Distributed Matching with Mixed Maximum-Minimum Utilities

Amos Azaria, David Sarne and Yonatan Aumann
Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel
{*azariaa1,sarned,aumann*}*@cs.biu.ac.il*

***Abstract*—In this paper we study distributed agent matching in environments characterized by costly exploration, where each agent's utility from forming a partnership is influenced by both the maximum and the minimum among the two agent's competence. This kind of utility function is somehow more applicable, compared to the one used in related work that takes the utility to be either the type of the agent partner or "standard" functions such as average or multiplication of the two types. The use of the hybrid min-max utility function is favorable whenever the performance of the agents forming a partnership is principally affected by the most (or least) competent among the two. This paper supplies a cohesive analysis for the min-max case, proving the equilibrium structure for the different min-max linear combination that may be used. We show that, in equilibrium, in any case that an agent sets its acceptance threshold below its own type it is guaranteed that any agent with a type between this threshold and its own will accept it (the agent) as a partner as well. This result substantially facilitates the calculation of equilibrium for such settings, e.g., when the set of types is finite.**

I. INTRODUCTION

Two-sided search is an important technique for modeling distributed matching processes 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 [8], [1] cannot be applied. Coalition formation mechanisms are also inapplicable, in particular in settings where the number of agents is substantial as discussed in this study.

The goal of each agent participating in the distributed matching process is to form a pairwise partnership that is optimally beneficial [11], [2], [4].Each agent is associated with a specific type that captures some property (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 (i.e., 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 benefits from continuing the exploration, potentially forming a better partnership in the future, and the costs associated with the future explorations.

During the past two decades, there has been a substantial progress in the analysis of distributed matching models (see [13] 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 (or simply the average of both agents forming a partnership) 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 [5], [6], [11]. 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 [12], [10].

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 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 mostly affected by the least competent player, as the other team will try to gain game points mostly by aiming the ball in his direction.

The paper presents an extensive analysis of the model with the new min-max utility function. The analysis results in an equilibrium characterization that for some cases is different from the one found for two-sided search with traditional utility functions. 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 that 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 function depends on both types, we manage to distinguish between three equilibrium patterns, where the first is assortative matching, the second is perfect segregation and we refer to the third as 'bumpy steps' where the acceptance threshold as a function of the agent type increases and decreases alternately. 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.

II. Related Work

The two-sided search for partnerships is a sub-domain of coalition formation. While coalition formation models usually consider general coalition-sizes [15], 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 [9], [14] (e.g., in the case of buyers and sellers or peer-to-peer applications). Various centralized matching mechanisms can be found in the literature [7], [3], [8]. However, in many MAS environments, in the absence of any reliable central matching mechanism, the matching process is completely distributed.

Two-sided 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 [11], [5], others assume a function defined over both types [2]. 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 [5] or additive explicit search costs [11], [6], [2]. Lastly, the models are distinguished by the nature of the utility earned by each of the agents (transferable [2] and "non-transferable" [5], [6]). 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.

III. General Model

The model used in this paper is a standard two-sided search model, and all assumptions given in this section are common in distributed matching literature [6]. The only differentiating element is the utility function used in this paper, which is a function of the minimum and maximum in this case.

We consider an environment populated with an infinite number of self-interested fully rational agents. Each agent is associated with a type x defined over 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 an agent associated with a type x obtains from forming a partnership with another agent associated with a type y is given by $u(x, y)$. This function is continuous and monotonic non-decreasing in x and y and is shared by agents of all types. In this work we focus on the case where the utility is of the form:

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

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 in two-sided search models has prior information concerning the type of specific other agents, it initiates interactions with other agents randomly. The nature of the two-sided search application suggests 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.

After two agents interact and learn the type of the other agents, 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 search activity (the 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 from the partnership formed minus the search costs accumulated along the search process.

A. Notation

We use $\phi(x)$ to denote the expected utility for an agent of type x from the search when using an optimal strategy.

Define $\underline{a}(x) = \min\{x' | \phi(x) \leq u(x, x')\}$ i.e. the minimum type for which the expected benefit from accepting it is greater or equal to the expected benefit of the agent if resuming the search. We later show that we can assume each agent of type x accepts all and only agents with types greater than or equal to $\underline{a}(x)$. Occasionally we will use $a(x)$ as a strategy for an agent of type x to indicate that the agent accepts all and only agents in $[a(x), 1]$; this might not be the optimal strategy for the agent (and therefore not in equilibrium), however, never do we consider more than a single agent deviating from equilibrium strategies. Similarly, we will use $\Phi(x, t)$ to denote the expected utility for an agent of type x when using a strategy according to which it accepts other agents only if their type is above threshold t (once again this threshold might not be optimal).

We use $A(x)$ to denote the group of agents that accept a match with an agent of type x and $\overline{A}(x)$ as the group of agents of type greater than x that accept a partnership with x: $\overline{A}(x) = \{y | y \in A(x) \wedge y > x\}$.

We use $M(x)$ to denote the matching group of an agent of type x, that is, the group of types with which x might end up forming a partnership: $M(x) = \{y | y \in A(x) \wedge x \in A(y)\}$. $\overline{M}(x)$ will denote the group of all agents with a type above x who are in the matching group of an agent of type x, i.e. $\overline{M}(x) = \{y | y \in M(x) \wedge y > x\}$.

IV. Individual Strategies and Expected Benefit

We begin with the analysis of individual expected-utility-maximizing strategies. Proposition 1 suggests that the expected-utility-maximizing strategies are reservation-value (i.e., threshold) based and Proposition 2 shows that the expected benefit of any individual agent when using such a strategy equals the reservation-value used.

Proposition 1: An agent of type x accepting all and only agents with types in $[\underline{a}(x), 1]$ will maximize its expected utility.

Proof: Assume agent of type x uses an optimal strategy P, 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 $x' < \underline{a}(x)$, the agent may reject it (regardless of P) since in the next iteration the agent may resume using strategy P and therefore gain $\phi(x)$ which, due to the monotonicity of u and definition of $\underline{a}(x)$, is greater than $u(x, x')$. If the agent interacts with a partner of type $x' \geq \underline{a}(x)$, the agent may accept it (regardless of P); the agent does not lose out by doing so, due to the monotonicity of u and definition of $\underline{a}(x)$, $u(x, x') \geq \phi(x)$.

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

The following is another known property which claims that if an agent of type x does not accept everyone, then its expected utility from the search is equal to its utility from forming a partnership with its threshold.

Proposition 2: If $\underline{a}(x) > 0$ then $\phi(x) = u(x, \underline{a}(x))$.

Proof: $\underline{a}(x) > 0$ therefore $\phi(x) > u(x, 0)$. From the monotonicity of u, $\phi(x) \leq u(x, 1)$. Since u is continuous, according to the intermediate value theorem there is a type x' where $\phi(x) = u(x, x')$. From the definition of $\underline{a}(x)$ we obtain that $\underline{a}(x) \leq x'$ and from the monotonicity of u we obtain that $\phi(x) = u(x, \underline{a}(x))$. ■

Next, we prove a theorem that claims that in any case where an agent's equilibrium acceptance threshold is below its own type, it is guaranteed that any agent with a type between this threshold and its own type will accept it (the agent) as a partner as well. Based on this theorem we will later propose an algorithm that finds the acceptance threshold for every agent. However, we first prove several lemmas which not only simplify the proof of the theorem, but also have significance on their own.

The following lemma claims that the higher the agent type is, the more agents accept it as a partner:

Lemma 1: For any x and $x' < x$, $A(x') \subset A(x)$.

Proof: Given $\tilde{x} \in A(x')$ by Proposition 1 $x' \geq \underline{a}(\tilde{x})$. Therefore $x \geq \underline{a}(\tilde{x})$ and again by Proposition 1 $\tilde{x} \in A(x)$. ■

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

Lemma 2: For any x and $x' < x$, $\phi(x') \leq \phi(x)$.

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

Theorem 1: For any x and $x' < x$, if $\underline{a}(x) \leq x'$ then $\underline{a}(x') \leq x$.

Proof: Assume by contradiction that $x < \underline{a}(x')$. From Proposition 2 and the definition of u we obtain that:

$$\phi(x) = \alpha \cdot \underline{a}(x) + (1 - \alpha) \cdot x \tag{1}$$

However, since $x < \underline{a}(x')$ and $\underline{a}(x) \leq x'$,

$$\alpha \cdot \underline{a}(x) + (1 - \alpha) \cdot x < \alpha \cdot x' + (1 - \alpha)\underline{a}(x') = \phi(x') \tag{2}$$

which contradicts Lemma 2 ■

Theorem 1 and Proposition 1 imply that if $\underline{a}(x) < x$ then the matching group $M(x)$ is $A(x) \cap [\underline{a}(x), 1] = [\underline{a}(x), x] \cup \overline{A}(x)$.

From Theorem 1 we conclude that agents can assume that any partners of lower types will accept a partnership if they accept it. Therefore, the acceptance threshold of each agent does *not* depend on the acceptance threshold of agents of lower types (than its own).

V. Minimum as the Utility Function

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

Intuitively, each agent accepts any other agent of a type equal to or greater than its own, since its utility is bounded by its own type. For similar reasons each agent of type x accepts any other agent with a value which is greater than $x - c$ (since resuming the search incurs a cost c).

Based on Theorem 1, the expected payoff for an agent of type x from the search, when using a threshold t, is given by:

$$\Phi(x,t) = -c + \int_t^x y f(y) dy + \int_{\overline{M}(x|a(x)=t)} x f(y) dy + \Phi(x,t)\big(1 - \int_{M(x|a(x)=t)} f(y) dy\big) \tag{3}$$

Where the first integrand is for a case where the agent forms a partnership with a partner of a lesser type (than

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

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 when the agent does not form a partnership and resumes the exploration. (Obviously, if $t \geq x$ then $\int_{a(x)}^{x} yf(y)dy = 0$).

From Equation 3 we obtain:

$$\Phi(x,t) = \frac{-c + \int_t^x yf(y)dy + \int_{\overline{M}(x|a(x)=t)} xf(y)dy}{\int_{M(x|a(x)=t)} f(y)dy} \quad (4)$$

The agent will use the threshold t that maximizes $\Phi(x,t)$ according to 4.

Figure 1 demonstrates the strategies used in equilibrium for the case where the utility function is the minimum, $f(x)$ is the uniform distribution and $c = 0.005$. All figures in this paper were generated using a discrete evaluation process, as described in section VIII.

In section VII 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 has a different threshold (depending on its type) and the higher the agent's type is the higher its threshold.

VI. Maximum as the Utility Function

We now consider the equilibrium in an environment where $\alpha = 0$ and therefore the utility from forming a partnership between any two agents is the maximum value among the two, i.e. $u(x, y) = \max\{x, y\}$.

Similar to when the utility is the minimum, the expected utility for each agent is given by:

$$\Phi(x,t) = \frac{-c + \int_t^x xf(y)dy + \int_{\overline{M}(x|a(x)=t)} yf(y)dy}{\int_{M(x|a(x)=t)} f(y)dy} \quad (5)$$

Clearly, an agent of type 1 will receive a utility 1 from any partnership, therefore it will accept any partner in its first search round ($\underline{a}(1) = 0$). Due to the search cost, agents with types near 1 accept any partner as well. Let $x^* = \inf\{x|\underline{a}(x) = 0\}$ (i.e., the agent with the lowest type which accepts any other agent).

Lemma 3: Agents with types smaller than x^* reject partners of types smaller than or equal to their own.

Proof: Given an agent with a type $x < x^*$, it rejects partnership with an agent with type 0, which would give it a utility of x. Therefore, by definition of $\underline{a}$(x), the expected utility of the agent is greater than x. Therefore, by Proposition 2, it rejects any other agent with a type smaller than or equal to x (which would give it a utility of x as well). ■

Lemma 4: No partnerships are formed between two agents with types smaller than x^*.

Proof: Suppose by contradiction that a partnership is formed by two agents with types $x' \leq x < x^*$. By Lemma 3, the agent with type x must reject an agent with type x' as a partner. ■

Lemma 5: All agents with types smaller than x^* set their threshold at x^*.

Proof: We first show that all agents with types smaller than x^* use the same threshold. Agents with type above x^* accept any other agent and by Theorem 1 an agent with type $x < x^*$ must accept an agent of type x^*. Therefore, combined with Lemma 4, all agents with type smaller than x^* have the same matching group. Therefore, they all have the same expected utility and by Proposition 2 the same threshold.

According to Theorem 1 an agent with type $x < x^*$ cannot set its threshold above x^*. Assume by contradiction that all the agents set their threshold at some point $t < x^*$. In this case an agent with type $t < x' < x^*$, must set its threshold at t as well, contradicting Lemma 3. ■

Summarizing the equilibrium, we obtain the following step function:

$$\underline{a}(x) = \begin{cases} 0 & \text{if } x^* \leq x \\ x^* & \text{otherwise} \end{cases} \quad (6)$$

Note that this solution has very interesting properties. First, any agent with a value greater than x^* accepts any partner. Second, partnerships are formed only if at least one side has a value greater than x^*.

Based on Proposition 2 agents of types smaller than x^* have an expected utility identical to the utility of forming a partnership with an agent of type x^* which equals x^*. Using this equilibrium structure, Equation 5 obtains:

$$\frac{-c + \int_{x^*}^1 yf(y)dy}{\int_{x^*}^1 f(y)dy} = x^* \quad (7)$$

and therefore:

$$c = \int_{x^*}^1 (y - x^*)f(y)dy \quad (8)$$

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

Figure 2 depicts the strategies used in equilibrium for the case where the utility function is the maximum, $f(x)$ is the uniform distribution and $c = 0.005$.

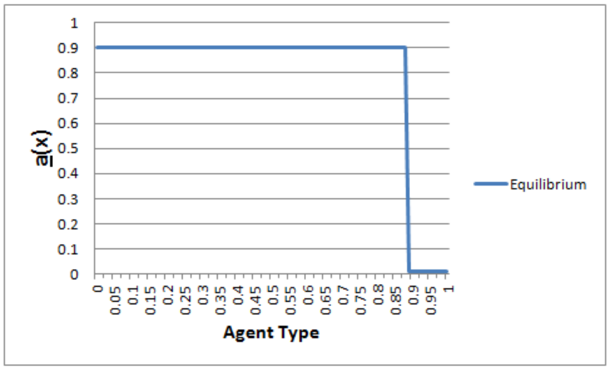

Figure 2. $\underline{a}(x)$ for $u(x,y) = \max\{x,y\}$

VII. Mixed Maximum-Minimum as the Utility Function

We now turn to analyze the equilibrium in environments where 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 agents of type x in this case is:

$$\arg\max_t \Big(-c+\alpha\big(\int_t^x yf(y)dy+(\int_{\overline{M}(x|a(x)=t)} f(y)dy\cdot x)\big) + (1-\alpha)\big((\int_t^x f(y)dy)\cdot x + \int_{\overline{M}(x|a(x)=t)} yf(y)dy\big)\Big)\cdot \frac{1}{\int_{M(x|a(x)=t)} f(y)dy} \tag{9}$$

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

The following theorem states that 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 that of 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 2: For any x, x', such that $x' < x$, $A(x) = A(x')$ and $0 < \underline{a}(x) \leq x'$:

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

Proof: Consider the case where $0 < \alpha < 0.5$. Since $0 < \underline{a}(x)$ then (from Proposition 2):

$$\phi(x) = u(x, \underline{a}(x)) \tag{10}$$

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

$$\phi(x) = \alpha\cdot\underline{a}(x)+(1-\alpha)\cdot x = \alpha\cdot\underline{a}(x)+(1-\alpha)\cdot x'+(1-\alpha)\cdot(x-x') \tag{11}$$

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

$$\phi(x) = \Big(-c+\int_{\underline{a}(x)}^{x'}(\alpha y+(1-\alpha)x)f(y)dy+ \int_{x'}^{x}(\alpha y+(1-\alpha)x)f(y)dy+\int_{\overline{A}(x)}(\alpha x+(1-\alpha)y)f(y)dy\Big)\cdot \frac{1}{\int_{M(x)} f(y)dy} \tag{12}$$

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

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

Subtracting Equation 13 from Equation 12 obtains:

$$\phi(x) - \Phi(x', \underline{a}(x)) = \Big(\int_{\underline{a}(x)}^{x'} f(y)dy\cdot(1-\alpha)(x-x') + \int_{x'}^{x}(\alpha y+(1-\alpha)x)f(y)dy - \int_{x'}^{x}(\alpha x'+(1-\alpha)y)f(y)dy + \int_{\overline{A}(x)} f(y)dy\alpha(x-x')\Big)\cdot\Big(\frac{1}{\int_{M(x)} f(y)dy}\Big) \tag{14}$$

Denote:

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

From Equation 14 we obtain:

$$\Phi(x', \underline{a}(x)) = \phi(x) - \zeta \tag{15}$$

Replacing $\phi(x)$ with $\alpha\cdot\underline{a}(x)+(1-\alpha)\cdot x'+(1-\alpha)\cdot(x-x')$ (Equation 11) obtains:

$$\Phi(x', \underline{a}(x)) = \alpha\cdot\underline{a}(x)+(1-\alpha)\cdot x'+(1-\alpha)\cdot(x-x')-\zeta \tag{16}$$

We now show that ζ is smaller than $(1-\alpha)\cdot(x-x')$. $\alpha < 0.5$ implies that $1-2\alpha > 0$, therefore:

$$\int_{x'}^{x}(((1-\alpha)x-\alpha x')-(1-2\alpha)\mathbf{y})f(y)dy < \int_{x'}^{x}(((1-\alpha)x-\alpha x')-(1-2\alpha)\mathbf{x'})f(y)dy = \int_{x'}^{x} f(y)dy(1-\alpha)(x-x') \tag{17}$$

Clearly:

$$\int_{\overline{A}(x)} f(y)dy\alpha(x-x') < \int_{\overline{A}(x)} f(y)dy(1-\alpha)(x-x') \tag{18}$$

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

$$(1 - \alpha) \cdot (x - x') > \zeta \tag{19}$$

Putting together Equations 16 and 19 we obtain:

$$\alpha \cdot \underline{a}(x) + (1 - \alpha) \cdot x' < \Phi(x', \underline{a}(x)) \tag{20}$$

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

Using the exact same proof (except for changes in the inequality starting at Equation 17), we obtain that if $\alpha > 0.5$ then $\underline{a}(x) > \underline{a}(x')$. We also obtain that when $\alpha = 0.5$ then $\underline{a}(x) = \underline{a}(x')$. ■

Inspired by Theorem 2, when examining the equilibrium in the mixed case we split the α's into three different cases:

- $0.5 < \alpha \leq 1$: where the minimum type has a greater impact (and when $\alpha = 1$ solely determines the utility).
- $\alpha = 0.5$: where both the minimum and the maximum types have equal effect on the agents' utility.
- $0 < \alpha < 0.5$: where the maximum type has a greater impact.

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

When the minimum type has a greater impact, we obtain assortative matching, where the higher the agent's type is the higher its threshold is, and all agents set their threshold beneath their own type. See Figure 3 for an example where $\alpha = 0.8$, uniform $f(x)$ and $c = 0.005$. Assortative matching is widely common in distributed matching, and occurs with many utility functions, such as $u(x, y) = xy$ [12]. The following Lemmas prove assortative matching for $0.5 < \alpha \leq 1$.

The following Lemma claims that the more agents accept a certain type, the higher its threshold becomes (or the more picky it can become).

Lemma 6: For any type x, if $S' \subset S$ then $\underline{a}(x|A(x) = S') \leq \underline{a}(x|A(x) = S)$.

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

The following Lemma claims that when $\alpha > 0.5$, $\underline{a}(x)$ is monotonously increasing in x.

[1] 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 14 and its sequels.

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

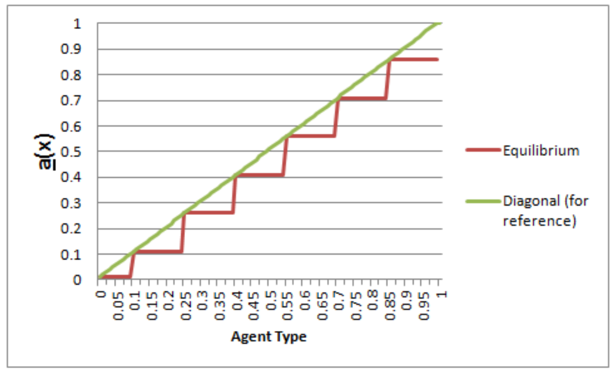

Figure 4. Simulation results for $\alpha = 0.5$

Lemma 7: If $\alpha > 0.5$, for every $x' < x$, if $0 < \underline{a}(x) < x$ then $\underline{a}(x') < \underline{a}(x)$.

Proof: From Theorem 2 and Lemma 6. ■

The following Lemma claims that when $\alpha > 0.5$, every agent sets its threshold below its own type.

Lemma 8: If $\alpha > 0.5$, then $\underline{a}(x) < x$ for every x.

Proof: Clearly, for every x, $\int_{M(x)} f(x)dx > 0$ (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 Lemma 7 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 $\int_{M(x)} f(x)dx = 0$, contradicting the above. ■

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

In the case of $\alpha = 0.5$, both the minimum and the maximum types have equal effects on the agents' utility, therefore the utility is simply the average between the two agents forming a partnership. This exact configuration is described in [6] where the entire domain is partitioned and partnerships are formed only within these partitions (termed perfect segregation). These results are intensified by Theorem 2. See Figure 4 for an example of the case where $\alpha = 0.5$, uniform $f(x)$ and $c = 0.005$.

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

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

The case in which $0 < \alpha < 0.5$ is both the most challenging and interesting case. Figure 5 gives an example for $\alpha = 0.2$, uniform $f(x)$ and $c = 0.005$. From the figure we observe the following properties when following the threshold as a function of the agent type from right to left (starting at $x = 1$):

- The threshold increases until there is an agent who sets its threshold at its own type (i.e. until the diagonal is reached).
- In any case where an agent sets its threshold above its own type (any point above the diagonal) the threshold decreases (as the type decreases).
- In several cases there is a sudden drop in the threshold. In our example this happens in $\{0.77, 0.54, 0.31, 0.1\}$. After the sudden drop, the threshold resumes climbing until once again there is an agent who sets its threshold at its own type.

To the best of our knowledge, this is the only equilibrium examined where, despite the utility function being monotonic, there exist types which set their threshold above their own type, and the threshold as a function of the agent type has segments which increase as the type decreases.

The first property is an immediate corollary of Theorem 2, 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 agents type is, the higher its threshold (Theorem 2).

The following lemmas prove the second property

The following Lemma claims that 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 9: Given x and $x' < x$ if $x \le \underline{a}(x)$ and $A(x) = A(x')$ then $\underline{a}(x') = \underline{a}(x)$.

Proof: Since $x \le \underline{a}(x)$:

$$\phi(x) = \alpha \cdot x + (1-\alpha) \cdot \underline{a}(x) = \alpha \cdot x' + \alpha \cdot (x - x') + (1-\alpha)\underline{a}(x) \quad (21)$$

On the other hand:

$$\phi(x) = \Phi(x', \underline{a}(x)) + \frac{\int_{M(x)} f(y)dy\alpha(x - x')}{\int_{M(x)} f(y)dy} = \Phi(x', \underline{a}(x)) + \alpha(x - x') \quad (22)$$

This implies:

$$\Phi(x', \underline{a}(x)) = \alpha \cdot x' + (1-\alpha) \cdot \underline{a}(x) = u(x', \underline{a}(x)) \quad (23)$$

and therefore $\underline{a}(x') = \underline{a}(x)$. ∎

The following Lemma claims that 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.

Lemma 10: Given x and $x' < x$, if $x \le \underline{a}(x)$ then $\underline{a}(x') \le \underline{a}(x)$.

Proof: From Lemma 1, $A(x') \subset A(x)$, hence the proof is immediate using Lemmas 9 and 6. ∎

We now provide an intuitive explanation for the third property. We do not provide a proof, since we believe this property depends on the density function ($f()$). Taking the example analyzed in Figure 5, consider for instance an agent with a type of 0.77. Note that the threshold set by the agents with a type of 1 is slightly above 0.77. Therefore, almost all agents with types greater than 0.77 reject partnerships with the agent. Therefore its expected utility drops drastically and so does its threshold. Since almost no agents with types above 0.77 accept it as a partner, all agents with types slightly under 0.77 are accepted by nearly the same group of agents ($A(\cdot) = [0, 0.77]$). Therefore according to Theorem 2 the threshold resumes climbing until once again there is an agent who sets its threshold at its own type.

VIII. Discrete Case

The analysis presented in the former sections completely unfolds the structure of the equilibrium in this important new class of two-sided search settings. 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.

Define s as the discrete step size. Under the discrete case, the expected outcome for agent x can be calculated using the following equation (a discretization of Equation 9):

$$\phi(x) = \Big(-c + \alpha\big(\sum_{y=\underline{a}(x)}^{x} yf(y)s + (\sum_{y \in \overline{M}(x)} f(y)s \cdot x)\big) + (1-\alpha)\big((\sum_{y=\underline{a}(x)}^{x} f(y)s) \cdot x + \sum_{y \in \overline{M}(x)} yf(y)s\big)\Big) \cdot \frac{1}{\sum_{y \in M(x)} f(y)s} \quad (24)$$

The algorithm performs a single pass on all types from type 1 and down. It finds the threshold for every type based on the thresholds calculated for greater types. Note that the order in which the threshold array ($\underline{a}$) is filled is crucial and

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

Input: Discrete step size (s), 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 $\lceil 1/s \rceil$ with 0's
  for each x, starting at $x = 1$ down by s **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 24 and choose the maximizer}
    Tie break using $\underline{a}[x] = \arg\min_y\{u(x, y) - \phi(x)\}$
  end for
  return $\underline{a}$.

must be started with the agents of type 1 and downwards, since every agent's threshold depends on the thresholds of all agents with greater types. However, based on Theorem 1, the agent's threshold does *not* depend on agents of lower types, which implies the correctness of the algorithm.

IX. Conclusions and Future Work

In this paper we study fully distributed matching with search costs, where each agent autonomously searches for a partner. We consider settings where the output function for a pair forming a partnership is a linear combination of the maximum and the minimum among the partners. This utility combination is much more applicable in real life (compared to traditional models), where usually either the best or the worst among the partners have greater influence on the resulting utility.

The analysis of the model introduces the unique properties of the equilibrium in the setting considered. In particular it includes a proof that in any case that an agent sets its acceptance threshold below its own type, it is guaranteed that any agent with a type between this threshold and its own, will accept it (the agent) as a partner as well. 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 for traditional utility functions, the equilibrium proved 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.

Naturally the attempt to integrate "search theory" techniques into day-to-day applications brings up the applicability question. Justification and legitimacy considerations for this integration were discussed in the wide literature we refer to throughout the paper. The current paper is not focused on re-arguing applicability, but rather on equilibrium analysis. 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 two-sided search model.

In future work we intend to study methods in which a system designer may increase the social welfare, defined as the integral on the expected utility of all participants. Consider for instance a University that can hand out scholarships to stronger students which form a partnership with weaker students. We will analyze the most efficient way to give out these scholarships so as to raise the social welfare.

X. Acknowledgments

We thank Chen Hajaj and Shira Abuhatzera for their help in this research. This work was supported by ISF/BSF grants 1401/09 and 2008-404.

References

[1] E. Anshelevich, S. Das, and Y. Naamad. Anarchy, stability, and utopia: Creating better matchings. *JAAMAS*, 2012. Forthcoming.

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

[3] G. Becker. A theory of marriage. *Journal of Political Economy*, 81:813–846, 1973.

[4] F. Bloch and H. Ryder. Two-sided search, marriages, and matchmakers. *International Economic Review*, 41:93–115, 2000.

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

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

[7] K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proc. of IJCAI*, 1997.

[8] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Math. Monthly*, 69:9–15, 1962.

[9] M. Klusch. Agent-mediated trading: Intelligent agents and e-business. *Journal on Data and Knowledge Engineering*, 36(3), 2001.

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

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

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

[13] L. Smith. Frictional matching models. *Annual Reviews in Economics*, 3:319–338, 2011.

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

[15] N. Tsvetovat, K. Sycara, Y. Chen, and J. Ying. Customer coalitions in electronic markets. In *Proceedings of AMEC 2000*, pages 121–138, Barcelona, 2000.