Distributed Matching with Mixed Maximum-Minimum Utilities

Amos Azaria

Computer Science Dept., Ariel University
Ariel 40700, Israel
amos.azaria@ariel.ac.il
Tel.:+1-929-293-8784

David Sarne

Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel

Yonatan Aumann

Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel

---

Abstract

In this paper we study distributed agent matching with search friction in environments characterized by costly exploration, where each agent's utility from forming a partnership is influenced by some linear combination of the maximum and the minimum among the two agents' competence. The paper provides a cohesive analysis for such case, proving the equilibrium structure for the different min-max linear combinations that may be used. The paper presents an extensive equilibrium analysis of such settings, proving three distinct resulting patterns of the acceptance thresholds used by the different agents. The first relates to settings where a greater emphasis is placed on the minimum type, or in the extreme case where the minimum type solely determines the output. In these cases, the assortative matching characteristic, holds, where all agents set their threshold below their own type and the greater is the agent type the greater is its threshold. When the utility from the partnership formation is solely determined by the maximum type, we show that there exists a type x^* where partnerships form if and only if one of the agents has a type equal or greater than x^*. When a greater emphasis is placed on the maximum type (but not only), we prove that assortative matching never holds, and the change in the agents' acceptance thresholds can frequently shift from an increase to a decrease.

Category. Computing methodologies Multi-agent systems;300

Keywords. Equilibrium, Distributed Matching, Two-Sided Search

1. Introduction

Distributed matching with search frictions (or "distributed matching" for short), is an important and highly applicable model of agents searching for partners in multi-agent systems. It is used in settings where no central information source can supply instant reliable information on the environment and on partnering opportunities within. In such settings, standard stable matching mechanisms [16, 2] cannot be applied.

The goal of each agent participating in the distributed matching process is to form a beneficial pairwise partnership [3, 13, 5, 27, 26]. Each agent is associated with a specific *type* that captures its properties (e.g., competence, wealth). During each stage of the process, agents randomly interact pairwise and learn each other's type. The process of initiating and maintaining an interaction is associated with a cost, termed *search cost*, incurred by both agents. In order for a partnership to be formed, it needs to be accepted by both agents. The agents thus need to consider, when deciding whether to commit to a given partnership, the tradeoff between the potential benefits from continuing the exploration, possibly forming a better partnership in the future, versus the costs associated with this exploration.

During the past two decades, there has been a substantial progress in the analysis of distributed matching models (see [38] for a survey). Works in this area typically differ in the assumption that they make about the utility that agents obtain from a partnership. The choice of the utility function affects the structure of the acceptance thresholds used by the different agents. For example, if the utility depends exclusively on the other agent's type the equilibrium can be characterized as a "perfect segregation", i.e., the agents form clusters, based on their type, in which every agent in a cluster is always willing to form a partnership with any other agent in the cluster [6, 8, 26]. For most common functions where the utility depends on both agent types, the resulting equilibrium can be characterized as "assortative matching", i.e., the acceptance thresholds used, increase in the agent's type [36, 22].

In this paper, we consider a different utility function for a match; one that is based on the minimum and maximum of the two types. This function is highly applicable to situations where the performance of the partnership formed depends mostly on its most or least competent member. For example, consider a group of students that need to form pairwise partnerships for the purpose of working on a course project. In this case, the grade that any team receives highly depends on the capabilities of the more competent student among the two. Alternatively, consider tennis players that seek partners when playing doubles. Here the players are rewarded exclusively based on the team's (rather than the individual's) performance. The performance of the team in this case will be primarily affected by the least competent player, as the other team will try to gain game points primarily by aiming the ball in his direction. Interestingly, in many cases the use of the new reward function leads to equilibrium matching structures that are different than those resulting from the traditional reward functions in distributed matching.

The paper presents an extensive analysis of the model with the new min-max utility function. We show that for the case where the utility function relies exclusively on the maximum type among the two, the equilibrium is characterized by a single threshold, where all agents of types greater than or equal to the threshold accept any agent, and all agents of types smaller than the threshold accept only agents of types greater than the threshold. For the case where the utility function relies exclusively on the minimum type, we show that the equilibrium is characterized by assortative matching. When the utility is a function of both the minimum and maximum, we manage to distinguish between three equilibrium patterns. Namely, when the minimum type has greater input on the agents outcome the equilibrium has a form of assortative matching. When the minimum and maximum types have equal impact, the equilibrium has a form of perfect segregation. When the maximum type has greater impact we observe an interesting pattern which we term 'bumpy steps' where the acceptance threshold as a function of the agent type increases and decreases alternately. We further show that, in equilibrium, an agent that sets its threshold below its own type is necessarily accepted by all agents with types between that threshold and the agent's type. These results facilitate the calculation of the agents' equilibrium strategies in discrete settings, as they enable the use of standard dynamic programming techniques for calculating each type's equilibrium strategy based on the strategies of higher types.

One important use of the new utility function is for mechanism design. We conclude this paper with an example in which this min-max utility function may be used to increase a designer's utility. As an example, we analyze a situation in which a professor requires her students to complete some task in pairs and would like to increase the students' learning rate (defined as the distance between the two students' types). We show that by applying a maximum function to determine the grades of her students the professor can achieve a much higher learning rate than by using other functions.

2. Related Work

Distributed matching is a sub-domain of coalition formation. While coalition formation models usually consider general coalition-sizes [41], the partnership formation model (often referred as matchmaking) considers environments where agents have a benefit only when forming a partnership and this benefit cannot be improved by extending the partnership to more than two agents [19, 40] (e.g., in the case of buyers and sellers or peer-to-peer applications). Various centralized matching mechanisms can be found in the literature [14, 4, 16]. Such mechanisms are indeed in use in some environments, such as the National Resident Matching Program (NRMP), which matches medical school students with residency programs [29, 30, 31, 42] and similarly, The ASHP Resident Matching Program which does so with pharmacists. However, in many multi-agent system environments, in the absence of any reliable central matching mechanism, the matching process is completely distributed.

The core analysis of distributed matching models relies on understanding the search strategy of each individual. These have emerged primarily from one-sided search models [21, 11, 24]. These models have developed to a point where their total contribution is referred to as "search theory". Over the years, many one-sided search model variants have been considered, differing in the decision horizon (finite versus infinite) [21], the presence of the recall option [24], the distribution of values, the assumptions made on the search costs [17], the certainty of findings [1, 12] and the extent to which findings remain valid along the search [20]. The model was even extended to the case of cooperative search, where the search is carried out for the benefit of a group of agents [23, 35] or when the agents are fully self-interested [28]. Within the framework of one-sided search theory, two main clusters of search models can be found: (a) the sequential search model, which is the primary model also in distributed matching, and (b) the fixed sample size model (also known as "parallel" model). In the sequential search model [32, 21], the searcher obtains a single value at a time, allowing multiple search stages. In the fixed sample size model, the searcher obtains a large set of values in a single search round [18, 39] and then chooses the best value from those obtained. Another well-known one-sided search problem, which is also a stopping problem, is the "Secretary Problem" [15]. In this problem a searcher observes candidates sequentially, and must decide when to end the search and hire the last candidate. However, there are additional substantial differences between that model and ours. These relate to the underlying search problem. In the secretary problem the searcher becomes acquainted with a candidate's relative ranking compared to those already interviewed whereas in our case the searcher obtains an actual (absolute) value (e.g., utility or profit). Also, in the secretary problem there are no search costs, hence the goal is typically to maximize the probability of stopping with the best candidate (or to minimize the expected rank), whereas in our case the goal is to maximize the process as a whole, taking the search costs into account.

In an effort to understand the effect of dual search activities in costly environments, distributed matching search models have been developed [3, 7, 33]. These search models are distinguished according to several assumptions they make. The first is the payoff utility each agent obtains from each partnership. While some of these models assume that the utility is exclusively a function of the other agent's type [26, 6], others assume a function defined over both types [36]. See Table 1 for a list of works and the utility functions which they have considered.

The second is the way according to which the search friction (cost) is modeled. This can be either the discounting of future flow of gains [6] or additive explicit search

Paper	Utility function
[6, 5, 26]	$u_x(x,y) = y$, $u_y(x,y) = x$
[8]	$u_x(x,y) = \alpha_1(x) + \alpha_2(y)$ $u_y(x,y) = \alpha_3(x) + \alpha_4(y)$ where $\alpha_{1\ldots4}$ are monotone strictly increasing functions
[36]	Describe sufficient conditions on the utility function to form assortative matching and discuss other utility functions
[22]	Product functions, such as: $u_x(x,y) = x\dot{y}$
This paper	$u_x(x,y) = u_y(x,y) = \alpha \min\{x,y\} + (1-\alpha)\max\{x,y\}$

Table 1: List of related work and the utility functions which they have considered. $u_x(\cdot)$ denotes the utility of the first agent, and $u_y(\cdot)$ of the second (these utilities must equalize in a single population environment).

costs [26, 8, 3]. Another distinction made in the models is whether partnerships can form only between two different populations (e.g. men and woman) [8], or a single population (e.g. students) [10]. Lastly, the models are distinguished by the nature of the utility earned by each of the agents (transferable [3] and non-transferable [6, 8]). Our model assumes non-transferable utilities, explicit search costs and a payoff that combines the minimum and maximum type in a pre-defined manner. To the best of our knowledge, such a model has not been investigated to date.

For most distributed matching models, it has been shown that, in equilibrium, the optimal strategy for agents is a reservation value strategy [7, 34, 37, 25, 8, 26]. In particular, it has been shown that with no costs of search nor discounting of gains the unique equilibrium strategy is *perfect assortative matching*, which means that every agent accepts only its own type and above, implying that partnerships are formed only between agents with the exact same type [3, 8, 26].[1] When search is costly and/or discounting of gains is applied, matching is no longer perfect assortative, as agents will inevitably widen the set of mates they accept. In this case, when utility is determined solely by the type of the other agent, it was shown that the equilibrium structure is of a "block segregation"-like structure: an interval of "highest" individuals match only with each other, the next highest match only with each other, and so on. This has been shown both for the case with explicit search costs [8], with discounting of gains [37]

[1]This can also be seen as an extension of Becker's classic result for the transferable utility matching market with no complementaries in production (supermodularity of the joint production function) [4].

and for a combination of the two [25].

3. The Model

The model used in this paper is a standard distributed matching model, and all assumptions given in this section are common in the distributed matching literature [8, 6, 26, 36, 22]. The differentiating element is the utility function used in this paper, which we assume to be a function of the minimum and maximum among the interacting types.

We consider an environment populated by an infinitely many self-interested, fully rational agents. Each agent is associated with a *type* x, which is a real number in the interval $[0, 1]$. The distribution of types in the environment is defined by a p.d.f. $f(x)$ (where $0 \leq f(x) < \infty$ for any x), and a c.d.f. $F(x)$ (where $F(1) = 1$). Each agent can form a partnership with any other agent in the environment. The utility that an agent of type x obtains from forming a partnership with agent of type y is $u(x, y)$. This function is continuous and monotonic non-decreasing in x and y. In this work we focus on the case where the utility is of the form:

$$u(x, y) = u(y, x) = \alpha \min\{x, y\} + (1 - \alpha) \max\{x, y\}, \tag{1}$$

for some $\alpha \in [0, 1]$.

The agents are assumed to be acquainted with the type distribution function $f(x)$. However, they cannot tell a-priori what is the type of any specific agent in their environment. The only way by which an agent can learn the type of another agent is by interacting with it. Since neither agent has prior information concerning the type of any specific other agent, it initiates interactions with other agents randomly.

After two agents interact and learn each other's type, each of them needs to decide whether to accept or reject a partnership with the other agent. If a dual acceptance is reached, the partnership is formed. Otherwise, both agents resume their search. The model assumes that the agents are satisfied with having a single partner. Thus, once a partnership is formed the two agents forming it terminate their search process and leave the environment and are replaced with two new agents of the same types. In other words, the size and the composition of the market are exogenously given.

The search activity/interaction is assumed to be costly. Each search stage incurs a cost c to each of the agents. We assume utilities and costs are additive and that the agents try to maximize their overall utility, defined as the utility obtained from the partnership minus the search costs accumulated along the search process.

Naturally the attempt to analyze day-to-day applications using "search theory" techniques brings up the applicability question. Justification and legitimacy considerations for the different assumptions of the "standard" model are widely discussed in the wide literature we refer to. The current paper is not focused on re-arguing applicability, but rather on the analysis of the core two-sided search model with the new (and in many cases more applicable) reward functions.

3.1. Notation

We use s_ζ to denote a strategy of agent ζ. Given all previous interactions of agent ζ, the agent which ζ is interacting with and its type, the strategy, S determines whether ζ is willing to commit to a partnership or not.

We use S to denote a *strategy profile*, that is, a set of strategies - one for each agent. We frequently call a strategy profile simply a strategy. Given $S_{-\zeta}$ - a set of strategies of all agents except ζ - strategy s_ζ of ζ is called *best-response* (or *optimal*), if no other strategy s'_ζ gives a strictly higher expected outcome to ζ. A strategy S is said to be *in equilibrium* if all agents follow their best response. We assume that an agent that is indifferent between accepting a partnership and resuming the search, accepts the partnership. We make this assumption to avoid the trivial equilibrium in which all agents reject all others. Since the distribution over the types remains identical between rounds and there is no limit over the decision horizon, agents best response, and hence equilibrium, strategies do not depend on previous rounds or on the type of interactions experienced in previous rounds (nor their outcomes). In fact, in equilibrium, agents' strategies differ only based on their type, since all agents of the same type have the same utility from a partnership with any other given type, their decision to terminate search and partner with the current potential partner is identical.

We use $\phi(x, S)$ to denote the expected utility from the search of an agent of type x assuming all agents adhere to their associated strategy in S. Once again, due to the memory-less nature of the problem, $\phi(x, S)$ equals the expected utility of the agent after rejecting a partner (or being rejected by it) and resuming the search (note that search costs incurred up to that point are "sunk costs" and do not influence further decisions).

We define $\underline{a}(x, S) = \min\{y|\phi(x, S) \leq u(x, y)\}$ as the minimum type for which the utility from accepting it by agent of type x is greater or equal to the expected utility of the agent if resuming the search. For exposition purposes, we occasionally omit S from the above notation whenever it represents a set S, which is in equilibrium. We later show that in equilibrium each agent of type x accepts all and only agents with types greater than or equal to $\underline{a}(x)$.

Occasionally, we will use $\Phi(x, S, t)$ to denote the expected utility of an agent of type x when using a strategy according to which it accepts other agents only if their type is above t (this might not be optimal) and all other agents adhere to their strategy in S. We will use $a(x)$ (rather than $\underline{a}(x)$) as a strategy for an agent of type x to indicate that the agent accepts all and only agents in $[a(x), 1]$; note that this might not be an optimal strategy for the agent (and therefore not in equilibrium).

We use $A(x, S)$ to denote the set of agent types that accept a match with an agent of type x, when following the strategies in S. We denote $\overline{A}(x)$ such that $\overline{A}(x, S) = \{y \mid y \in A(x, S) \wedge y > x\}$.

We use $M(x, S)$ to denote the set of types with which an agent of type x can end up forming a partnership with, when following strategies in S, that is, $M(x, S) = \{y \mid y \in A(x, S) \wedge x \in A(y, S)\}$. We call this set the "matching set" of x. $\overline{M}(x, S)$ denotes the set of all types above x which are in the matching set of x, i.e. $\overline{M}(x, S) = \{y \mid y \in M(x, S) \wedge y > x\}$, and $\underline{M}(x, S)$ denotes the set of all types below x which are in the matching set of x, i.e. $\underline{M}(x, S) = \{y \mid y \in M(x, S) \wedge y < x\}$.

4. Individual Strategies and Expected Utility

We begin with the analysis of individual expected-utility-maximizing strategies. All claims of this section are not limited to the min-max utility function, but hold for any case of strictly monotone and symmetric u, i.e., any u in which $u(y,x) = u(x,y) < u(x+\epsilon_1, y+\epsilon_2)$ for any $\epsilon_1 > 0, \epsilon_2 > 0$. Proposition 1 establishes that the expected-utility-maximizing strategies are reservation-value (i.e., threshold) based and that the threshold for an agent of type x is exactly $\underline{a}(x)$. Our assumption that an agent that is indifferent between accepting and rejecting a partner must accept, entails that, in equilibrium, all agents follow the threshold based strategy described in Proposition 1. Proposition 2 shows that the expected utility of any individual agent when using such a strategy equals the reservation-value used.

Proposition 1. *Given a strategy profile S (not necessarily in equilibrium), an agent of type x accepting all and only agents of types in $[\underline{a}(x,S), 1]$ will maximize its expected utility.*

Proof. Assume that agent of type x uses an optimal strategy s, which indicates which partners to accept/reject. On iteration i of the search, if the partner rejects forming a partnership, then no partnership will be formed, and therefore the expected utility of the agent does not depend on the agent's strategy. Assuming the partner accepts the partnership: if the agent interacts with a partner of type $x' < \underline{a}(x,S)$, the agent may reject it (regardless of s) since in the next iteration the agent may resume using s as its strategy[2], and therefore gains $\phi(x,S)$ which, due to the monotonicity of u and definition of $\underline{a}(x,S)$, is greater than $u(x,x')$. If the agent interacts with a partner of type $x' \geq \underline{a}(x,S)$, the agent may accept it (regardless of s); the agent does not lose out by doing so, due to the monotonicity of u and definition of $\underline{a}(x,S)$, $u(x,x') \geq \phi(x,S)$.

This is true for any iteration i, therefore the agent can always reject agents with a type less than $\underline{a}(x,S)$ and accept agents with a type greater than or equal to $\underline{a}(x,S)$.

□

[2]Once giving up on the current value, the agent faces the exact same problem as in the beginning, hence the expected benefit on-wards is $\phi(x,S)$.

Proposition 2. *Given a strategy profile S, and assuming an agent follows its optimal strategy, if $\underline{a}(x, S) > 0$ then $\phi(x, S) = u(x, \underline{a}(x, S))$.*

Proof. $\underline{a}(x, S) > 0$ therefore $\phi(x, S) > u(x, 0)$. Since u is monotone, $\phi(x, S) \leq u(x, 1)$. According to the intermediate value theorem, there exists a type y where $\phi(x, S) = u(x, y)$. From the definition of $\underline{a}(x, S)$ we obtain that $\underline{a}(x, S) \leq y$. Assume by contradiction that $\underline{a}(x, S) = y' < y$; due to the monotonicity of u, $u(x, y') < u(x, y) = \phi(x, S)$, however, from Proposition 1, agent of type x accepts partnership with agent of type y', but, agent with type x, is better off rejecting partnership with an agent with type y' and resume the search in which the expected utility is $\phi(x, S)$. $\square$

Now that we have established the use of the reservation-value based strategies by all agents and the indifference property of the reservation value, we can prove that an agent that sets its threshold below its own type is necessarily accepted by all agents with types between that threshold and the agent's type. Based on this theorem we will later propose an algorithm that, in a discrete case, finds the acceptance threshold for all agent types. Before proving the main theorem, we first prove two lemmas that simplify the proof of the theorem.

The following lemma establishes that, in equilibrium, the higher the agent type, the larger the set of agent types that accept it as a partner:

Lemma 1. *Given strategy profile S* **in equilibrium***, for any x and $x' < x$, $A(x', S) \subset A(x, S)$.*

Proof. Given $\tilde{x} \in A(x', S)$ (i.e. an agent with type $\tilde{x}$ accepts x' as a partner) according to Proposition 1 $\underline{a}(\tilde{x}, S) \leq x'$. Therefore $\underline{a}(\tilde{x}, S) \leq x$ and again according to Proposition 1 $\tilde{x} \in A(x, S)$. $\square$

The following lemma establishes that, in equilibrium, the higher the agent type, the higher its expected utility:

Lemma 2. *Given a strategy profile S* **in equilibrium***, for any x and $x' < x$, $\phi(x', S) \leq \phi(x, S)$.*

Proof. From Lemma 1 $A(x', S) \subset A(x, S)$, therefore x can guarantee $M(x, S) = M(x')$ by rejecting any agent not in $M(x', S)$. Now, for every $\tilde{x} \in M(x')$ by the monotonicity of u, $u(x', \tilde{x}) \leq u(x, \tilde{x})$, therefore, $\phi(x', S) \leq \phi(x, S)$. $\square$

At this point, we have everything needed in order to prove Theorem 1.

Theorem 1. *Given a strategy profile S* **in equilibrium***, for any x and $x' < x$, if $\underline{a}(x, S) \leq x'$ then $\underline{a}(x', S) \leq x$.*

Proof. If $\underline{a}(x, S) > 0$, by Proposition 2 we get that $\phi(x, S) = u(\underline{a}(x), x)$. However, if $\underline{a}(x, S) = 0$ we know that $\phi(x, S) \leq u(0, x)$, otherwise (i.e. $\phi(x, S) > u(0, x)$) the agent would reject the partnership and continue its search. Therefore:

$$\phi(x, S) \leq u(\underline{a}(x, S), x).$$

For a contradiction assume that $x < \underline{a}(x', S)$. Since $x < \underline{a}(x', S)$ and $\underline{a}(x, S) \leq x'$, and due to the strictly monotonicity and symmetry of u:

$$u(\underline{a}(x), x) < u(x', \underline{a}(x', S)).$$

$x < \underline{a}(x', S)$ entails that $0 < \underline{a}(x', S)$, therefore by Proposition 2 we get:

$$u(x', \underline{a}(x', S)) = \phi(x', S),$$

resulting in:

$$\phi(x, S) < \phi(x', S),$$

which contradicts Lemma 2.

$\square$

From Theorem 1 we conclude that, in equilibrium, it is guaranteed that all agents of lower types that an agent accepts will accept it as well. Therefore, in equilibrium, the acceptance threshold of each agent does *not* depend on the acceptance threshold of agents of types lower than its own. This entails that, in equilibrium, $\underline{M}(x) = [\underline{a}(x), x)$, which facilitates the analysis given in the following sections.

5. Minimum as the Utility Function

We begin by analyzing the equilibrium in settings where $\alpha = 1$, that is, the utility from forming a partnership between any two agents is the minimum value among the two: $u(x, y) = \min\{x, y\}$.

Clearly, when the utility is the minimum, each agent accepts all other agents of type equal to or greater than its own, since its utility is anyhow bounded by its own type. In-fact, an agent of type x accepts any other agent of type at least $x - c$ (since resuming the search incurs a cost c).

An immediate corollary resulting from Theorem 1 and Proposition 1 is that, in equilibrium, if $\underline{a}(x) < x$ (i.e. an agent sets its reservation value beneath its own type) then the matching set, $M(x, S)$, according to Proposition 1 equals $A(x) \cap [\underline{a}(x, S), 1]$ which in turn according to Theorem 1 equals $[\underline{a}(x, S), x] \cup \overline{A}(x, S)$. Therefore, when all other agents are assumed to be using strategies which are in equilibrium, when finding the threshold which maximizes the expected utility for agent of type x, it is sufficient to assume that $\underline{M}(x, S) = [\underline{a}(x, S), x]$. Given a strategy profile S, in equilibrium, the expected payoff for an agent of type x from the search, when using a threshold t, (and $\alpha = 0$, i.e. $u(x, y) = \min\{x, y\}$) is given by:

$$\begin{aligned}\Phi(x, t, S) = &- c + \int_{\underline{M}(x,S)} y f(y) dy + \int_{\overline{M}(x,S)} x f(y) dy + \\ &\Phi(x, t)\big(1 - \int_{M(x,S)} f(y) dy\big).\end{aligned} \tag{2}$$

Where the first integrand is for a case where the agent forms a partnership with a partner of a lesser type (than its own), the second is for the case where the agent forms a partnership with a partner with a greater type than its own and the third is for the case where the agent does not form a partnership and resumes the exploration.

From Equation 2 and Theorem 1 we obtain, that the agent will set its threshold which maximizes:

$$\underline{a}(x, S) = \arg\max_t \frac{-c + \int_{y=t}^{x} y f(y) dy + \int_{\overline{M}(x,S)} x f(y) dy}{\int_{M(x,S)} f(y) dy}. \tag{3}$$

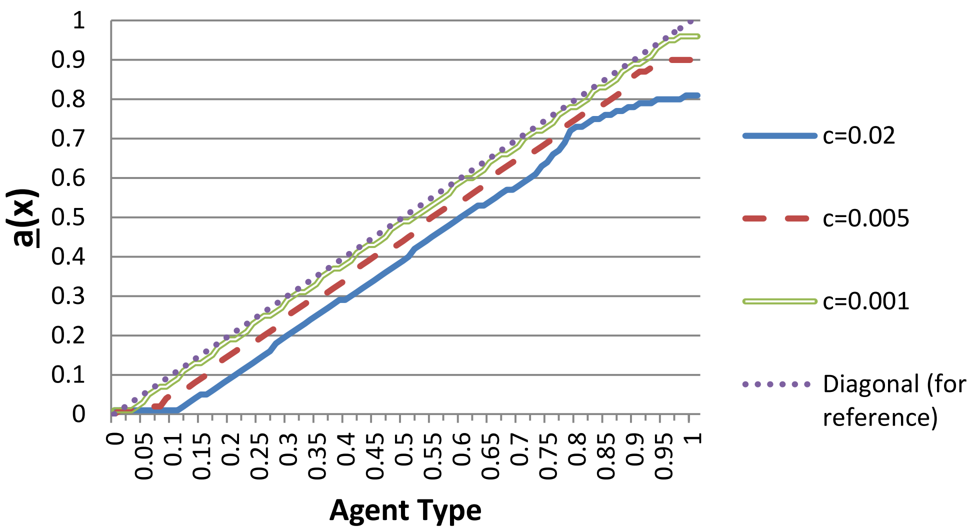

Figure 1: $\underline{a}(x)$ for $u(x, y) = \min\{x, y\}$

In section 7 we prove that in the minimum case (i.e. when $\alpha = 1$) the acceptance pattern is of the form of assortative matching, i.e., each agent type has a different threshold and the higher the type the higher the threshold. Assortative matching is common in distributed matching and occurs with many utility functions.

Figure 1 demonstrates the strategies used in equilibrium for the cases in which the utility function is the minimum of the two agent types, $f(x)$ is the uniform distribution and the search cost is 0.02, 0.005 or 0.001. The figure, as all other figures in this paper, was generated using a discrete evaluation process, as described in section 8. As can be seen in the figure, the equilibrium line is monotone increasing, that is, the higher the agent's type the higher the threshold. The figure also demonstrates that all agents set their threshold beneath their own type since the equilibrium line is always under the diagonal line. The distance between the two lines becomes larger as c increases.

6. Maximum as the Utility Function

We now consider the equilibrium in an environment where $\alpha = 0$; that is, the utility from forming a partnership between any two agents is the maximum value among the two, $u(x, y) = \max\{x, y\}$. Similar to when the utility function is the minimum, the optimal threshold of an agent of type x is given by:

$$\underline{a}(x,S) = \arg\max_t \frac{-c + \int_{y=t}^{x} x f(y)dy + \int_{\overline{M}(x,S)} y f(y)dy}{\int_{M(x,S)} f(y)dy}. \tag{4}$$

We prove that in the equilibrium solution in this case there exists a threshold x^* such that all types greater than x^* accept all types, and all types below or equal to x^* accept only agents of types greater than x^*.

Clearly, an agent of type 1 will receive a utility 1 from any partnership, therefore it will accept any partner ($\underline{a}(1) = 0$). Let $x^* = \inf\{x | \underline{a}(x) = 0\}$ (i.e., the agent with the lowest type that accepts all other agents).

Theorem 2. *For $\alpha = 0$ (i.e., $u(x,y) = \max\{x,y\}$), the following step function reflects a strategy profile that is in equilibrium:*

$$\underline{a}(x) = \begin{cases} 0 & \text{if } x^* \leq x \\ x^* & \text{otherwise.} \end{cases} \tag{5}$$

Proof. The expected utility of an agent of type x not setting its threshold at 0, must be greater than x (otherwise it would have accepted forming a partnership with agents of type 0). Therefore, by Proposition 2, it rejects any other agent with a type smaller than or even equal to x (which would give it a utility of x as well). This entails that partnerships can only be formed between agents which at least one of them has a threshold of 0.

Assume that $\underline{a}(x) > x$. Since no partnerships form between agents which both set their threshold above 0, any $x' \leq x$ could set its threshold at $\underline{a}(x)$ and yield an expected utility identical to $\phi(x)$, which is optimal according to Proposition 2 and is higher than what it would have obtained had it set its threshold to 0. This proves the first part of the step function, (i.e., agents of type x^* or greater set their threshold at 0).

We now prove that agents with type $x < x^*$ (where x^* satisfies the definition above) set their acceptance threshold at exactly x^*. By definition, an agent of type $x < x^*$ rejects partnerships with agents of type 0, therefore, no partnerships are formed between agents with types less than x^*. Consider $x < x^*$. Let t_x be the acceptance threshold of agents of type x. Agents of type $\geq x^*$ accept x. So, by Theorem 1, $t_x \leq y$ for all $y \geq x^*$. Hence $t_x \leq x^*$. Suppose $t_x < x^*$. However, since we know that no partnerships are formed between agents with types smaller than x^*, setting t_x to x^* will not modify the resulting partnerships, nor the expected utility of any agent. Therefore, for every $x < x^*$, $\phi(x) = u(x, x^*) = x^*$, and thus $\underline{a}(x) = x^*$.

□

Note that this structure has a very interesting property, namely, partnerships are formed if and only if one side has a value greater than or equal to $x*$.

Based on Proposition 2 the expected utility of agents of types smaller than x^* equals the utility of forming a partnership with an agent of type x^*, which equals x^*. Using this equilibrium structure (as S), Equation 4 obtains:

$$\Phi(x < x^*, x^*, S) = \frac{-c + \int_{x^*}^{1} y f(y) dy}{\int_{x^*}^{1} f(y) dy} = x^*, \tag{6}$$

which, after some mathematical manipulations, becomes:

$$c = \int_{x^*}^{1} (y - x^*) f(y) dy. \tag{7}$$

Clearly, the higher the search cost is, the lower x^* is and, similarly, $\lim_{c \to 0} x^* = 1$.

Figure 2 depicts the strategies used in equilibrium for the case in which the utility function is the maximum of the two types, $f(x)$ is the uniform distribution and the search cost is 0.02, 0.005 or 0.001. As can be seen in the figure, when $c = 0.02$, all agents with types $x > 0.8$ accept any other agents, while the agents with the lower types ($x \leq 0.8$) accept only agents with types greater than 0.8.

7. Mixed Maximum-Minimum as the Utility Function

We now turn to analyze the equilibrium in the more general case, in which the utility from the partnership is given by: $u(x, y) = \alpha \min\{x, y\} + (1 - \alpha) \max\{x, y\}$ with $0 < \alpha \leq 1$.

The maximization problem of an agent of type x in this case is:

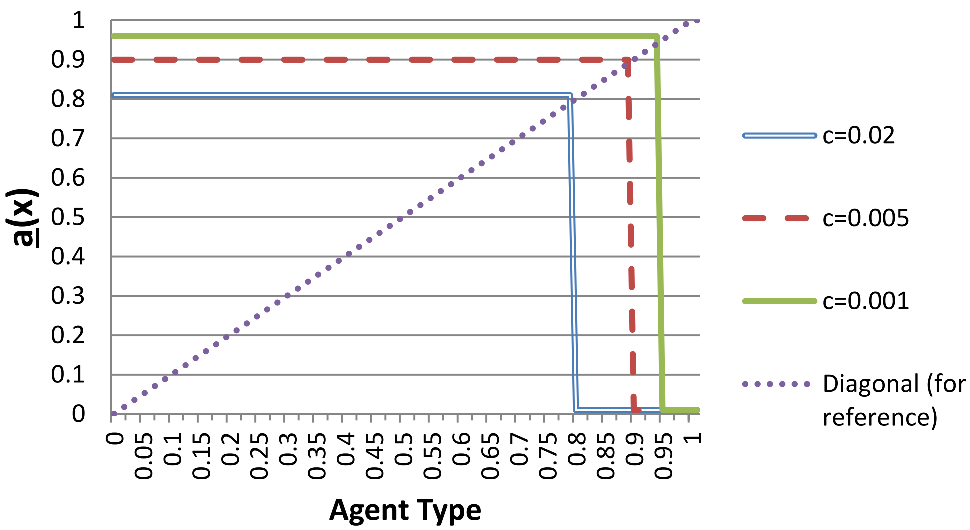

Figure 2: $\underline{a}(x)$ for $u(x, y) = \max\{x, y\}$

$$\arg\max_t \Big(-c + \alpha\big(\int_{y=t}^{x} yf(y)dy + (\int_{\overline{M}(x,S} f(y)dy \cdot x)\big) + (1-\alpha)\big((\int_{\underline{M}(x,S)} f(y)dy) \cdot x + \int_{\overline{M}(x,S)} yf(y)dy\big)\Big) \cdot \frac{1}{\int_{M(x,S)} f(y)dy}. \quad (8)$$

We now turn to prove a key theorem which, along with its corollaries, allows us to determine the equilibrium in this mixed case.

The following theorem states that, given a strategy profile S *in equilibrium*, when $0 < \alpha < 0.5$ (more weight is given to the maximum), if two agents are accepted by the same group, then the agent of the lower type will set its threshold above the one used by the higher type. The opposite happens when $\alpha > 0.5$ (more weight to the minimum). When $\alpha = 0.5$ the two agents will set the same threshold.

Theorem 3. *Given a strategy profile S* in equilibrium*, for any x, x', such that $x' < x$, $A(x, S) = A(x', S)$ and $0 < \underline{a}(x, S) \leq x'$:*

- *If $0 < \alpha < 0.5$ then $\underline{a}(x, S) < \underline{a}(x', S)$.*
- *If $0.5 < \alpha \leq 1$ then $\underline{a}(x, S) > \underline{a}(x', S)$.*

- *If* $0.5 = \alpha$ *then* $\underline{a}(x,S) = \underline{a}(x',S)$.

Proof. Consider the case where $0 < \alpha < 0.5$.

Since $0 < \underline{a}(x,S)$ then (from Proposition 2):

$$\phi(x,S) = u(x, \underline{a}(x)), \tag{9}$$

which, given $\underline{a}(x,S) < x$, implies:

$$\phi(x,S) = \alpha \cdot \underline{a}(x) + (1-\alpha)\cdot x = \alpha \cdot \underline{a}(x) + (1-\alpha)\cdot x' + (1-\alpha)\cdot(x-x'). \tag{10}$$

On the other hand, calculating $\phi(x)$ explicitly, obtains:

$$\begin{aligned}\phi(x,S) = \Big(-c + \int_{\underline{a}(x)}^{x'} (\alpha y + (1-\alpha)x) f(y)dy + \\ \int_{x'}^{x} (\alpha y + (1-\alpha)x) f(y)dy + \int_{\overline{A}(x,S)} (\alpha x + (1-\alpha)y) f(y)dy\Big)\cdot \\ \frac{1}{\int_{M(x,S)} f(y)dy}.\end{aligned} \tag{11}$$

Recall that $A(x) = A(x')$, therefore, if an agent of type x' sets its threshold at $\underline{a}(x,S)$, both agents will have the same matching set (implying that $\underline{a}(x',S) = [\underline{a}(x), x')$), therefore:

$$\begin{aligned}\Phi(x', \underline{a}(x), S) = \Big(-c + \int_{\underline{a}(x)}^{x'} (\alpha y + (1-\alpha)x') f(y)dy + \\ \int_{x'}^{x} (\alpha x' + (1-\alpha)y) f(y)dy + \int_{\overline{A}(x,S)} (\alpha x' + (1-\alpha)y) f(y)dy\Big)\cdot \\ \frac{1}{\int_{M(x,S)} f(y)dy}.\end{aligned} \tag{12}$$

Subtracting Equation 12 from Equation 11 obtains:

$$\phi(x,S) - \Phi(x',\underline{a}(x),S) = \Big(\int_{\underline{a}(x)}^{x'} f(y)dy \cdot (1-\alpha)(x-x')$$
$$+ \int_{x'}^{x} (\alpha y + (1-\alpha)x) f(y)dy - \int_{x'}^{x} (\alpha x' + (1-\alpha)y) f(y)dy +$$
$$\int_{\overline{A}(x,S)} f(y)dy\alpha(x-x')\Big) \cdot (\frac{1}{\int_{M(x,S)} f(y)dy}). \quad (13)$$

Denote:

$$\eta = \Big(\int_{\underline{a}(x)}^{x'} f(y)dy \cdot (1-\alpha)(x-x') +$$
$$\int_{x'}^{x} (((1-\alpha)x - \alpha x') - (1-2\alpha)y) f(y)dy +$$
$$\int_{\overline{A}(x,S)} f(y)dy\alpha(x-x')\Big) \cdot \frac{1}{\int_{M(x,S)} f(y)dy}.$$

From Equation 13 we obtain:

$$\Phi(x',\underline{a}(x),S) = \phi(x,S) - \eta. \quad (14)$$

Replacing $\phi(x,S)$ with $\alpha \cdot \underline{a}(x) + (1-\alpha) \cdot x' + (1-\alpha) \cdot (x-x')$ (Equation 10) obtains:

$$\Phi(x',\underline{a}(x),S) = \alpha \cdot \underline{a}(x) + (1-\alpha) \cdot x' + (1-\alpha) \cdot (x-x') - \eta. \quad (15)$$

We now show that η is smaller than $(1-\alpha) \cdot (x-x')$.

$\alpha < 0.5$ implies that $1 - 2\alpha > 0$, therefore:

$$\int_{x'}^{x} (((1-\alpha)x - \alpha x') - (1-2\alpha)\mathbf{y}) f(y)dy <$$
$$\int_{x'}^{x} (((1-\alpha)x - \alpha x') - (1-2\alpha)\mathbf{x}') f(y)dy = \int_{x'}^{x} f(y)dy(1-\alpha)(x-x'). \quad (16)$$

Clearly:

$$\int_{\overline{A}(x,S)} f(y)dy\alpha(x-x') < \int_{\overline{A}(x,S)} f(y)dy(1-\alpha)(x-x'). \quad (17)$$

Since $M(x,S) = [\underline{a}(x), x'] \cup [x', x] \cup \overline{A}(x,S)$, we obtain:

$$(1-\alpha)\cdot(x-x') > \eta. \tag{18}$$

Putting together Equations 15 and 18 we obtain:

$$\alpha \cdot \underline{a}(x) + (1-\alpha)\cdot x' < \Phi(x', \underline{a}(x), S). \tag{19}$$

However $\alpha \cdot \underline{a}(x) + (1-\alpha)\cdot x' = u(\underline{a}(x), x')$ and $\Phi(x', \underline{a}(x), S) \leq \phi(x', S)$ (because $\phi(x', S)$ is the expected utility when using an optimal threshold), therefore $u(\underline{a}(x), x') < \phi(x, S)$ and therefore (by definition of $\underline{a}$ and monotonicity and continuousness of u) $\underline{a}(x) < \underline{a}(x')$. [3]

Using the exact same proof (except for changes in the inequality starting at Equation 16), we obtain that if $\alpha > 0.5$ then $\underline{a}(x) > \underline{a}(x')$. We also obtain that for $\alpha = 0.5$, $\underline{a}(x) = \underline{a}(x')$. $\square$

Theorem 3 provides a logical division of the equilibrium structure according to the three different cases, which will be further discussed in the following subsections:

- $0.5 < \alpha \leq 1$: where the lower type among the two agents in a match has a greater impact, in which higher types set their threshold higher (assuming the same group of agents accepting them).
- $\alpha = 0.5$: where both types have equal effect on the agents' utility, in which agents with the same accepting group set their threshold identically.
- $0 < \alpha < 0.5$: where the higher type among the two agents in a match has a

[3] The same proof can also be used when $x' < \underline{a}(x) < x$ (rather than $\underline{a}(x) \leq x'$) by replacing the term

$$\int_{\underline{a}(x)}^{x'} f(y)dy \cdot (1-\alpha)(x-x') + \int_{x'}^{x} (\alpha y + (1-\alpha)x) f(y)dy - \int_{x'}^{x} (\alpha x' + (1-\alpha)y) f(y)dy,$$

by

$$\int_{\underline{a}(x)}^{x} (\alpha y + (1-\alpha)x) f(y)dy - \int_{\underline{a}(x)}^{x} (\alpha x' + (1-\alpha)y) f(y)dy,$$

in Equation 13 and its sequels.

greater impact, in which higher agent types set their threshold *lower* (assuming the same group of agents accepting them).

The latter case, in which higher agent types set their threshold *lower*, is very interesting and the equilibrium behavior may seem counter-intuitive. The following discussion may provide some intuition behind such behavior. When $\alpha = 0$, as shown above, agents with types above some x^* set their threshold at 0, as it is not worthwhile for them to resume the search and pay the search cost again. Suppose we slightly increase α. In this case, an agent of type 1 will still accept any other agent, as the maximum difference in utility (i.e., the difference between forming a partnership with an agent of type 0 and another agent of type 1) equals α, which may still be less than the search cost. However, an agent of a lower type, x, (but still above x^*), will have a greater difference between forming a partnership with an agent of type 0 and forming a partnership with an agent of type 1 (this difference equals: $\alpha \cdot x + (1-\alpha) \cdot 1 - (1-\alpha) \cdot x = \alpha + (1-x)(1-2\alpha) > \alpha$). In fact, an agent of type x has two different ways of expecting to increase its utility by resuming the search. The first is simply by finding a different agent with a higher type (though still lower than itself). If this happens, the agent will increase its utility by the difference between these two agent types multiplied by α. This opportunity for increasing its utility is shared by all agents with high types. However, the second way an agent of type x can expect to benefit from resuming the search is by finding an agent with a type higher than its own. If this happens the agent's utility will be increased by the difference between its own type and the new agent's type multiplied by $(1-\alpha)$ which is large. The lower the agent's type the higher its expectation from this second form of revenue. Therefore, it is more likely to resume the search if it interacts with agents of lower types, which implies having a higher threshold.

We further illustrate this behavior using the following toy example. Assume three agent types $0, \frac{1}{2}, 1$, and a uniform distribution over them. Assume $\alpha = 0.2$ and $c = 0.08$. Table 2 depicts the utilities each of the agent types would receive when forming a partnership with each of the other agent types, along with the agent types each agent would accept in equilibrium and each agent type's expected utility. It is most interesting

Agent Type	Utility of partnership if formed with			agent-type accepting	expected utility
	1	0.5	0		
1	1	0.9	0.8	{0.5,1}	0.83
0.5	0.9	0.5	0.4	1	0.66
0	0.8	0.4	0	{0,0.5,1}	-0.24

Table 2: Toy example for $\alpha = 0.2$ and $c = 0.08$, with agents of types in $\{0, 0.5, 1\}$ only.

to note that agents of type 0.5 accept only agents of type 1 and not agents of type 0.5 while agents of type 1 accept both agents of type 1 and of type 0.5. If agents of type 0.5 were to accept both agents of type 1 and of type 0.5, their expected utility would be reduced from 0.66 to 0.58.

7.1. $0.5 < \alpha \leq 1$ *(Greater Minimum Impact)*

When the minimum type has a greater impact, we obtain assortative matching, where the higher the agent's type the higher its threshold, and all agents set their threshold beneath their own type. See Figure 3 for an example where $\alpha = 0.8$, uniform $f(x)$ and $c = 0.02$, $c = 0.005$ and $c = 0.001$. As discussed earlier in the paper, assortative matching is common in distributed matching. The Section proves assortative matching for $0.5 < \alpha \leq 1$.

The following Lemma shows that the larger the set of agents accepting an agent of a certain type, the higher the threshold that the agent will set. This lemma is true for any monotonic u.

Lemma 3. *Given two strategy profiles S and S', with $A(x, S') \subset A(x, S)$, the optimal strategy for an agent of type x satisfies: $\underline{a}(x, S') \leq \underline{a}(x, S)$.*

Proof. An agent of type x can reject any agent which is not in $A(x, S')$, therefore $\phi(x, S') \leq \phi(x, S)$, and consequently, based on Proposition 2, $u(x, \underline{a}(x, S')) \leq u(x, \underline{a}(x, S))$. Now, due to the monotonicity of u, $\underline{a}(x, S') \leq \underline{a}(x, S)$. $\square$

The intuitive interpretation of Lemma 3 is that the larger the set of agents accepting an agent, the more "picky" that agent can be, and consequently the higher its reservation value.

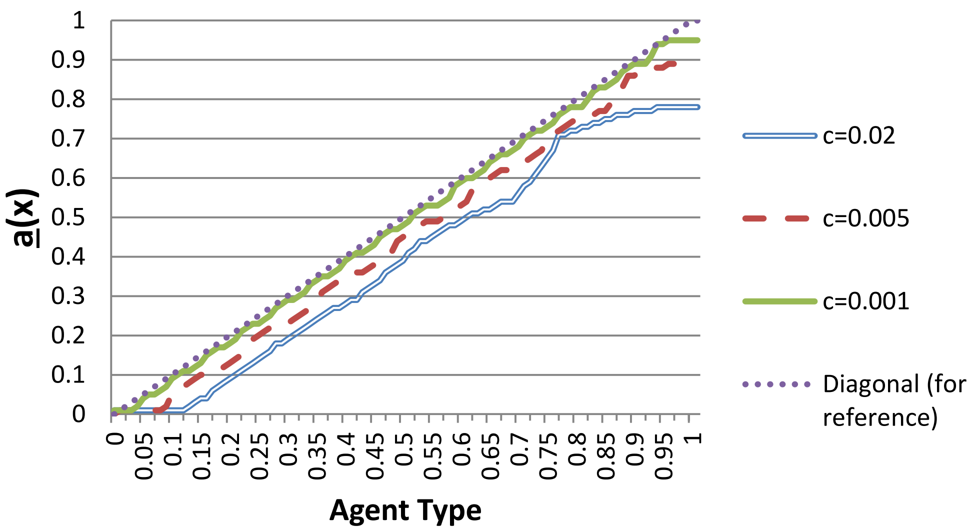

Figure 3: Simulation results for $\alpha = 0.8$

Theorem 4. *When* $0.5 < \alpha \leq 1$*, in equilibrium, we obtain assortative matching, i.e. the higher the agent's type, the higher its threshold, and all agents set their threshold beneath their own type.*

Proof. From Theorem 3 and Lemma 3 we obtain that if $\alpha > 0.5$, for every $x' < x$, if $0 < \underline{a}(x) < x$ then $\underline{a}(x') < \underline{a}(x)$, i.e., $\underline{a}(x)$ is monotonously increasing in x.

We now show that every agent sets its threshold beneath its own type, i.e, If $\alpha > 0.5$, then $\underline{a}(x) < x$ for every x. Clearly, for every x, $\int_{M(x,S)} f(x)dx > 0$ (otherwise the search will go on forever). Assume by contradiction that there exists $\tilde{x}$ such that $\underline{a}(\tilde{x}) \geq \tilde{x}$ and let x' be the maximum type such that $\underline{a}(x') \geq x'$. For every $x > x'$ $0 < \underline{a}(x) < x$, which according to the monotonicity of the thresholds which was proved above, implies that $\underline{a}(x') < \underline{a}(x)$, implying $x' < \underline{a}(x)$. Therefore, every agent of type $x > x'$ rejects an agent of type x', but x' rejects any agent of types smaller than its own, implying that $\int_{M(x,S)} f(x)dx = 0$, contradicting the above. $\square$

7.2. $\alpha = 0.5$ *(Equal Minimum and Maximum Impact)*

In the case of $\alpha = 0.5$, the utility is simply the average between the two agents forming a partnership. This configuration is described in [8] where it is shown that, in

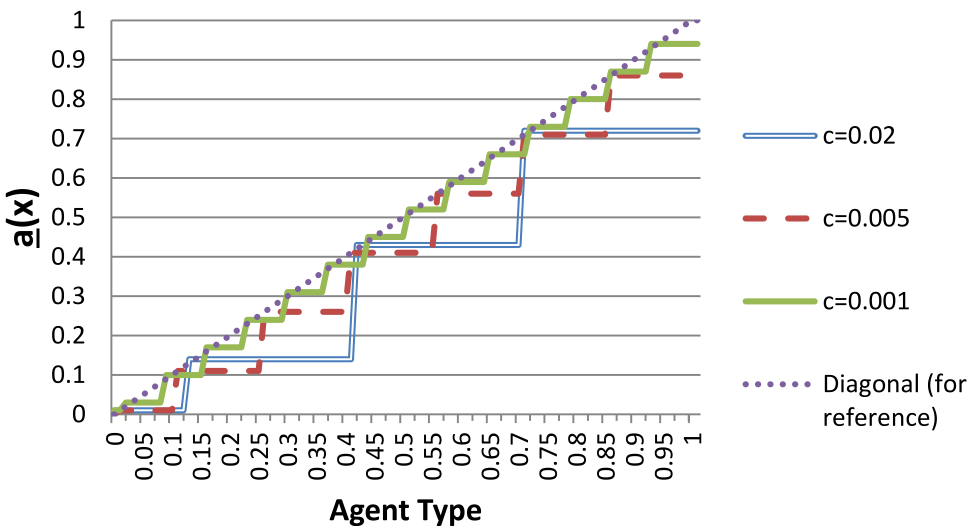

Figure 4: Simulation results for $\alpha = 0.5$

equilibrium, the entire domain is partitioned and partnerships are formed only within these partitions (a structure termed "perfect segregation"). The larger the search cost the less partitions are created. See Figure 4 for an example of the case where $\alpha = 0.5$, uniform $f(x)$ and the search cost is 0.02, 0.005 or 0.001. As can be seen in the figure, the entire domain is divided into a number of partitions in which agents form partnerships only with agents in their partition (see for example the segment [0.43,0.7] in Figure 4, when $c = 0.02$—all agents with types in that segment set their threshold to 0.43). The higher the cost, the less partitions there are. As illustrated in the figure, when $c = 0.001$ there are 15 such partitions, when $c = 0.005$ there are 7 such partitions, and when $c = 0.02$ there are only 4 such partitions.

7.3. $0 < \alpha < 0.5$ *(Greater Maximum Impact)*

The case in which $0 < \alpha < 0.5$ is both the most challenging and interesting one. In this section we will show that if $0 < \alpha < 0.5$, when following the threshold as a function of the agent type from right to left (starting at $x = 1$):

1. The threshold increases until there is an agent that sets its threshold at its own type (i.e. until the diagonal is reached).

2. In any case where an agent sets its threshold above its own type (any point above the diagonal) the threshold does *not* increase (as the type decreases).
3. If c is small enough but not zero, the threshold as a function of the agent type is non-monotonic (for some f).

To the best of our knowledge, this is the only equilibrium described in the literature where, despite the utility function being monotonic, there exist types which set their threshold above their own type, and that the threshold as a function of the agent type has segments which increase as the type decreases. The following theorem proves these three properties.

Theorem 5. *Let* $\hat{x} = \sup\{x | x \leq \underline{a}(x)\}$. *When* $0 < \alpha < 0.5$ *the following hold:*

1. *For any* x, x' *such that* $\hat{x} < x' < x$ *we get that* $\underline{a}(x) < \underline{a}(x')$.
2. *For any* x, x' *such that* $x' < x$ *and* $x < \underline{a}(x)$ *we get that* $\underline{a}(x') \leq \underline{a}(x)$.
3. *For some* c, f *there exist* $x'' < x' < x$ *such that* $\underline{a}(x) < \underline{a}(x')$ *and* $\underline{a}(x'') < \underline{a}(x')$.

Proof. The first property is an immediate corollary of Theorem 3, since the group of agents with the higher types who set their threshold beneath their own type are all accepted by all of the other agents (i.e. $A(\cdot) = [0, 1]$). Therefore the smaller the agent's type, the higher its threshold (Theorem 3).

We now address the second property. The following lemma claims that in equilibrium, when $\alpha < 0.5$, if an agent sets its threshold above its own type, then any agent with a value beneath its type will use the exact same threshold, if it is accepted by the same group of agents.

Lemma 4. *Given an equilibrium profile* S, x *and* $x' < x$ *if* $x \leq \underline{a}(x, S)$ *and* $A(x, S) = A(x', S)$ *then* $\underline{a}(x', S) = \underline{a}(x, S)$.

Proof. Since $x \leq \underline{a}(x, S)$:

$$\phi(x, S) = \alpha \cdot x + (1 - \alpha) \cdot \underline{a}(x, S) = \alpha \cdot x' + \alpha \cdot (x - x') + (1 - \alpha)\underline{a}(x, S) \quad (20)$$

On the other hand:

$$\phi(x,S) = \Phi(x',\underline{a}(x,S),S) + \frac{\int_{M(x,S)} f(y)dy\alpha(x-x')}{\int_{M(x,S)} f(y)dy} = \Phi(x',\underline{a}(x,S),S) + \alpha(x-x') \tag{21}$$

This implies:

$$\Phi(x',\underline{a}(x,S),S) = \alpha \cdot x' + (1-\alpha) \cdot \underline{a}(x,S) = u(x',\underline{a}(x,S)) \tag{22}$$

and therefore $\underline{a}(x',S) = \underline{a}(x,S)$. □

Lemma 4 enables proving the second property, i.e. if an agent sets its threshold above its type, then any agent with a value beneath its type will not set its threshold above the former agent's threshold.

Lemma 5. *Given x and $x' < x$, if $x \leq \underline{a}(x)$ then $\underline{a}(x') \leq \underline{a}(x)$.*

Proof. From Lemma 1, $A(x',S) \subset A(x,S)$, hence the proof is immediate using Lemmas 4 and 3. □

For the third property we will first show that there exist c and f such that $0 < \underline{a}(1) < 1$. For f we assume uniform distribution (though a similar proof can be used for any distribution that gives some positive weight to any interval of types).

If $c > 0$ we get that $\underline{a}(1) < 1$, otherwise, $\phi(1) = -\infty$. Let $c = 0.01 \cdot \alpha$, an agent of type 1, will reject forming a partnership with an agent of type 0.1 (and receiving only $0.1 \cdot \alpha + (1-\alpha)$), as it has an 50% chance of interacting with an agent of type 0.5 or above next round and receiving $0.5 \cdot \alpha + (1-\alpha) - 0.1 \cdot \alpha$, and if setting its threshold above 0.1 it would be required to reject forming a partnership only 0.1 times on average.

Set $x = 1, x'' = 0$ and $x' = \hat{x}$. By property 1, we get that $\underline{a}(x) < \underline{a}(x')$. Also from property 1, we know that for all agents y with types in $(\hat{x},1)$, $\underline{a}(y) > \underline{a}(1)$, therefore, $(\hat{x},1) \notin M(0)$. Hence, $\underline{a}(0) < \hat{x}$, otherwise $\phi(0) = -\infty$. Recall that (by definition) $\hat{x} \leq \underline{a}(\hat{x})$, therefore $\underline{a}(x'') < \underline{a}(x')$.

□

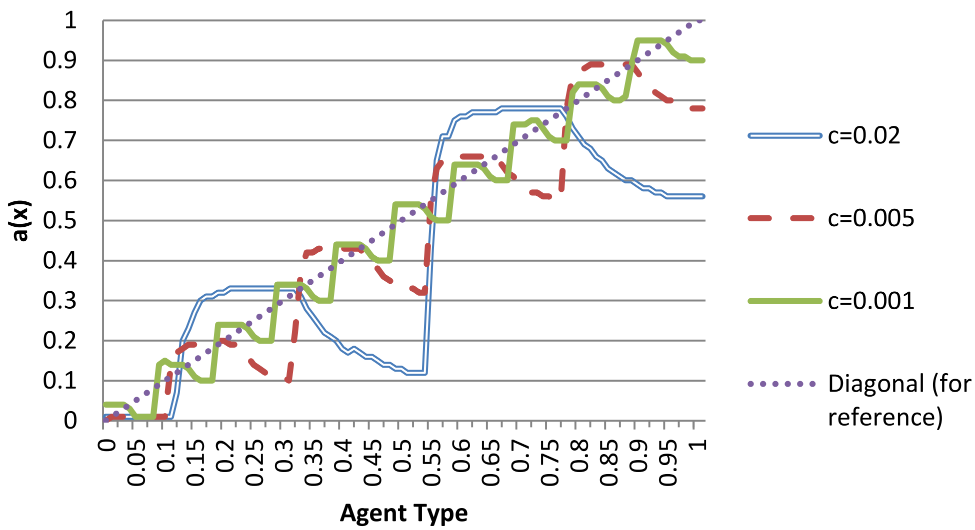

Figure 5: Simulation results for $\alpha = 0.2$

Figure 5 gives some examples for $\alpha = 0.2$, uniform $f(x)$ and when the cost is 0.02, 0.005 or 0.001. For $c = 0.02$, note the segment $(0.77, 1)$ in which $\underline{a}(x)$ is monotone decreasing as the type increases.

Before concluding this section, we highlight an interesting property reflected from Figure 5. In several cases there is a sudden drop in the threshold, resulting from a small decrease in the type. When $c = 0.02$ this happens next to 0.54 (when $c = 0.005$ this happens next to $0.77, 0.54, 0.31$ and 0.1). After the sudden drop, the threshold resumes climbing until once again there is an agent that sets its threshold at its own type. We provide a possible explanation for this phenomenon based on the example analyzed in Figure 5 when the cost is 0.02. Consider for instance an agent with a type of 0.54. Note that the threshold set by the agents with a type of 1 is slightly above 0.54. Therefore, almost all agents with types greater than 0.54 reject partnerships with the agent. Therefore its expected utility drops drastically and so does its threshold. Since almost no agents with types above 0.54 accept it as a partner, all agents with types slightly under 0.54 are accepted by nearly the same set of agents ($A(\cdot) = [0, 0.54]$). Therefore according to Theorem 3 the threshold resumes climbing until once again there is an agent who sets its threshold at its own type.

8. Discrete Case

The analysis presented in the former sections completely unfolds the structure of the equilibrium in this new class of distributed matching. In addition, it facilitates the calculation of the equilibrium through discretization of types (or whenever the types are a priori inherently discretized). In this section we show how, based on the above analysis, the equilibrium thresholds of the different agents types can be calculated through dynamic programing, using Algorithm 1.

Let X be a finite set of possible agent types. $f(x)$ is the fraction of the agents of type x. In this case, the expected outcome for agent x can be calculated using the following equation (a discretization of Equation 8):

$$\phi(x,S) = \Big(-c + \alpha\big(\sum_{y=\underline{a}(x)}^{x} yf(y) + (\sum_{y \in \overline{M}(x,S)} f(y) \cdot x)\big) + (1-\alpha)\big((\sum_{y=\underline{a}(x)}^{x} f(y)) \cdot x + \sum_{y \in \overline{M}(x,S)} yf(y)\big)\Big) \cdot \frac{1}{\sum_{y \in M(x,S)} f(y)} \quad (23)$$

Algorithm 1 Finding the threshold for all agents in the discrete case

Input: A finite set of all agent types X, α, search cost c and a density function $f(x)$.
Output: An array which holds $\underline{a}(x)$, for every x.
Initialize array $\underline{a}$ of size $|X|$ with 0's
for each $x \in X$, in descending order **do**
Set $A(x) = \{y|\underline{a}(y) < x\}$
Set $\underline{a}(x) = \arg\max_y\{\Phi(x,y,A(x))\}$ {calculate $\Phi(x,y,A(x))$ for every y using Equation 23 and choose the maximizer}
Tie break using $\underline{a}[x] = \arg\min_y\{u(x,y) - \phi(x)\}$
end for
return $\underline{a}$.

The algorithm performs a single pass on all types. It finds the threshold for every type based on the thresholds calculated for the greater types. Note that the order in which the threshold array ($\underline{a}$) is filled is crucial and must start with the agent with the highest type going downwards, since every agent's threshold depends on the thresholds of all agents with greater types. Based on Theorem 1, the agent's threshold does *not* depend on agents of lower types, which implies the correctness of the algorithm.

9. Mixed Maximum Minimum Utility Based Mechanism Design

The mixed maximum-minimum utility function is not only very common in real-life situations, but may also be artificially imposed by a mechanism designer in order to influence the behavior of a population in a way that better aligns with a desired one. In this section we demonstrate the benefit which may be obtained by a mechanism designer if she were to use a Mixed Maximum-Minimum Utility function (or more specifically a maximum utility function, i.e. $\alpha = 0$).

Assume a professor would like to provide an assignment to her class which must be completed in groups of two and presented in class. As part of the presentation, each student is asked questions, and the assignment's grade is set based on the knowledge reflected by the two. It is assumed that a student's knowledge regarding the assignment is fully correlated with her type (i.e., her "competence"), hence stronger students will generally exhibit greater knowledge and will perform better individually. While students aim to maximize their utility, which is a linear tradeoff of the time spent searching for a partner and the grade they eventually receive, the professor may be interested in maximizing a different measure. For example, the professor may be interested in having the weaker students matched with the stronger ones to encourage knowledge transfer between the two. In such case, setting the students' grade based on a maximum-minimum function of the type analyzed in this paper, rather than using the traditional functions (e.g., grading each student according to the knowledge exhibited individually or based on the average of the individual grades) may result in better performance from the professor's point of view.

Formally, we denote the assignment grade to student of type x if pairing with a student of type y using $u(x, y)$. The students are assumed to use distributed matching to find their partner for the assignment, taking into consideration: (a) the function $u(x, y)$ set by the professor; (b) the search cost c (e.g., the cost of time communicating with other students and reasoning about their type/competence), which is assumed to be additive to the grade parameter; and (c) the distribution of types in the population of students, which in our case is taken to be uniform ($[0, 1]$). The performance measure, from the professor's point of view, as discussed above, denoted $\ell(x, y)$, is taken to be linearly proportionate to the expected distance between the weaker student and the stronger student in the pair, i.e. $\ell(x,y) = E[K\cdot|x-y|]$. In the following paragraphs we evaluate three common grading functions the professor may use, demonstrating that for a large portion of the possible settings a function of the type analyzed in this paper is the dominating one. The three functions are: (a) $u(x,y) = x$ (each student obtains her own grade); (b) $u(x,y) = u(y,x) = \frac{x+y}{2}$ (each student obtains the average grade); and (c) $u(x,y) = u(y,x) = \max\{x,y\}$. Obviously more sophisticated functions could have been used, e.g., ones that correlate grade with the distance between the types of the two students, however these do not necessarily guarantee the desired intuitive property according to which a partnership between any two types $x' > x$ and $y' > y$ implies $u(x',y') > u(x,y)$, i.e., if both students in a pair are better than the students in another pair then the first should get a higher grade.

We begin by analyzing $u(x,y) = x$. In this case, a student's grade does not depend on the grade of his partner, therefore every student will accept any other student and the search will terminate after a single search iteration. The expected performance, from the professor's point of view, in this case is:

$$K \cdot \int_0^1 \Big(\int_0^t (t-t')dt' + \int_t^1 (t'-t)dt' \Big) dt = K \cdot (\frac{1}{2}\cdot t - \frac{t^2}{2} + \frac{t^3}{3})\Big|_0^1 = K \cdot \frac{1}{3} \tag{24}$$

Next we analyze the case where $u(x,y) = \frac{x+y}{2}$. This is the function analyzed by Chade [8], in which perfect segregation is observed. When the density function is uniform and $c < \frac{1}{2}$, the size of each cluster is $\sqrt{2c}$ (where the last cluster is possibly of a smaller size). Therefore, by computing the expected performance from the professor's point of view, in a manner similar to Equation 24, we obtain an upper bound of $K \cdot \frac{1}{3} \cdot \sqrt{2c}$ (as the last cluster is equal to or smaller than $\sqrt{2c}$).

Finally we analyze the case where $u(x,y) = \max\{x,y\}$.[4] Recall, that in this case, there exists some $0 < x^* < 1$, and all pairs contain at least one member with

[4]The maximum function is an instance of the min-max function when $\alpha = 0$. Clearly when $\alpha > 0.5$, the expected performance will be very low since, as we have shown, $\alpha > 0.5$ implies assortative matching and thus the partnerships formed will be very close to each other. Intuitively it seems that using $\alpha = 0$ yields the best performance and we use it since $\alpha = 0$ may be solved analytically.

type greater than x^*. The smaller the search cost c, the larger x^* (see section 6). The expected performance, from the professor's point of view, in this case is:

$$
\begin{aligned}
&K \cdot \frac{1}{1-x^*} \cdot \int_{x^*}^{1} \Big(\int_0^t (t-t')dt' + \int_t^1 (t'-t)dt' \Big) dt \\
=&K \cdot \frac{1}{1-x^*} \cdot (\frac{1}{2} \cdot t - \frac{t^2}{2} + \frac{t^3}{3})\Big|_{x^*}^{1} = K \cdot \frac{1}{1-x^*} \cdot (\frac{1}{3} - \frac{1}{2}x^* + \frac{x^{*2}}{2} - \frac{x^{*3}}{3})
\end{aligned}
\tag{25}
$$

When $c = \frac{3}{8}$, $x^* = \frac{1}{2}$ and $\frac{1}{1-x^*} \cdot (\frac{1}{3} - \frac{1}{2}x^* + \frac{x^{*2}}{2} - \frac{x^{*3}}{3}) = \frac{1}{3}$. Equation 25 is monotonically increasing in x^* in $(0.5, 1)$. Therefore, as long as the search cost is smaller than $\frac{3}{8}$ (which is a very high search cost), the expected performance from the professor's point of view using $u(x, y) = \max\{x, y\}$ is greater than when using $u(x, y) = x$ and $u(x, y) = \frac{x+y}{2}$. The lower the search cost, the higher x^* and the better the expected performance with $u(x, y) = \max\{x, y\}$.

10. Conclusions and Future Work

The settings in which the output function for a pair forming a partnership is a linear combination of the maximum and the minimum among the partners types are highly applicable in real life. The analysis of the model introduces the unique properties of the equilibrium in the setting considered. In particular it includes a proof that an agent that sets its threshold below its own type is necessarily accepted by all agents with types between that threshold and the agent's type (Theorem 1). These properties substantially simplify the calculation of the equilibrium strategies. While the structure of the agents' strategies in equilibrium for the case of $\alpha \geq 0.5$ resembles those obtained by traditional utility functions, the equilibrium for the case of $\alpha < 0.5$ is substantially different from related work. The acceptance threshold as a function of the agent type increases and decreases alternately and, although the utility function is monotonic, some agents set their threshold above their own type. We present an example in which the min-max utility can actually be used as a mechanism design in order to imply certain behavior on the partnerships formed.

In this work we assumed that when a partnership is formed, the two agents leave the environment and are replaced with two new agents of the same type, and all the proofs in this paper rely on this assumption. Despite not being realistic, this "replenishment assumption" appears in most literature that deals with two-sided search (e.g. [26, 8, 25, 5]), as it allows the type distribution to remain constant and not change over time. If we relax this assumption, and assume that partnerships are not replaced at all[5], our analysis capabilities become very limited, as $f(x)$ becomes $f_t(x)$ and there is no convergence guarantee. However, while the general solution is too difficult to analyze, it does result with a simple solution for the case in which $\alpha = 0$ (maximum). In this case, clearly, any agent with a type greater than $1 - c$, will accept any partnership at the first round (since the agent cannot reach a better result next round), and so will their partners. Therefore, in the second round, there should remain no agents with types greater than $1 - c$. Knowing this, any agent with type greater than $1 - 2c$ must also accept forming a partnership at the first round, and so must their partners. By repeating this reasoning, we conclude that all agents must accept their first partner, regardless of their types. In [8] Chade analyzes the case of $\alpha = 0.5$ (along with its equivalents). Chade argues, that when the "replenishment assumption" is relaxed, while multiple equilibria arise, they all should exhibit the perfect segregation property.

We see great importance in future research that will combine bargaining as part of the interaction process. We believe such research can result in many rich variants of our distributed matching model. Another interesting variant of our model, which will bring it closer to reality, may assume that the agents have some prior knowledge on other agents and that the pairing and interaction between the agents is not completely random. For this variant the search may be composed of two phases, in the first step (which may be analogous to going to a bar) the agents are randomly paired with a pool of k other agents, and obtain some noisy observation of their types (see [9] for a noisy observation model). In the second phase (which may be analogous to offering a drink), the agents may decide whether they would like to approach any of these agents in order to reveal their true type (with an additional search cost paid solely by the approaching

[5] A different approach to relaxing this assumption was taken by [6] in which, the authors assume that there is a steady stream of incoming agents with a specific distribution over their types which is a-priori known to all agents. However, a steady state requires that the flow of incoming agents equals the flow of partnerships formed.

agent), or resume the search.

11. Acknowledgments

Preliminary results of this work appeared in Proceedings of the 2012 IEEE/WIC/ACM International Conference on Intelligent Agent Technology. This work was partially supported by ISF grant 1083/13 and ISF-NSFC joint research program 2240/15.

References

[1] S. Alkoby, D. Sarne, S. Das, Strategic free information disclosure for search-based information platforms, in: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS'15), 2015.

[2] E. Anshelevich, S. Das, Y. Naamad, Anarchy, stability, and utopia: creating better matchings, Autonomous Agents and Multi-Agent Systems 26 (1) (2013) 120–140.

[3] A. Atakan, Assortative matching with explicit search costs, Econometrica 74 (3) (2006) 667–680.

[4] G. S. Becker, A theory of marriage: Part i, Journal of Political Economy 81 (4) (1973) 813–846.

[5] F. Bloch, H. Ryder, Two-sided search, marriages, and matchmakers, International Economic Review 41 (2000) 93–115.

[6] K. Burdett, M. G. Coles, Marriage and class, The Quarterly Journal of Economics 112 (1) (1997) 141–68.

[7] K. Burdett, R. Wright, Two-sided search with nontransferable utility, Review of Economic Dynamics 1 (1998) 220–245.

[8] H. Chade, Two-sided search and perfect segregation with fixed search costs, Mathematical Social Sciences 42 (1) (2001) 31–51.

[9] H. Chade, Matching with noise and the acceptance curse, Journal of Economic Theory 129 (1) (2006) 81–113.

[10] H. Chade, L. Smith, Simultaneous search, Econometrica 74 (5) (2006) 1293–1307.

[11] M. Chhabra, S. Das, D. Sarne, Competitive information provision in sequential search markets, in: International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14, Paris, France, May 5-9, 2014, 2014.

[12] M. Chhabra, S. Das, D. Sarne, Expert-mediated sequential search, European Journal of Operational Research 234 (3) (2014) 861–873.

[13] S. Das, E. Kamenica, Two-sided bandits and the dating market, in: Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05), 2005.

[14] K. Decker, K. Sycara, M. Williamson, Middle-agents for the internet, in: Proc. of IJCAI, 1997.

[15] T. S. Ferguson, Who solved the secretary problem?, Statistical science (1989) 282–289.

[16] D. Gale, L. Shapley, College admissions and the stability of marriage, American Math. Monthly 69 (1962) 9–15.

[17] N. Hazon, Y. Aumann, S. Kraus, D. Sarne, Physical search problems with probabilistic knowledge, Artificial Intelligence 196 (2013) 26–52.

[18] J. Huang, A. Kazeykina, Optimal strategies for reviewing search results, in: National Conference on Artificial Intelligence, 2010.

[19] M. Klusch, Agent-mediated trading: Intelligent agents and e-business, J. on Data and Knowledge Engineering 36 (3) (2001) 59–76.

[20] M. Landsberger, D. Peled, Duration of offers, price structure, and the gain from search, Journal of Economic Theory 16 (1) (1977) 17–37.

[21] S. Lippman, J. McCall, The economics of job search: A survey, Economic Inquiry 14 (1976) 155–189.

[22] X. Lu, R. P. McAfee, Matching and expectations in a market with heterogeneous agents, Advances in Applied Microeconomics 6 (1996) 121–156.

[23] E. Manisterski, D. Sarne, S. Kraus, Cooperative search with concurrent interactions, Artificial Intelligence Research (JAIR) 32 (2008) 1–36.

[24] J. McMillan, M. Rothschild, Search, in: R. Aumann, S. Amsterdam (eds.), Handbook of Game Theory with Economic Applications, 1994, pp. 905–927.

[25] J. M. McNamara, E. J. Collins, The Job Search Problem as an Employer-Candidate Game, Journal of Applied Probability 27 (4) (1990) 815–827.

[26] P. Morgan, A model of search, coordination, and market segmentation, mimeo, SUNY Buffalo (1995).

[27] Y. Nahum, D. Sarne, S. Das, O. Shehory, Two-sided search with experts, Autonomous Agents and Multi-Agent Systems 29 (3) (2015) 364–401.

[28] I. Rochlin, D. Sarne, M. Mash, Joint search with self-interested agents and the failure of cooperation enhancers, Artificial Intelligence 214 (2014) 45–65.

[29] A. E. Roth, The evolution of the labor market for medical interns and residents: A case study in game theory, Journal of Political Economy 92 (6) (1984) 991.

[30] A. E. Roth, The NRMP as a labor market: understanding the current study of the match, Journal of the American Medical Association 275 (1996) 1054–1056.

[31] A. E. Roth, E. Peranson, The effects of the change in the NRMP matching algorithm, JAMA 278 (9) (1997) 729–732.

[32] M. Rothschild, Searching for the lowest price when the distribution of prices is unknown, Journal of Political Economy 82 (1961) 689–711.

[33] D. Sarne, S. Kraus, Time-variant distributed agent matching applications, in: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), 2004.

[34] D. Sarne, S. Kraus, Managing parallel inquiries in agents' two-sided search, Artif. Intell. 172 (4-5) (2008) 541–569.

[35] D. Sarne, E. Manisterski, S. Kraus, Multi-goal economic search using dynamic search structures, Autonomous Agents and Multi-Agent Systems 21 (2010) 204–236.

[36] R. Shimer, L. Smith, Assortative matching and search, Econometrica 68 (2) (2000) 343–370.

[37] L. Smith, The marriage model with search frictions, Journal of Political Economy 114 (6) (2006) 1124–1146.

[38] L. Smith, Frictional matching models, Annual Reviews in Economics 3 (2011) 319–338.

[39] G. Stigler, The economics of information, Journal of Political Economy 69(3) (1961) 213–225.

[40] K. Sycara, S. Widoff, M. Klusch, J. Lu, Larks: Dynamic matchmaking among heterogeneous software agents in cyberspace, AAMAS 5 (2002) 173–203.

[41] N. Tsvetovat, K. Sycara, Y. Chen, J. Ying, Customer coalitions in electronic markets, in: Proceedings of AMEC 2000, Barcelona, 2000.

[42] K. J. Williams, A reexamination of the NRMP matching algorithm. national resident matching program., Academic Medicine 70 (6) (1995) 470–6.