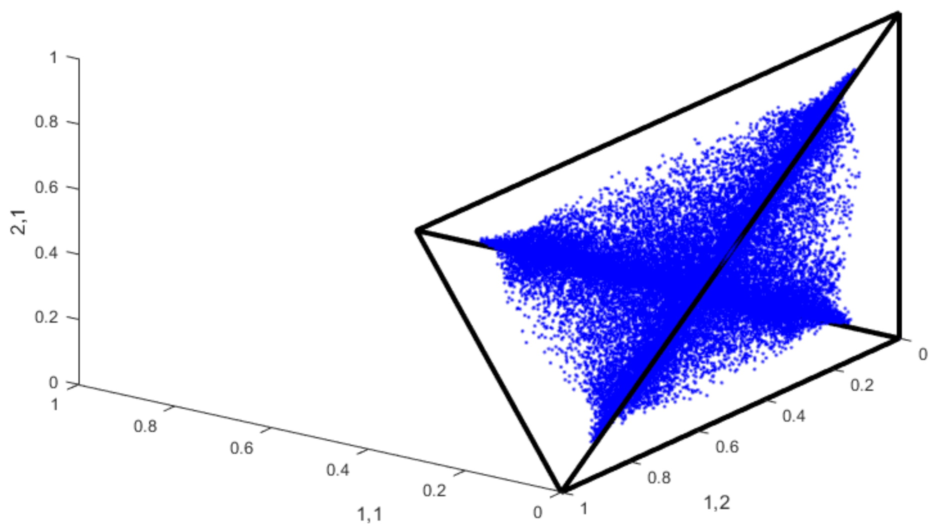

Figure 3: Simplex

We find that large white spaces remain in the simplex, implying that many shapes of matching rates simply 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 all prospective matches to just one of the males, and generate 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.

Given the non-invertibility of the mapping between the surplus and the matching rates, how can we test the model and estimate the surplus from data on matching rates? For any shape of the surplus, we can compute the matching rates implied by the model. We can then search for a shape of the surplus that minimizes some measure of distance between the predicted and observed matching rates. One natural measure of distance is the likelihood function of the data given the predicted matching rates. Maximization of such a likelihood function efficiently minimizes the properly weighted sum of distances between the the data and the model's prediction, and should have nice properties. The results of such estimation can be treated as an upper bound on the 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

In order 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 men and women, as well as newly married couples comes from IPUMS for the year 2001.[11] For computational transparency we attribute both men and women to three equally sized bins, which we refer to as low (L), medium (M) and high (H) types. We consider three dimensions along which men and women evaluate

[11] We thank the authors of Gayle and Shephard (2015) for kindly sharing the IPUMS data with us.

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.

Income is a continuous characteristic, so the three bins correspond to people with low, medium and high incomes. 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 which 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 that never went to college, those that are currently in college, and those that have graduated from college.

For each of the three cases, we estimate the shape of the surplus using the maximum likelihood methodology described earlier. We assume that all currently unmarried men and women are searching, and the number of matches is proxied by the number of couples that got married in the past 12 months, as indicated by answers to the questionnaire. The data contains roughly 93599 unmarried males, 82673 unmarried females, and 23572 newly married couples above the age of 21. Whenever our model produces multiple equilibria we select the one that fits the observed matching rate best.

The matching rate for the case of income is presented in the left panel of Figure 4. The estimate of the underlying surplus is shown in the right panel of the same Figure. A notable property of the surplus 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 surplus have a correlation of 0.72.

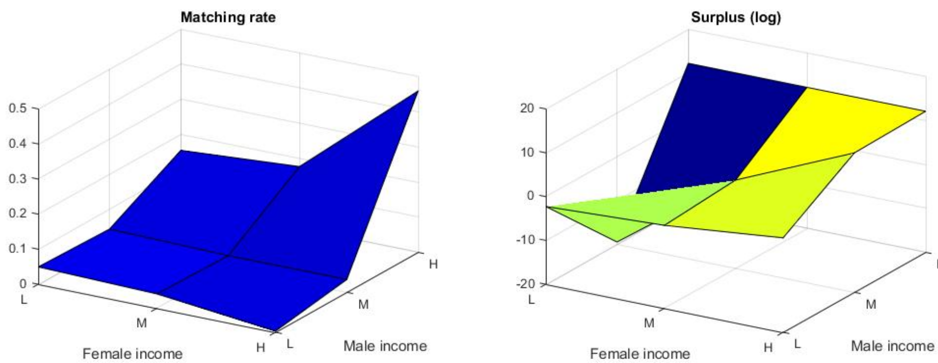

Figure 4: Sorting by income

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

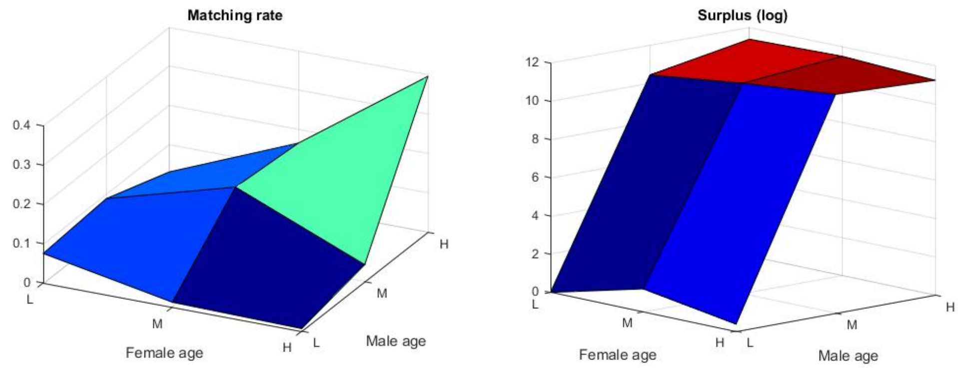

Figure 5: Sorting by age

The matching rate for sorting by education is presented in the left panel of Figure 6. In this case the surplus exhibits a combination of absolute and relative advantage. Low educated people and high educated people prefer someone with their same level of education, displaying 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 surplus have a correlation of 0.52.

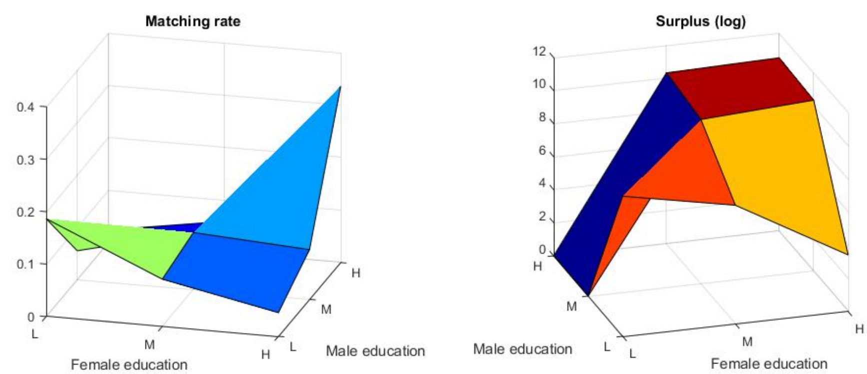

Figure 6: Sorting by education

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 men and women reaches an implicit agreement on the matching rate among themselves, matching (or staying single) is a voluntary decision. In their model, the surplus 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 surplus there is a unique matching rate, and for any matching rate one can invert the relationship to compute the surplus.

The second notable property is that the matching rate only depends 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 surplus. 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 surplus. 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 surplus is 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 surplus and the equilibrium matching rate obtained from our model - on the horizontal axis, and the correlation between the same matching rate and the corresponding surplus recovered by the model of Choo and Sioux on the vertical axis. We find that in many cases, the shape of the true surplus and of the matching rate goes all the way down 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 surpluses 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 surpluses with absolute advantage.

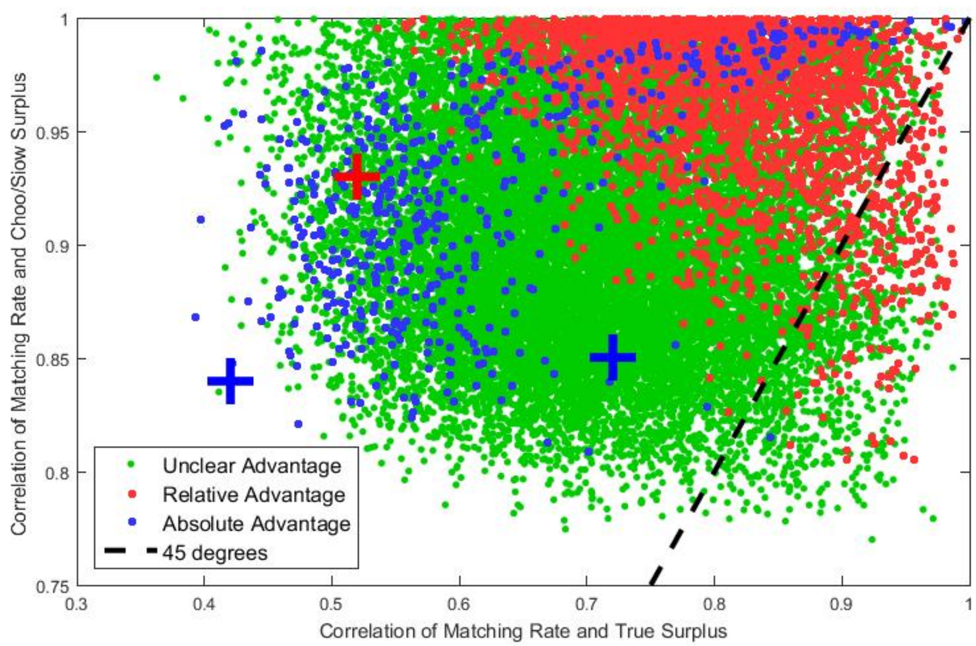

Figure 7: Monte Carlo results and Data

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 have to pay a cost to reduce the noise level associated with distinguishing among potential matches. If they pay a higher cost, 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 information choice 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 acquire information. 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 of surplus transfers. 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] **Becker, G. S., (1973).** "A Theory of Marriage: Part I." Journal of Political Economy, 81(4), pp. 813-846.

[2] **Burdett, K. and M. G. Coles, (1999).** "Long-Term Partnership Formation: Marriage and Employment." Economic Journal, Royal Economic Society, 109(456), pp. F307-F334.

[3] **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.

[4] **Choo, E. and A. Siow, (2006).** "Who Marries Whom and Why." Journal of Political Economy, 114(1), pp. 175-201.

[5] **Diamond, P. A., (1982).** "Wage Determination and Efficiency in Search Equilibrium." *Review of Economic Studies*, 49(2), pp. 217-227.

[6] **Galichon, A. and B. Salanié, (2012).** "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models." Working Paper hal-01053710

[7] **Gayle, G.-L. and A. Shephard, (2015).** "Optimal Taxation, Marriage, Home Production, and Family Labor Supply." October 2015, University of Pennsylvania Working Paper.

[8] **Lester, B. (2011).** "Information and Prices with Capacity Constraints," *American Economic Review*, 101(4), pp. 1591-1600.

[9] **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.

[10] **Menzio, G. (2007).** "A Theory of Partially Directed Search," *Journal of Political Economy*, 115(5), pp. 748–769.

[11] **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.

[12] **Shimer, R. and L. Smith, (2000).** "Assortative Matching and Search." *Econometrica*, 68(2), pp. 343-370.

[13] **Shimer, R. (2005).** "The Assignment of Workers to Jobs in an Economy with Coordination Frictions." *Journal of Political Economy*, 113, pp. 996-1025.

[14] **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.

[15] **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.

[16] **Thomas, C. and J. Cover, (1991)**. "Elements of Information Theory". John Wiley & Sons, Inc.

[17] **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:

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

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

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

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

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

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

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

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

Then, by construction,

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

and

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

Denote them a_y and a_x respectively.

It then follows that:

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

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

$$(i) \qquad M \exp\left(1 + \tfrac{\tilde{\lambda}_y}{a_y/\ln 2}\right) = \sum_{x=1}^{M} \exp\left(\tfrac{\Phi_{xy}p_x(y)}{a_y/\ln 2}\right)$$

$$(ii) \qquad M \exp\left(1 + \tfrac{\lambda_y}{a_y/\ln 2}\right) = \sum_{x=1}^{M} \exp\left(\tfrac{\varepsilon_{xy}(x)p_x(y)}{a_y/\ln 2}\right)$$

$$(iii) \qquad \left(\Phi_{xy} - \varepsilon_{xy}\right)p_x\left(y\right) = \tilde{\lambda}_y - \lambda_y \qquad \text{for all } x$$

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

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

Hence,

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

Therefore, either:

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

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

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

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

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

Cases b) and d) have been ruled out by the assumptions of the theorem. Cases a) and b) 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 information processing constraints. 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 pursuing male y. It is also the female's distribution of attention. We assume that each female can rationally choose her strategy facing a trade-off between a higher payoff and a higher cost of processing information.

A female's cost of searching is given by $c_x(\kappa_x)$. This cost is a function of the amount of information processed by a female measured in bits, κ_x. Once the optimal distribution $p_x(y)$ is chosen, each female draws from it to determine which male to contact.[12]

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 surplus in a match with male y conditional on matching with that male. She also incurs a search cost, which depends on the information processing capacity defined as follows:

$$\kappa_x = \sum_{y=1}^{M} p_x(y) \log_2 \frac{p_x(y)}{1/M} \tag{6}$$

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 in information capacities. Hence, each problem has a unique solution. □

[12]As 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.[13] Rearranging the first order conditions for the buyer, we obtain:

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

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

The equilibrium condition (7) 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 of pursuing 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 information processing constraints that we place on females. Thus, the social planner maximizes the following welfare function:

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

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

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

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

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

subject to the information constraint (6) 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:

$$p_x^o(y) = \exp\left(\frac{\Phi_{xy} q_y}{\frac{1}{\ln 2} \left.\frac{\partial c_x(\kappa_x)}{\partial \kappa_x}\right|_{p_x^o}}\right) / \sum_{y'=1}^{M} \exp\left(\frac{\Phi_{xy'} q_y}{\frac{1}{\ln 2} \left.\frac{\partial c_x(\kappa_x)}{\partial \kappa_x}\right|_{p_x^o}}\right). \tag{8}$$

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 pursuing 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 surplus 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 surplus to the females, $\varepsilon_{xy} = \Phi_{xy}$, effectively assigning them a bargaining power 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 costs of information are absent, the equilibrium of the model is socially optimal. When costs of information 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.