

# Matching, Cardinal Utility, and Social Welfare

ELLIOT ANSHELEVICH

Rensselaer Polytechnic Institute

and

SANMAY DAS

Rensselaer Polytechnic Institute

---

Categories and Subject Descriptors: J.4 [Computer Applications]: Social and Behavioral Sciences

General Terms: Algorithms, Economics

Additional Key Words and Phrases: Matching, Game Theory, Utility

---

## 1. INTRODUCTION

Matching markets have historically been an important topic in economics research. On the positive (descriptive) side, researchers have modeled everything ranging from marriage markets to labor markets using the framework of matching. Matching was also one of the first areas in which market design made a name for itself, perhaps most famously in the redesign of the market that matches graduating M.D.s to their first residency programs in the United States. The arrival of computer scientists to the field of market design in general can be traced to many of the reasons suggested recently by Conitzer [2010] (in a broader context than just market design) in an article in *Communications of the ACM*, including the effects of new markets that have been made possible by advances in networking and Internet technology, a more computational mindset in general, and also the ability to view problems from a different perspective. In the case of matching markets in particular, in addition to the (often) constructive nature of computational approaches, there is also the historical fact that computer scientists have studied matching from many different perspectives, perhaps because matching markets have a very natural representation in the language of graphs.

While it is impossible to characterize the entire body of work on matching, some of the major themes of research have been:

- (1) The modeling of matching markets with individually rational agents [Becker 1983; Jovanovic 1979, e.g.]. This is classic positive economics, modeling observed phenomena in mathematical language, which in turn allows one to make predictions. For example, models of matching can predict phenomena like unraveling, where matchings start occurring earlier and earlier [Roth and Xing 1994].
- (2) Market design for centralized methods of engineering stable matchings; the

---

Authors' addresses: {eanshe1,sanmay}@cs.rpi.edu. This material is based partially on work supported by NSF grant CCF-0914782 and BSF grant 2008404.

key questions typically involve possibilities for manipulation and the design of truthful mechanisms [Roth and Sotomayor 1990; Immorlica and Mahdian 2005, e.g.]

- (3) Analysis of the convergence properties of decentralized mechanisms [Ackermann et al. 2008; Das and Kamenica 2005, e.g.].
- (4) Characterizing the outcomes of search processes involving friction in matching scenarios in terms of whether or not they result in assortative matchings [Shimer and Smith 2000, e.g.].

The majority of work on matching has focused on settings with ordinal preferences, and often the even more restrictive setting of binary (compatible/non-compatible) preferences. This makes sense on many levels. First, it is hard to put reasonable numbers on utilities in matching scenarios; these are often major decisions about life: whom to marry, which job to take, whether or not a kidney will be compatible. Second, in matching scenarios, stability is often the be-all and end-all of design goals, and doesn't need the overhead imposed by thinking about cardinal utilities. Third, the question of what an appropriate social utility function would be is difficult in many situations, since additive social utility may not be well-suited. However, each of these arguments breaks down at some point. As our understanding and our ability to engineer useful markets expands, consideration of cardinal utilities will inevitably become more and more important. In this note, we discuss several settings from our own research where cardinal utilities play a key role for research in matching.

## 2. MATCHING MODELS WITH CARDINAL UTILITIES

For the rest of this note, it should be assumed that we work in settings with nontransferable utilities, not allowing participants in matching markets to make payments to each other. Such settings arise in many cases where social norms or legal restrictions make it impossible to explicitly use monetary transfers to improve one's position in a matching. For example, it is unusual to directly pay someone in a dating/marriage market in order to win their affections, or to pay a competitor in order for them to stop pursuit of the same potential spouse (which is not to say that everyone's preferences are unaffected by monetary realities). It is also typically illegal to directly pay for organs. Even in many labor markets, employers do not compete on salary: medical residents typically get payed somewhat similarly.

### 2.1 Social Welfare

Market design is essentially a normative discipline: given some allocative desiderata, how can we design a system that achieves them when we have no control over participant behavior, but we may be able to influence their incentives in certain ways? The tendency of the field has been to focus on objectives that are relatively uncontroversial: in particular, stability is often the goal [Roth and Sotomayor 1990; Roth and Xing 1994; Immorlica and Mahdian 2005, *inter alia*]. This is natural for many of the reasons discussed above. In particular, it is hard to come up with meaningful utilities in many matching situations, it is unclear if simple measures of social welfare like additive utility are the best to use, and, perhaps most importantly, it is not clear if there are any reasonable ways to incentivize agents to

participate in unstable matchings, even if they would be socially preferable.

There definitely exist matching situations where we *care* about social welfare. Consider two examples. Pair programming is a central practice of the software engineering methodology known as Extreme Programming [Dawande et al. 2008]. If a software firm uses this methodology, the utility of a matching is a function of the productivity of a pair of programmers working together. Maximizing additive social utility in this case would be optimal for the firm, but individual programmers may not be happy in this situation.

While kidney *exchange* in the United States is typically analyzed in a 0-1 preference framework based on compatibility [Roth et al. 2004; 2005], both the European school of thought on kidney exchange and the literature on cadaveric kidney transplants evaluate success in terms of the quality of the match produced [Segev et al. 2005; Held et al. 1994]. The operations research literature on cadaveric kidney donation is often explicit in attempting to optimize measures like quality-adjusted life years for all recipients [Su and Zenios 2005]. Recent literature on kidney exchange coming from the computer science community is beginning to consider optimizing overall social welfare instead of the number of compatible matches [Abraham et al. 2007; Awasthi and Sandholm 2009]. It is not clear if the resistance to analyzing kidney exchange in terms of utilities was because (1) we really could not improve outcomes, (2) analyzing matching in terms of cardinal utilities is hard, or, (3) the ethical questions this could raise are extraordinarily difficult to deal with.<sup>1</sup> With respect to (3), we should note that similarly difficult questions have been engaged in cadaveric kidney donation [Zenios et al. 2001], and in many other areas of medical ethics.

Some of our recent work focuses on extending our understanding of matching problems in situations where we are concerned with social welfare in terms of utility, instead of just stability and choice among stable outcomes (we work with simple additive measures of social welfare). In “Anarchy, Stability, and Utopia: Creating Better Matchings,” Anshelevich, Das, and Naamad [2009] address two major questions of relevance: (1) How bad are stable matchings when compared with socially optimal ones? (2) How can we incentivize agents to participate in matchings that are socially desirable?

We study, both theoretically and in simulation, the effects of different network structures and utility distributions on the price of anarchy: the worst-case ratio of social utilities achieved by stable and optimal matchings respectively. While theoretically, the price of anarchy can be quite bad, we find that under some natural preference structures and simple distributions of utilities, stable matchings attain close to the optimal social welfare (generally above 90%). We characterize some situations where the loss in social welfare can be more substantial, and then study a potential means of incentivizing good stable matchings. We consider *approximate stability*, which corresponds to the addition of a switching cost to the mechanism, so that an agent would have to pay in order to deviate from the current matching. We show theoretical bounds on the price of stability, the ratio of the social utility of the best (approximately) stable matching to the socially optimal matching. Our proof is constructive: we present an algorithm that constructs an approximately

<sup>1</sup>In reality, of course, it is probably a combination of all the above.

stable matching that achieves a good ratio. A matchmaker could use this algorithm to suggest a near-optimal approximately stable matching to all agents. If they participated in the mechanism under the knowledge that they would have to pay a switching cost to deviate from the suggested matching, this would lead to matchings that were close to socially optimal. Additionally, simulation results show that the algorithm typically achieves even better performance (in terms of closeness to social optimality) than guaranteed by the theoretical bounds.

Our paper also verifies another insight related to cardinal utilities: as preferences in a matching “society” become more diverse, the social utility of stable matchings increases, even though the price of anarchy may also increase. The intuition is that a matchmaker is more likely to be able to make everyone happy when they don’t all want the same things, and thus diversity can greatly increase overall societal happiness.

## 2.2 Preference Learning

The matching literature typically assumes that agents know their preferences in advance of entering the matching mechanism. In many contexts this assumption is overly restrictive. For example, employers often employ elaborate, multi-round interview processes with potential employees in order to learn their preferences over the employees. In dating situations, people learn about their compatibility by going out on dates, or finding out more about each other on dating websites or speed-dating situations. While this can again be analyzed in terms of ordinal rankings of options, much of the framework of learning is more readily applied when agents are learning about the actual intensity of their preferences rather than just the ranking (compare also analysis of the classical secretary problem with the “search” problem [Gilbert and Mosteller 1966]).

Das and Kamenica [2005] explore the consequences of agents having to learn their preferences sequentially through interactions with each other, in the context of a dating game. They find that the actual matching mechanism used can significantly affect agents’ ability to converge to a stable matching. In particular, there is an interesting interaction between learning the actual value of going out with a particular person and learning the probability that a match will be successfully made, depending on the mechanism being used. For example, there may be a considerable risk in asking out the “most preferred” person when the periodwise matching mechanism involves take-it-or-leave-it offers, leading to worse learning and less stable outcomes. But if the periodwise mechanism involves submitting preference lists to a centralized Gale-Shapley type matchmaker, agents may be better off submitting their true preferences.

This was one of the early papers to look at having to learn preferences in the context of matching markets, but there has since been a significant amount of work that views preference learning in terms of information acquisition costs. In particular, work on costly interviewing in matching markets (see [Lee and Schwarz 2009] and references therein) looks at the effects on the whole matching system of firms making rational decisions about whom to interview. Again, this is a situation where the actual costs of interviewing have important consequences (just consider the budget for interviewing faculty candidates in your department!).

### 2.3 Effort Allocation in Social Networks

Cardinal utilities and social welfare arise in more general matching problems as well, such as many-to-many matching and fractional matching. For example, if we think of a link between two people in a social network as representing a joint project or a friendship, then a fractional matching could tell us the strength of this link. Although this leads to more complex interactions and matching structures than traditional matching problems, a lot of the basic ideas and techniques used for “integral” matching can be used to analyze and form good mechanisms for “fractional” matching.

In “Contribution Games in Social Networks,” Anshelevich and Hoefer [2010] propose and study a framework for *network contribution games* (NCGs) that model collaboration and coordination. Each agent in the network has a budget of effort that it can contribute to different collaborative projects. Budgets and contributions are non-negative numbers that model the different ways and degrees by which actors can contribute to a bilateral project, e.g., by allocating time, money, and personal energy to maintaining a relationship or a collaboration, the development of a new product, or the installation or standardization of a new technology. Depending on the contribution of the involved agents a project will flourish or drown, and success is measured using a cardinal reward function. Every agent is trying to maximize the reward from all projects that it is involved in. Thus, it may be in the best interest of an agent to form several different partnerships/matches of different strengths. A framework like this requires cardinal utility functions, since every agent can be involved in many projects at once, and the quality of each project changes depending on how much effort both participants put into it. NCGs encompass the well-studied *minimum effort coordination game* [van Huyck et al. 1990], where the outcome depends solely on the minimum contribution of any agent, as a special case. It is known that the Nash equilibria in this game exhibit a simple structure; behavior in laboratory experiments, however, led to sometimes surprising patterns [Dufwenberg and Gneezy 2005; Bornstein et al. 2002; Fatas et al. 2006; Riechmann and Weimann 2008; Chaudhuri et al. 2008].

We characterize the existence, computational complexity, and quality of pairwise equilibrium [Jackson and Wolinsky 1996] in NCGs based on the types of reward functions involved. We demonstrate (constructively) the existence of a pairwise equilibrium when all reward functions are concave, and show that the price of anarchy is at most 2. For convex functions the same holds only under some special but very natural conditions. We also show tight bounds for approximate equilibria and convergence of dynamics in these games. The overall insight is that for “fractional” matching systems, good stable solutions exist and can often be found efficiently, even though these systems are much more complex than traditional 0-1 matching.

## 3. CONCLUSION

One of the contributions of the computational viewpoint to economics in general and market design in particular is the focus on quantification, ranging from quantifying the cost of computation to measures like the price of anarchy. It is natural to be somewhat skittish about doing the same in matching, for good reasons. For example, how can we assign utilities to outcomes when money is not involved? Is it

meaningful to use additive notions of social welfare common in computer science, or do we need to consider different welfare functions? Is it even possible to come up with strong results when cardinal utilities are involved? And perhaps there is a reticence to make suggestions for systems when political philosophy becomes involved – who decides which mechanism to use? However, not confronting these questions is not a solution. Market design has already made a tremendous impact in many fields. When ethical questions arise, they need to be brought to light and addressed by those with relevant expertise. Moreover, although measuring the utility of an outcome can be difficult, data can certainly be collected about the quality of matches, and experiments to measure the quality of mechanisms can be designed. Matching with cardinal utilities is a significant challenge, but we believe it is a promising one.

## REFERENCES

- 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*. ACM, 304.
- ACKERMANN, H., GOLDBERG, P., MIRROKNI, V., ROGLIN, H., AND VOCKING, B. 2008. Uncoordinated two-sided markets. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*. 256–263.
- ANSHELEVICH, E., DAS, S., AND NAAMAD, Y. 2009. Anarchy, stability, and utopia: Creating better matchings. In *Proceedings of the Second International Symposium on Algorithmic Game Theory*. 159–170.
- ANSHELEVICH, E. AND HOEFER, M. 2010. Contribution games in social networks. *Arxiv preprint arXiv:1004.1854*.
- 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*. Morgan Kaufmann Publishers Inc., 405–411.
- BECKER, G. 1983. A Treatise On The Family. *Family Process* 22, 1, 127–127.
- BORNSTEIN, G., GNEEZY, U., AND NAGEL, R. 2002. The effect of intergroup competition on group coordination: An experimental study. *Games and Economic Behavior* 41, 1, 1–25.
- CHAUDHURI, A., SCHOTTER, A., AND SOPHER, B. 2008. Talking ourselves to efficiency: Coordination in inter-generational minimum effort games with private, almost common and common knowledge of advice. *The Economic Journal* 119, 534, 91–122.
- CONITZER, V. 2010. Making decisions based on the preferences of multiple agents. *Commun. ACM* 53, 3, 84–94.
- DAS, S. AND KAMENICA, E. 2005. Two-sided bandits and the dating market. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*. Edinburgh, UK, 947–952.
- DAWANDE, M., KUMAR, S., MOOKERJEE, V., AND SRISKANDARAJAH, C. 2008. Maximum Commonality Problems: Applications and Analysis. *Management Science* 54, 1, 194.
- DUFWENBERG, M. AND GNEEZY, U. 2005. Gender and coordination. In *Experimental Business Research*, A. Rapoport and R. Zwick, Eds. Vol. 3. Kluwer, 253–262.
- FATAS, E., NEUGEBAUER, T., AND PEROTE, J. 2006. Within-team competition in the minimum effort coordination game. *Pacific Econ. Review* 11, 2, 247–266.
- GILBERT, J. AND MOSTELLER, F. 1966. Recognizing the maximum of a sequence. *Journal of the American Statistical Association* 61, 35–73.
- HELD, P., KAHAN, B., HUNSICKER, L., LISKA, D., WOLFE, R., PORT, F., GAYLIN, D., GARCIA, J., AGODOA, L., AND KRAKAUER, H. 1994. The impact of HLA mismatches on the survival of first cadaveric kidney transplants. *The New England journal of medicine* 331, 12, 765.
- IMMORLICA, N. AND MAHDIAN, M. 2005. Marriage, Honesty, and Stability. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*.
- ACM SIGecom Exchanges, Vol. 9, No. 1, June 2010.

- JACKSON, M. AND WOLINSKY, A. 1996. A strategic model of social and economic networks. *Journal of economic theory* 71, 1, 44–74.
- JOVANOVIĆ, B. 1979. Job Matching and the Theory of Turnover. *The Journal of Political Economy* 87, 5, 972.
- LEE, R. AND SCHWARZ, M. 2009. Interviewing in two-sided matching markets. *NBER Working Paper*.
- RIECHMANN, T. AND WEIMANN, J. 2008. Competition as a coordination device: Experimental evidence from a minimum effort coordination game. *European J. Political Economy* 24, 2, 437–454.
- ROTH, A., SÖNMEZ, T., AND ÜNVER, M. 2004. Kidney Exchange. *Quarterly Journal of Economics* 119, 2, 457–488.
- ROTH, A., SÖNMEZ, T., AND ÜNVER, M. 2005. A kidney exchange clearinghouse in New England. *American Economic Review* 95, 2, 376–380.
- ROTH, A. E. AND SOTOMAYOR, M. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monograph Series. Cambridge University Press, Cambridge, UK.
- ROTH, A. E. AND XING, X. 1994. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *The American Economic Review* 84, 4, 992–1044.
- SEGEV, D., GENTRY, S., WARREN, D., REEB, B., AND MONTGOMERY, R. 2005. Kidney paired donation and optimizing the use of live donor organs. *Jama* 293, 15, 1883.
- SHIMER, R. AND SMITH, L. 2000. Assortative matching and search. *Econometrica* 68, 2, 343–369.
- SU, X. AND ZENIOS, S. 2005. Patient choice in kidney allocation: A sequential stochastic assignment model. *Operations research* 53, 3, 443–455.
- VAN HUYCK, J., BATTALIO, R., AND BEIL, R. 1990. Tacit coordination games, strategic uncertainty and coordination failure. *American Econ. Review* 80, 1, 234–248.
- ZENIOS, S., WOODLE, E., AND FRIEDMAN ROSS, L. 2001. Primum non nocere: Avoiding harm to vulnerable wait list candidates in an indirect kidney exchange. *Transplantation* 72, 4, 648.