On the Social Welfare of Mechanisms for Repeated Batch Matching*

Elliot Anshelevich
eanshel@cs.rpi.edu
Rensselaer Polytechnic Institute

Meenal Chhabra and **Sanmay Das**
{meenal,sanmay}@cs.vt.edu
Virginia Tech

Matthew Gerrior
gerrior.matthew@gmail.com
GreaneTree Technology

Abstract

We study hybrid online-batch matching problems, where agents arrive continuously, but are only matched in periodic rounds, when many of them can be considered simultaneously. Agents not getting matched in a given round remain in the market for the next round. This setting models several scenarios of interest, including many job markets as well as kidney exchange mechanisms. We consider the social utility of two commonly used mechanisms for such markets: one that aims for stability in each round (greedy), and one that attempts to maximize social utility in each round (max-weight). Surprisingly, we find that in the long term, the social utility of the greedy mechanism can be higher than that of the max-weight mechanism. We hypothesize that this is because the greedy mechanism behaves similarly to a soft threshold mechanism, where all connections below a certain threshold are rejected by the participants in favor of waiting until the next round. Motivated by this observation, we propose a method to approximately calculate the optimal threshold for an individual agent to use based on characteristics of the other agents participating, and demonstrate experimentally that social utility is high when all agents use this strategy. Thresholding can also be applied by the mechanism itself to improve social welfare; we demonstrate this with an example on graphs that model pairwise kidney exchange.

Introduction

Many matching scenarios operate in a hybrid online/batch mode, where agents arrive and wait until the next market clearing period. In any given clearing period, all candidates currently waiting are considered for a match. Those who are successfully matched leave the market, while others wait for the next clearing period. This describes scenarios ranging from kidney exchange (which clear every few weeks) to academic job markets (typically once a year).

There has been a lot of work on analyzing single rounds of markets that clear in batches as well as on designing and analyzing online matching algorithms. Typically, when thinking about a single round, issues of individual rationality and incentive compatibility dominate. There is now a significant literature on market design for achieving various different goals in such single-shot matching scenarios (Abraham et al. 2007). This work has been applied to the design of many job markets (Roth and Peranson 1999), assigning students to public schools (Abdulkadiroglu, Pathak, and Roth 2009), and to kidney exchanges, where incompatible donor-recipient pairs are matched (or placed into a longer chain) with others who have a compatible kidney available (Su and Zenios 2005; Roth, Sönmez, and Ünver 2004).

The goals for market design differ by domain. Typically, in labor markets, the focus is on finding stable matchings and preventing unraveling (Roth and Xing 1994). Recently there has been a focus on understanding the social utility of different matching mechanisms that could be employed in different markets (Anshelevich and Das 2010). For kidney exchange, there are general frameworks that take into account possible different welfare functions for society, but most of the work so far has focused on maximizing the number of compatible matches found in a round (Abraham, Blum, and Sandholm 2007). Ashlagi, Jaillet, and Manshadi (2013) find the appropriate batch size to run the exchange on so as to increase the number of matches without incurring negative consequences. Ünver (2010) studies barter mechanisms where the goal is to maximize the additive utility accounting for waiting time. In studies of assignment of students to public schools, the focus has been on incentive compatibility for parents when making their preference lists, and also on the welfare properties of different mechanisms in terms of student priorities (Abdulkadiroglu, Pathak, and Roth 2009).

The body of work on online weighted matching is large, usually using models where nodes arrive one at a time (Birnbaum and Mathieu 2008; Khuller, Mitchell, and Vazirani 1994), data stream models (McGregor 2005; Epstein et al. 2011), and others (Awasthi and Sandholm 2009). In the hybrid batch/online matching domain, there have typically only been two approaches: either form a stable matching in each round (as in many job markets), or form a maximum weighted matching in each round. The long-term social utility of these methods has not been analyzed in detail. Dickerson, Procaccia, and Sandholm (2012a) take a different approach, introducing a learning-based method for informing myopic algorithms in the context of kidney exchange. They

*A preliminary version appeared as an extended abstract in Proc. AAMAS (Anshelevich et al. 2012). Parts of this work were completed while all the authors were at Rensselaer.

focus on maximizing the number of donors and recipients matched rather than social welfare.

We are interested in understanding the social utility of different matching mechanisms that get used in every round of a repeated batch matching scenario. In this paper we perform an initial investigation, focusing on matching graphs with a particular utility structure: the utility received by each of two members of a pair is the same, and thus can be represented by the weight of the (undirected) edge between them. This corresponds to a world where the utility of a match is completely a function of the "compatibility" between the two matched agents. However, this model can also be used for mechanism design on graphs where the utilities received by the two members of a link are different, by allowing the weight of the edge to represent the sum of the two utilities. The social welfare maximization problem remains the same, although individual incentives may change.

With such preference structures it is known that the greedy matching is stable in a given round (Anshelevich, Das, and Naamad 2013). However, the stable matching does not maximize social utility. In many job markets, the greedy mechanism is a reasonable abstraction of the actual matching process in a given round, so we can gain understanding of long-term outcomes by studying the long-term social utility in repeated batch matching using the greedy mechanism. In other domains, like kidney matching, an explicit maximum weighted assignment is performed at each round (Abraham, Blum, and Sandholm 2007) .

Our Contributions. Surprisingly, we find that a mechanism that forms a greedy allocation in each round can lead to higher social utility in the long-term than the mechanism that forms a maximum weighted matching in each round. The intuition is that maximum weighted matching inefficiently matches agents with low compatibilities in a given round who would be better off waiting for the next round. Based on this intuition, we propose a new algorithm for agents to compute approximately optimal thresholds below which they should refuse to accept a match in any given period. We demonstrate experimentally that this algorithm comes close to yielding the best social utility of any threshold-based strategy, and significantly higher social utility than simply performing either greedy or maximum weighted allocations at each round. We also show that, once thresholds are picked according to our algorithm and all links between agents whose compatibility is below the threshold are removed, then even completely ignoring the quality of compatibility between agents while forming the matching still yields high social welfare. This provides experimental support for using unweighted matching mechanisms after picking appropriate thresholds to label each pair of agents "compatible" or "not compatible". Finally, we evaluate the benefits of our proposed thresholding mechanism on a network with many different (9) types of nodes, based on real-world probabilities of donor-recipient pairs in kidney exchange, and find that our thresholding mechanism can lead to a significant increase in social welfare over using no thresholds.

The Model

Time is discrete, and at each unit of time (a *round*) all agents who are thus far unmatched participate in a batch matching. At time $t = 0$ there are n agents, and at each future time period r new agents arrive. Agents can be thought of as nodes on a graph. The existence of an edge between two nodes means that the agents are compatible; there is a non-zero utility to both from being matched with each other. If there is no edge then neither agent gets any utility from being matched with the other.

Agents are of types $1, \ldots, k$. Type i is defined by a vector $(p_{i1}, p_{i2}, \ldots p_{ik})$ where p_{ij} is the probability that a type i agent is compatible with a type j agent (and $p_{ij} = p_{ji}$). The number of agents of each type that arrive at the beginning of each round is denoted by a vector $(r_1, r_2, \ldots r_k)$, where r_i is the number of arriving agents of type i.

When agents arrive (or at the start), a random graph is generated using the above probabilities to determine edge formation. If an edge is formed, it is associated with a weight u_{ij} that determines the utility of that matching. The distribution of weights u_{ij} is independent of the types of agents i and j and remains stationary; i.e., u_{ij}'s are i.i.d draws from a distribution $f(x)$ irrespective of the type of the agents to be connected and the time at which the edge is formed. Agents lose utility from waiting. An agent i that is matched with agent j after t rounds of waiting receives utility $\delta^t u_{ij}$, for a fixed discount factor $\delta \in (0, 1)$. The utilities received by the two agents need not be symmetric if they wait a different number of rounds. Social utility is additive: $U = \sum_{i,j \in \text{Matches}} u_{ij}(\delta^{t-t_i} + \delta^{t-t_j})$, where t_i and t_j are the arrival times of i and j respectively, and t is the time at which they are matched. While an agent suffers in personal utility from being unmatched (because of the discount factor associated with waiting), the utility of matching with that agent could still be high for a new agent entering the market who has not suffered the discounting penalty.

At each round, the matching mechanism (which has access to the entire graph, the utility structures, and knows the value of δ), assigns the agents to a particular matching, in which some agents may be unmatched. Before the mechanism decides on a matching, each individual agent can report its set of acceptable neighbors to the mechanism, and the mechanism can only match an agent with an acceptable neighbor. After reporting, however, agents cannot change their minds: they are obedient and accept any link that they are assigned by the mechanism. Agents that remain unmatched are eligible to be matched again in the next round.

Market Mechanisms and Agent Strategies

We consider situations where agents must accept any potential matching (which could be a good model for scenarios like organ matching, where any compatible matching should be acceptable), and situations where they pre-specify that some matchings are unacceptable (a better model for job markets and problems like marriage and working in teams).

Matching Mechanisms

We restrict our attention to mechanisms we call *non-retrospective*: in any given round, they only evaluate newly arrived agents for matches. In doing so, they consider matching newly arrived agents with agents left over from previous rounds, but do not re-evaluate matches between two agents who are both left over from previous rounds.

The *greedy mechanism* constructs a maximal greedy matching where pairs of agents are added to the matching starting with highest utility u_{ij}, until there are no more unmatched pairs with an edge between them. It is greedy on the edge weights and not on the actual utilities (which may be discounted). This mechanism is guaranteed to yield a stable matching in any given round, that is, there will not exist a pair a and b such that a and b would both rather be matched with each other than with their current partners. This has been proved for a single round when the weight of the edge represents the utility to both agents who share that edge (Anshelevich, Das, and Naamad 2013). In our context, the only change is that if an agent i gets matched with agent j at time t, it receives a utility of $u_{ij}\delta^{t-t_i}$, where t_i is the arrival time of agent i. The matching will not be roundwise stable if there exists an edge (i, j) not included in the matching, such that i and j would rather match with each other instead of their assigned matches (call these assigned matches k and ℓ). This implies that $u_{ij}\delta^{t-t_i} > u_{ik}\delta^{t-t_i}$ and $u_{ij}\delta^{t-t_j} > u_{j\ell}\delta^{t-t_j}$. Thus, $u_{ij} > \max(u_{ik}, u_{j\ell})$, which is a contradiction since then the greedy algorithm would match agents i and j with each other. (The argument for the case where one or both of i and j are unmatched is similar.)

The *max weight mechanism* constructs a maximum weighted matching at each round, taking actual utilities (including the discount factor) into account. This is social welfare maximizing in any given round, but it is not forward looking in that it does not consider future rounds.

In *threshold* mechanisms, the mechanism chooses thresholds τ_{ij} for every type combination (i, j) of agents, and forms either a max-weight or a greedy matching on the pairs of agents that are above the threshold. No matches are made by such mechanisms between i and j with $u_{ij} < \tau_{ij}$.

A Strategy For Rational Agents

We now describe a methodology for calculating (approximately) the optimum threshold to use, first motivating it from the perspective of individual agents making decisions about which potential matches to deem unacceptable prior to the mechanism being run. Such agents will want to choose a threshold exactly equal to their (discounted) expected value from being unmatched and remaining for the next round (where they again get to make the same decision). We then show that this method can be used by the mechanism designer to form a matching that achieves high social utility.

Suppose the mechanism assigns each agent either zero or one possible matches in each round. The only decision an agent has to make is, prior to any given round, which of its possible partners are acceptable. Specifically, the agent has to provide the mechanism with a function mapping from the space of possible utilities of its partners to a yes or no decision, specifying whether a partner of that utility is acceptable. We can show that, under certain conditions, an agent's optimal strategy is the same in any round, and can be characterized by a reservation value t^* such that the agent should (pre-)reject all potential matches with utility less than t^* and be ready to accept any match with utility greater than t^*.

The conditions involve one major assumption: that the expected number of agents of each type left unmatched by the mechanism in any given round is constant (we call this the **well-mixed assumption**). The assumption is not unreasonable: consider the whole dynamic mechanism as a long-running system. If the number of agents of any type being left unmatched were declining steadily, the system would reach an equilibrium where no agents of that type were being left unmatched. If the number of agents of any type being left unmatched were increasing, there would have to be natural exiting of the market (death or leaving for an alternative market), maintaining some number from previous rounds in equilibrium. Other possibilities are that the number of any type left unmatched follow either a well-defined cyclical pattern or a chaotic pattern, but we leave consideration of those possibilities to future work. We note that this is an assumption of the state of the system and arrival/departure of agents, rather than on the mechanism itself. For any particular mechanism, the natural state of the system will equilibrate to well-mixedness (although the particular well-mixed state may differ depending on the mechanism) because agents will enter/leave at different rates.

Theorem 1 *Suppose that at each round the same non-retrospective matching mechanism is used, and that the well-mixed assumption holds: the expected number of agents left unmatched is constant for each type. At the beginning of each round, an agent can report to the mechanism a mapping from utilities to the set {yes, no} that specifies which potential matches are acceptable. Then, the optimal strategy for any agent i is characterized by a reservation value t^*, such that the agent should reject all potential matches with utility less than t^* and accept any potential match with utility greater than t^*. t^* is the same at any round and equals the discounted value of the expected utility U_i the agent i gets in any round other than the round at which it arrives, if it doesn't get matched in the current round.*

Proof sketch: The proof involves first showing that the arrival process is stationary. Then, Bellman's Optimality Principle implies that the optimal policy involves making the same decision at each round. Finally, it is then easy to show that the optimal policy must be of a reservation value form. Details are omitted due to space considerations. □

t^* is the solution to the Bellman equation $t^* = \delta(t^* \Pr(\neg\mathrm{M}) + \mathbb{E}(\mathrm{Utility}|\mathrm{M}) \Pr(\mathrm{M}))$ where M represents the event that the agent is matched with another with the utility of the match being greater than t^*.

Approximately Optimal Thresholds Quantifying $\Pr(\mathrm{M})$ (i.e., the probability of getting matched with utility greater than t^*) is difficult in a random graph model. Instead, we propose an approximate method to calculate this threshold and demonstrate empirically that it yields close to optimal results. Let T represent the event that the maximum value of the weight $u_m = \max_j(u_{ij})$, among all the connections of the agent i, is greater than t^*.

Clearly $\Pr(\mathrm{M}) \leq \Pr(\mathrm{T})$ because a match may not be formed between the two links even though the weight is greater than t^*. We modify the Bellman equation above to:

$$t^* = \delta(t^* \Pr(\neg\mathrm{T}) + \mathbb{E}(u_m|\mathrm{T}) \Pr(\mathrm{T})) \quad (1)$$

The threshold t^* is calculated using the optimistic assumption that, in the next round, if an agent has any links greater than t^*, she will surely be matched to the highest of all links. The probability of making a connection with weight greater than t^* and the expected maximum utility can be calculated using order statistics. If an agent makes K connections in any given round, then the p.d.f. of the edge with highest weight is given by $f_m(y|K)\,dy = K[F(y)]^{K-1} f(y)\,dy$.

The agent i connects to each of the r agents who arrive in each time period with the type-appropriate probability (this is a Bernoulli trial). The number of connections K made by the agent is a Binomial random variable as it is a sum of independent Bernoulli trials. Therefore, $F_m(y) = \mathbb{E}_K(F_m(y|K)) = \prod_k (1 - p_{ik} + p_{ik}F(y))^{r_k}$, where r_k represents the number of agents of type k who arrive in next round. Then we can rewrite the approximate Bellman equation as $t^* = \delta\left(t^* + \int_{t^*}^{\infty} (1 - F_m(y))\,dy\right)$.

Although the proof of existence of a single true optimal threshold is only valid when the well-mixed assumption holds, the approximation can still be used when the system violates the assumption. The net impact of a growing number of agents remaining in the system would simply be that the approximation becomes a little more optimistic. In the next section, we demonstrate empirically that this method computes thresholds close to the optimal reservation value.

Experimental Results

In our experiments, for a variety of thresholds τ_{ij}, we calculate the resulting social utility when using the following set of mechanisms (in each case, any thresholding is enforced by the mechanism at the time of edge formation, i.e., before the matching process; agents cannot reject any matching after the final matching has been formed by the mechanism):

Online Maximum Weight Matching: The matching at each round is formed using the max-weight matching algorithm, using only edges such that $u_{ij} > \tau_{ij}$.

Online Greedy Algorithm: Similar, except that matchings are formed using the greedy algorithm in each round.

Online Maximal Matching: The mechanism removes all edges below τ_{ij}, then picks an arbitrary maximal matching.

Omniscient Matching: This mechanism knows the entire future in terms of which nodes and edges will be added to the matching graph, and calculates the maximum weight matching over all rounds simultaneously (taking discount factors into account when assigning utility weights). This is the optimal solution to the offline version of the problem, and therefore no online algorithm can perform better than this; thus, it is an upper bound.

OmniThresh Matching: This mechanism also has foresight; however the matching is calculated only after pruning the edges below τ_{ij}, providing an upper bound on the overall social utility of any threshold-based offline algorithm.

While the underlying process is a continuing one, we simulate performance by running the mechanism for a fixed number of rounds T; however only agents who arrive in the first T_m rounds contribute to measured social utility (the online mechanisms are unaware of T or T_m, but the upper-bound mechanisms know their values). Agents who stay in the system for at least $T - T_m$ rounds are assumed to have spent enough time in the system for their utility to be a good approximation of their utility in a continuing system.

Homogeneous Populations: Single Type

All agents connect to each other with the same propensity (we consider two probabilities of connection $p = 0.02, 0.06$). We set $n = 51$ (initial population), $r = 50$ (number of new agents per round), $\delta = 0.9$ (discount factor), $T = 40$ (rounds) and $T_m = 30$. We use the following distributions for individual utilities u_{ij}: (1) Exponential with rate parameter $\lambda = 1$; (2) Uniform on [0, 1]; (3) Lognormal with location (μ) and scale (σ) parameters (0,1).

First, we verify the approximate optimality of the threshold determination scheme. Figure 1 demonstrates that, if all other agents are using a single fixed threshold, no matter what that threshold is, there is a single best threshold for any individual agent to use. Moreover, that threshold is well approximated by Equation 1. The computed threshold is typically slightly higher because it is calculated using the optimistic assumption that, in the next round, the agent gets matched to the best link with utility better than t^*.

Threshold mechanisms can significantly improve social welfare. Figure 2 shows the improvement in social welfare due to using threshold mechanisms. The vertical blue line represents the approximately optimal threshold, and is close to the best threshold for maximizing social welfare (as well as to the best threshold for rational agents to use, as shown above). For example, with lognormal utilities ($\mu = 0, \sigma = 1$), using the appropriate threshold improves the competitive ratio of online mechanisms from about 0.4 to about 0.75.

In addition to the improvement in social welfare, Figure 2 shows several interesting properties of threshold mechanisms. First, we observe the unimodal behavior of the competitive ratio w.r.t threshold. This occurs because initially, adding a threshold helps in removing low quality links, thus minimizing the online effect and producing a matching which yields social welfare close to the offline optimal matching. However, if the threshold is too high, it will remove high-quality links and produce a sub-optimal matching. This effect can be further quantified by looking at the OmniThresh mechanism: for low thresholds it behaves as well as the Omniscient mechanism, since not being able to use low-quality links is not a large constraint on the matching quality produced by OmniThresh. Once the thresholds become large, however, OmniThresh starts doing worse than the unconstrained Omniscient mechanism. This indicates that the social welfare exhibited by our thresholded online mechanisms is most likely a combination of two counteracting effects: (1) Having a high threshold removes some of the "online" nature of the mechanism, since it no longer matches pairs on low-quality edges, and instead waits to

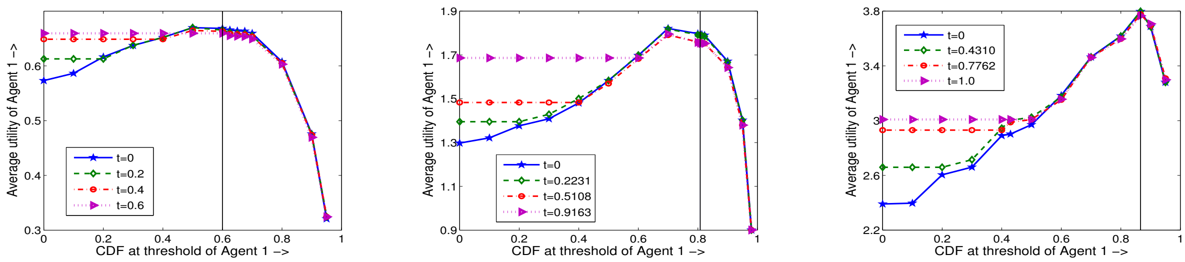

Utilities $\sim U[0\,1]$ Utilities $\sim$ Exponential($\lambda = 1$) Utilities $\sim$ Lognormal($\mu = 0, \sigma = 1$)

Figure 1: Utility received by Agent 1 when the mechanism is roundwise greedy as a function of its threshold. All other agents use a constant threshold t (see legend). The threshold computed using Equation 1 (the vertical line) approximately maximizes the utility of Agent 1, no matter what threshold other agents use. Here $p = 0.02, r = 50, n = 51$.

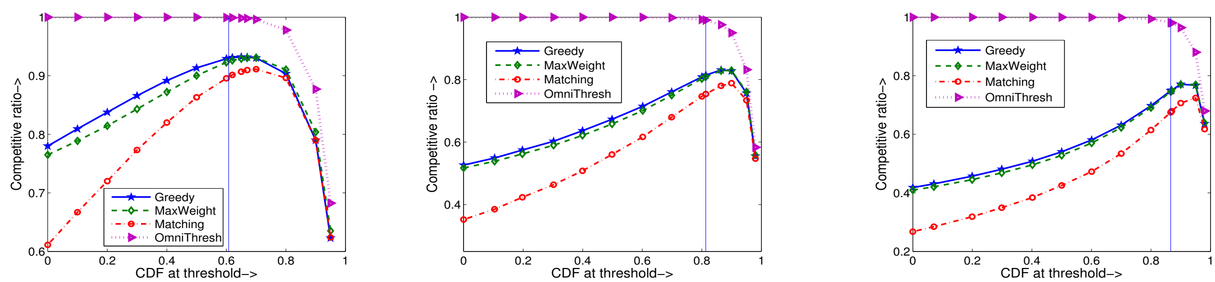

U1:Uniform($U[0\,1]$, $p = 0.02$) E1:Exponential ($\lambda = 1, p = 0.02$) L1:Lognormal(($\mu = 0, \sigma = 1$), $p = 0.02$)

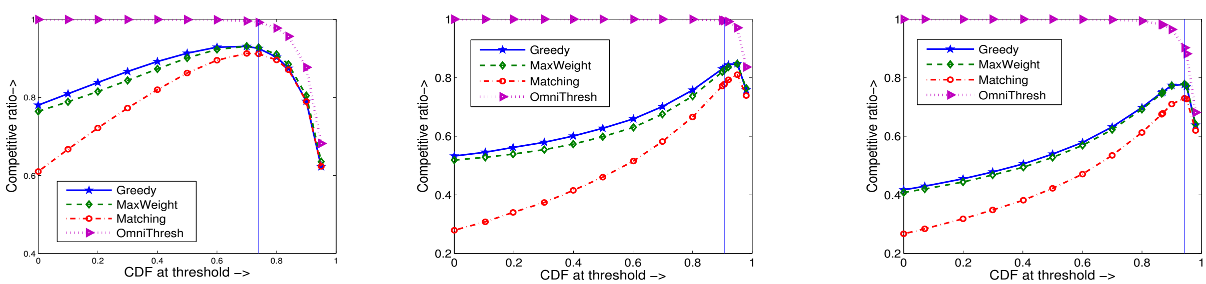

U2:Uniform ($U[0\,1]$, $p = 0.06$) E2:Exponential ($\lambda = 1, p = 0.06$) L2:Lognormal (($\mu = 0, \sigma = 1$), $p = 0.06$)

Figure 2: Competitive ratio of social welfare (compared with the Omniscient matching) as a function of threshold. Surprisingly, greedy matching yields higher social welfare than max weight matching at lower thresholds. All the curves are unimodal which shows that there exists only one optimal threshold. The vertical line represents the computed threshold (from Equation 1). Here $r = 50, n = 51$.

match them in future rounds, and (2) Having a high threshold removes high-quality edges from consideration, thus making a matching worse.

Greedy performs better than Max-Weight. Another interesting property apparent from Figure 2 is that the Greedy mechanism consistently performs better than the Max-Weight mechanism. For an individual round, the Greedy mechanism guarantees stability, while the Max-Weight mechanism maximizes social welfare. Thus, it seems surprising that, in aggregate, the Greedy mechanism is superior. The intuition is that the maximum weighted matching inefficiently matches agents with low-quality edges who would be better off waiting for the next round. The greedy matching instead tends to use the same high-quality edges as the maximum-weight matching, but uses fewer low-quality edges, allowing agents who would have received a low quality match to remain unmatched, and receive higher utility in future rounds. Only when the threshold is almost optimal does the greedy matching stop surpassing the maximum-weight matching in quality.

Thresholds matter more than edge weights: support for unweighted matching. Figure 2 shows that, while the Online Maximal Matching mechanism performs worse than mechanisms that explicitly consider edge weights, it still performs well (within just a few percent of the other online mechanisms) when the threshold is picked appropriately. Therefore, once thresholds are picked according to our algorithm and all links between agents whose compatibility is below the threshold are removed, even completely ignoring the quality of compatibility between agents while forming the matching yields high social welfare.

This is intriguing: suppose that exact utilities of each edge were noisy or difficult to evaluate. Then, the above mechanism is useful and intuitive. It only requires evaluating each pair as compatible or incompatible, and as many compatible pairs are matched in every round as possible. In fact, there is a potential relation to current practice in kidney exchange in the United States: while there has been discussion of more complex utility functions, the standard practice is to label pairs as compatible or incompatible, and assign a utility of 1 for compatible pairs: then the maximum weighted matching is just the maximum cardinality matching (although kidney exchange often considers longer cycles, rather than just pairwise matchings) (Abraham, Blum, and Sandholm 2007; Delmonico 2004). While there are ongoing disputes about whether HLA matching yields useful information about expected life of the transplanted kidney for live donors (Opelz and Döhler 2007), and the degree of HLA mismatch is used in cadaveric kidney allocation (Su and Zenios 2005), it is possible that the implicit thresholding used in making decisions about compatibility is good enough that maximum cardinality exchanges may also be close to socially optimal.

Heterogeneous Populations

Two Types Suppose there are two types of agents A and B such that type A agents are pickier, or more difficult to find partners for, than those of type B. Let $p_{A\text{-}A}$, $p_{B\text{-}B}$ and $p_{A\text{-}B}$ represent the probabilities of formation of the link types A-A, B-B and A-B respectively. This is the *only* difference between agents of different types: once a link is formed, the utility of that edge is independent of agent types.

As the probability of connection is different for the two types of agents, the threshold which maximizes the individual expected profit will also be different across agent types. The optimal threshold for an agent of type A will be less than that of an agent of type B because type A agents are less likely to connect and thus have fewer options available in the future. We ran three sets of experiments with different values for the connection probabilities, and different lognormal distributions of utilities (details of the parameters omitted due to space considerations). In each of our experiments, thresholds computed using Equation 1 turn out to satisfy certain useful conditions when the mechanism uses the max of the thresholds of the two nodes forming an edge as a cutoff. The resulting social welfare for type B agents is close to the optimal welfare for any threshold combination. Simultaneously, the threshold selected by type A agents is an approximate best response to the strategy of type B agents. Since type B agents are more "powerful" since they are better-connected, our choice of thresholds results in them receiving high social welfare. When the mechanism chooses the *minimum* of the individual thresholds as a cutoff, overall social welfare is often increased without significantly hurting the highly connected agents. This suggests that imposing a preference for "less connected" agents can result in higher social welfare, without significantly hurting the "more connected" agents.

Multiple Types We build a stylized model of a paired (two-way) kidney exchange. Each node represents a donor-patient pair. In each round 100 new donor-patient pairs arrive and links are formed between them based on compatibility. Let (X, Y) represent one donor-patient pair, where X and Y are the blood group of the donor and the patient respectively. Donors and recipients may be incompatible due to either blood-group incompatibility or positive cross-match (sensitization). (X_1, Y_1) and (X_2, Y_2) are linked if X_1 is compatible with Y_2 and X_2 is compatible with Y_1. There are 4 blood groups O, A, B, AB. AB patients are rare (3.85% of the population), plus they are universal recipients, so it is even rarer for them to need a match. Conversely, they are not often useful as donors, so we reduce the model to pairs formed from the other three types (leaving 9 types of nodes). We use real-world probabilities for different blood types (%ages of O, A, and B types are 48.1%, 33.7%, and 14.3% respectively) and positive crossmatches (11%) to generate simulated networks (Zenios, Woodle, and Ross 2001; Dickerson, Procaccia, and Sandholm 2012b; Saidman et al. 2006).

The utility from kidney exchange depends on several factors like age, race, gender, transplant history, PRA etc. However, it is known that the number of HLA matches is strongly correlated with graft-survival (Opelz and Döhler 2007), so following Su and Zenios (2005), we set the weight on each link to the number of HLA matches between the donor and the patient. We generate HLA mismatch probabilities from the frequency tables of proteins on each of the six HLA loci for Caucasians provided by Zenios (1996). We also assume, following Su and Zenios (2005), that all mismatches have equal effect and occur independently of each other.

Utilities can be different for both nodes on an edge. So the mechanism assigns an edge the sum of the utilities of the two agents. It computes the distribution for the sum of utilities (based on the combined number of HLA matches in the two donor-recipient pairs), and uses this to calculate the threshold for each link using Equation 1. We find that this mechanism achieves a 15% increase in social welfare compared with using no threshold. While still stylized, this demonstrates the potential of our method to increase social welfare in realistic networks with many types of nodes.

Discussion

It is surprising that the round-wise greedy mechanism outperforms the round-wise max-weight mechanism in our setting. Based on this insight, we propose a threshold-based mechanism, and find that it performs very well in several settings. While the theoretical optimality of thresholds for individuals only holds under the well-mixed assumption, our simulation results are all in settings where well-mixedness does not hold, and this provides further evidence of the applicability of the results and the algorithms in this paper. Many interesting questions remain: in particular, how the superiority of greedy or threshold mechanisms is affected by the valuation setting and the number of types of agents. Also, while we focus mostly on additive utility, other social welfare functions could be used in future analyses.

Acknowledgements

The authors acknowledge support from NSF grants IIS-0952918/1303350, CCF-1101495 and CCF-0914782, and US-Israel BSF grant no. 2008404.

References

Abdulkadiroglu, A.; Pathak, P.; and Roth, A. 2009. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review* 99(5):1954–1978.

Abraham, D.; Levavi, A.; Manlove, D.; and O'Malley, G. 2007. The stable roommates problem with globally-ranked pairs. *Internet and Network Economics* 431–444.

Abraham, D.; Blum, A.; and Sandholm, T. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, 295–304.

Anshelevich, E., and Das, S. 2010. Matching, cardinal utility, and social welfare. *ACM SIGECom Exchanges* 9(1):4:1–4:7.

Anshelevich, E.; Chhabra, M.; Gerrior, M.; and Das, S. 2012. On the social welfare of mechanisms for repeated batch matching. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*, 1321–1322.

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

Ashlagi, I.; Jaillet, P.; and Manshadi, V. H. 2013. Kidney exchange in dynamic sparse heterogenous pools. *arXiv preprint arXiv:1301.3509*.

Awasthi, P., and Sandholm, T. 2009. Online stochastic optimization in the large: Application to kidney exchange. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, 405–411.

Birnbaum, B., and Mathieu, C. 2008. On-line bipartite matching made simple. *ACM SIGACT News* 39:80–87.

Delmonico, F. 2004. Exchanging kidneys-advances in living-donor transplantation. *New England Journal of Medicine* 350(18):1812–1814.

Dickerson, J.; Procaccia, A.; and Sandholm, T. 2012a. Dynamic matching via weighted myopia with application to kidney exchange. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*.

Dickerson, J.; Procaccia, A.; and Sandholm, T. 2012b. Optimizing kidney exchange with transplant chains: Theory and reality. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, 711–718.

Epstein, L.; Levin, A.; Mestre, J.; and Segev, D. 2011. Improved approximation guarantees for weighted matching in the semi-streaming model. *SIAM Journal on Discrete Mathematics* 25(3):1251–1265.

Khuller, S.; Mitchell, S.; and Vazirani, V. 1994. On-line algorithms for weighted bipartite matching and stable marriages. *Theoretical Computer Science* 127:255–267.

McGregor, A. 2005. Finding graph matchings in data streams. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*. Springer. 170–181.

Opelz, G., and Döhler, B. 2007. Effect of human leukocyte antigen compatibility on kidney graft survival: Comparative analysis of two decades. *Transplantation* 84(2):137.

Roth, A., and Peranson, E. 1999. The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review* 89(4):748–780.

Roth, A. E., and Xing, X. 1994. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *American Economic Review* 84(4):992–1044.

Roth, A.; Sönmez, T.; and Ünver, M. 2004. Kidney Exchange. *Quarterly Journal of Economics* 119(2):457–488.

Saidman, S.; Roth, A.; Sönmez, T.; Ünver, M.; and Delmonico, F. 2006. Increasing the opportunity of live kidney donation by matching for two-and three-way exchanges. *Transplantation* 81(5):773.

Su, X., and Zenios, S. 2005. Patient choice in kidney allocation: A sequential stochastic assignment model. *Operations Research* 53(3):443–455.

Ünver, M. U. 2010. Dynamic kidney exchange. *The Review of Economic Studies* 77(1):372–414.

Zenios, S.; Woodle, E.; and Ross, L. 2001. Primum non nocere: Avoiding harm to vulnerable wait list candidates in an indirect kidney exchange. *Transplantation* 72(4):648.

Zenios, S. 1996. Health care applications of optimal control theory. *Ph.D. thesis, Operations Research Center, Massachussets Institute of Technology, Cambridge, MA*.