Research Statement — Elliot Anshelevich

April 20, 2017

My research interests center in the design and analysis of algorithms for systems composed of independent agents, and studying the properties resulting from their behavior. Much of my work concerns algorithmic game theory (as agents are often self-interested), the study of large networks of agents, and markets with many buyers and/or sellers. More recently, I have begun studying and developing algorithms for settings where only a limited amount of information is known about the input: this often occurs because the agents which report their state may not tell the truth, or may simply report only a small amount of information instead of everything that the algorithm designer would like to know. The emphasis of all my work is usually on algorithms and network properties with provable guarantees, and especially on approximation algorithms. In this research statement, I discuss each of these topics in detail, focusing especially on my recent work.

Ordinal Approximation Algorithms. (See [1] and presentation slides at [2].) In this work we design algorithms which are only given a limited amount of information (specifically *ordinal* information), and yet must compete with algorithms which know the "ground truth" numerical information. We call these *ordinal approximation algorithms*.

To illustrate when such constraints can arise, consider a simple matching scenario in which agents are located in some metric space, and the goal is to form a high-quality matching of the agents, where the distance between two agents represents how much those agents want to be matched together. However, instead of knowing the exact distances, all we know are the ordinal preferences: each agent tells us which would be their first choice to be matched with, which would be their second choice, etc. Based only on this information, how well can we compete with algorithms which know the true distances?

The same issues arise in social choice settings: as in spacial preference literature suppose that a set N of agents (voters) and a set of of alternatives (candidates) A correspond to points in some metric space. The distance from an agent i to alternative a represents the cost that agent i assigns to alternative $a \in A$ occurring,[1] and the goal is to find the optimum alternative (for example, the one which minimizes the total distance to all voters). If we knew the exact distances this would be easy to do, but once again this is complicated by the fact that while it may often be reasonable to assume some underlying utility structure on the preferences of the agents, it is usually *un*reasonable to assume that we know these numerical values exactly; it is much more likely that we only know the preference ordering over the candidates for each voter. In other words, just as in the matching setting, we only know the *ordinal* preferences for each agent i over alternatives in A; we do not know the actual numerical values and thus the *strengths* of agent i's preferences.

Perhaps surprisingly, we were able to show that in many settings ordinal information is enough to form algorithms which perform almost as well as ones with access to the full numerical information. In [3] we showed that several classic social choice mechanisms are good ordinal approximation algorithms as well. For example, the Copeland mechanism always returns an alternative which is *guaranteed* to be within a factor of 5 in quality as compared to the optimum alternative, *no matter what* the underlying numerical truth is; since we also proved a lower bound of 3 for the approximation factor of any deterministic algorithm which uses only ordinal knowledge, this is somewhat close to the best possible. In follow-up work [4,5], we designed randomized algorithms with much better approximation factors. For the matching setting described above, in [6] we designed an ordinal 1.6-approximation algorithm for the maximum-weight metric matching problem; this algorithm uses a careful mix or greedy and random matchings in order to form matchings with provably high quality. Moreover, in [8]

[1] See e.g., [1,3,7] for discussion and history of this utilitarian view of social choice.

we specifically focused on designing *truthful* algorithms: ones which result in a good matching while using only ordinal information, and in addition do not provide any incentive for the agents to lie about their preferences. More generally, other graph and clustering problems allow good approximations based only on ordinal information; for example in [9] we give such approximations for Max k-sum, Densest k-subgraph, and Maximum Traveling Salesman problems.

Our results so far show that in many settings, limited ordinal information is enough to form algorithms which perform almost as well as ones with access to the full numerical information. Thus, if obtaining full numerical information is too costly, it may be better to collect ordinal information and apply approximation algorithms. More generally, the same questions can be asked for other types of limited information: how much information about the true input is enough to produce good results? I believe this is a promising area of research, and am excited to continue pursuing it.

Agents in Networks – Network Formation Games and Group Formation. (See presentation slides at [10].) I continue to study various aspects of networks which are created or administered by selfish independent agents, following up on my work on network formation games from [11,12]. (This work popularized now-standard concepts such as potential games and price of stability, and has resulted in a huge amount of followup work; see for example [13, Chapters 17,19] and references therein.) In recent years, I have mostly been interested in network contribution games, where agents not only form local links to their neighbors, but also determine the strength of these links. More specifically, each node/agent has a budget of effort that it can allocate to different incident edges representing its friendships, relationships, collaborations, etc. The amount of effort allocated to a link by its endpoints determines the strength of this relationship, as well as the happiness of the participants with this relationship. Together with Martin Hoefer, I was able to show the existence of various types of coalitional equilibria in such games, as well as to quantify the price of anarchy, and prove various convergence properties of agent interactions [14]. These games have a close relationship to the classic problem of stable matching with cardinal preferences, which arises in such applications as matching medical residents to hospitals, online dating, and kidney exchange [15]. I was able to use similar techniques to provide price of anarchy bounds for both exact and approximate stable matching [16,17]. I was especially interested in seeing how externalities affect the properties of networks formed in this manner: for example when agents are somewhat altruistic or care about the status of their friends in the network [18,19], or when agents receive benefit not just from their immediate friends, but also from friends-of-friends (e.g., I may decide to become friends with person X not because of X themselves, but because of X's circle of friends) [20]. Further questions in this area include determining effective ways to influence the agents of these games, and considering a combination of agent incentives (including average distance to other nodes, betweenness, etc.), as well as applying these techniques to specific network settings (see [30] below).

Games in which agents form groups are very related to network formation games, both in the techniques which can be used and in the questions commonly asked about them [21]. In addition to work on price of anarchy using semi-smoothness [22], I was recently able to analyze the quality of various equilibrium concepts when player utilities from joining a group are non-monotone in the size of that group [23], as well as show how to construct high-quality approximate equilibria by giving the players small amounts of incentives [24, 25]. The latter is especially nice since for settings where exact equilibria do not exist, we were able to create good stable solutions by using a small injection of incentives.

Agents in Networks – Autonomous Systems and the Internet. The Internet is composed of tens of thousands of sub-networks called *Autonomous Systems* (AS), each under a single administrative authority with its own distinct goals in controlling the traffic entering and leaving its network. The complex system of business relationships and routing agreements between various Internet entities (e.g., Autonomous Systems, Internet Exchange Points, enterprize networks, residential customers) is at the heart of Internet connectivity. In addition to looking at abstract network formation games, I also consider specifically the games played by Autonomous Systems and ISP's when forming contracts in the Internet [26, 27], as well as when pricing those contracts [28, 29], and when purchasing connections from Internet Exchange Points (IXP's) [30]. In all of the above contexts, I was able to show that the stable solutions of the game have good quality, ways to influence the agents in the network in order to form good global solutions, or pricing schemes which yield high-quality solutions.

Item Pricing for Large Markets and Multiple Objectives. Perhaps the most classical setting with many independent agents is that of economic markets. Together with PhD student Shreyas Sekar, I have designed simple non-discriminatory item pricing algorithms for various types of markets [31–34]. In such algorithms, a seller chooses prices for each good that they are selling, and the buyers receive their most preferred bundle of goods at these prices. While this is a fundamental economic setting that has received a lot of attention, much remains unknown about forming efficient pricing algorithms which provably result in good outcomes for both the buyers and the seller (or sellers). Some of the things that set our work apart from most of the literature are as follows. (1) *Large Markets:* We especially focus on markets in which each buyer can be assumed to be infinitesimal as compared to the overall market size. No single buyer makes much difference to the market, and such markets often have very different properties from "small" markets with only a few buyers. (2) *Optimizing Multiple Simultaneous Objectives:* We have had significant success designing pricing schemes which form solutions *simultaneously* optimizing several objectives, especially both Revenue (payments received by the seller) and Social Welfare (total utility of the agents). (3) *Production Costs:* Many real markets include *production,* instead of having a fixed supply. Unfortunately, despite the fact that production costs often require very different techniques from limited supply, and in fact generalize both limited and unlimited supply, markets with non-linear production costs have received too little attention in the item pricing literature.

For example, in [31] we studied envy-free pricing in large markets with unit-demand valuations: each buyer wants a single item from a set of satisfactory items, but both these sets and the values that the buyers assign to obtaining an item can differ for different buyers. This problem is known to be notoriously hard (no approximation algorithm better than $O(\log n)$ is possible in general), and has received a lot of attention during the last decade. In [31], we gave a new pricing scheme such that for large markets with the demand curve obeying the Monotone Hazard Rate condition, this scheme is guaranteed to obtain a 1.88-approximation for the maximum revenue. We were later able to greatly generalize these results for other demand types beyond Monotone Hazard Rate, as well as to multi-minded (not just unit-demand) buyers, and provide simple pricing schemes which guarantee both good revenue (e.g., e-close to optimum) for the seller and good social welfare (e.g., 2-close to optimum) for the buyers simultaneously [31–33].

More generally, we are especially interested in black-box reductions. For example, these include general procedures converting pricing schemes with good revenue or good welfare into pricing schemes which guarantee good revenue *and* welfare at the same time. We were able to provide such reductions for a variety of settings [32, 33], in which a pricing scheme approximating maximum possible revenue is converted into a scheme with slightly less revenue, but is guaranteed to give a constant approximation to the optimum social welfare for the buyers. Most results for these problems have been concerned only with revenue: our results show that in this setting, it is possible to also achieve high social welfare without sacrificing much revenue. Moreover, we were able use our techniques to create new approximation algorithms for sequential buyers with XoS valuations; ours is the first known item pricing algorithm with a polylogarithmic approximation for revenue for such general classes of valuations [33]. Another type of black-box reduction which we have worked on are techniques to convert pricing schemes which are known to work well in the presence of limited supply to ones which have good properties in the presence of production costs. We believe that the ideas involved in our techniques can be used to design a variety of new efficient pricing algorithms; this is something that we are actively working on. We also continue to work on quantifying the tradeoffs between revenue and welfare, especially for large markets, and for markets with a network structure [34].

References

[1] Elliot Anshelevich. Ordinal Approximation in Matching and Social Choice. *Newsletter of the ACM Special Interest Group on E-commerce (SIGecom Exchanges),* Volume 15.1, July 2016.

[2] See presentation slides at `http://www.cs.rpi.edu/~eanshel/slides/Distortion.pptx` and `http://www.cs.rpi.edu/~eanshel/slides/OrdinalMatching.pdf`

[3] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. Approximating Optimal Social Choice under Metric Preferences. *Proc. of 29th Conference on Artificial Intelligence (AAAI 2015).*

[4] Elliot Anshelevich and John Postl. Randomized Social Choice Functions Under Metric Preferences. *Proc. of 25th International Joint Conference on Artificial Intelligence (IJCAI 2016).*

[5] Stephen Gross, Elliot Anshelevich, and Lirong Xia. Vote Until Two of You Agree: Mechanisms with Small Distortion and Sample Complexity. *Proc. of 31st Conference on Artificial Intelligence (AAAI 2017).*

[6] Elliot Anshelevich and Shreyas Sekar. Blind, Greedy, and Random: Algorithms for Matching and Clustering using only Ordinal Information. *Proc. of 30th Conference on Artificial Intelligence (AAAI 2016)*

[7] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 197–214. ACM, 2012.

[8] Elliot Anshelevich and Shreyas Sekar. Truthful Mechanisms for Matching and Clustering in an Ordinal World. *Proc of 12th Conference on Web and Internet Economics (WINE 2016).*

[9] Elliot Anshelevich and Shreyas Sekar. Blind, Greedy, and Random: Ordinal Approximation Algorithms for Graph Problems. *arXiv:1512.05504*

[10] See presentation slides at `http://www.cs.rpi.edu/~eanshel/slides/ContributionGames.ppt` and `http://www.cs.rpi.edu/~eanshel/slides/Friendship.pdf`

[11] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-Optimal Network Design with Selfish Agents. In *Theory of Computing,* Volume 4 (2008), pp. 77-109. Conference version appeared in *Proc. 35th ACM Symposium on Theory of Computing* (STOC 2003).

[12] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *SIAM Journal on Computing,* Volume 38, Issue 4 (November 2008), pp. 1602-1623. Conference version appeared in *Proc. 45th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2004).

[13] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani (eds.), *Algorithmic Game Theory,* Cambridge University Press.

[14] Elliot Anshelevich and Martin Hoefer. Contribution Games in Social Networks. *Algorithmica,* Volume 63, Issue 1-2 (June 2012), pages 51-90. Conference version appeared in *ESA 2010.*

[15] Elliot Anshelevich, Meenal Chhabra, Sanmay Das, and Matthew Gerrior. On the Social Welfare of Mechanisms for Repeated Batch Matching. *Proc. 27th Conference on Artificial Intelligence (AAAI 2013).*

[16] Elliot Anshelevich and Sanmay Das. Matching, Cardinal Utility, and Social Welfare. *ACM SIGecom Exchanges,* Volume 9.1, June 2010.

[17] Elliot Anshelevich, Sanmay Das and Yonatan Naamad. Anarchy, Stability, and Utopia: Creating Better Matchings. *Journal of Autonomous Agents and Multi-Agent Systems,* Volume 26, Issue 1 (January 2013), pages 120-140. Conference version appeared in *SAGT 2009*.

[18] Elliot Anshelevich, Onkar Bhardwaj, and Martin Hoefer. Stable Matching with Network Externalities. *Algorithmica,* to appear. Conference version appeared in ESA 2013.

[19] Elliot Anshelevich, Onkar Bhardwaj, and Martin Hoefer. Friendship and Stable Matching. *Proc. 21st European Symposium on Algorithms (ESA 2013).*

[20] Elliot Anshelevich, Onkar Bhardwaj, and Michael Usher. Friend of My Friend: Network Formation with Two-Hop Benefit. *Theory of Computing Systems,* Volume 57, Issue 3 (2015), pages 711-752. Conference version appeared in *SAGT 2013*.

[21] Elliot Anshelevich and Bugra Caskurlu. Exact and Approximate Equilibria for Optimal Group Network Formation. *Theoretical Computer Science,* Volume 412, Issue 39 (September 2011), pp. 5298-5314. Conference version appeared in *ESA 2009*.

[22] Elliot Anshelevich, John Postl, and Tom Wexler. Assignment Games with Conflicts: Robust Price of Anarchy and Convergence Results via Semi-Smoothness. *Theory of Computing Systems,* Volume 59, Issue 3 (October 2016), Pages 440–475.

[23] Elliot Anshelevich and John Postl. Profit Sharing with Thresholds and Non-monotone Player Utilities. *Theory of Computing Systems,* Volume 59, Issue 4 (2016), Pages 563–580. Conference version appeared in *SAGT 2014*.

[24] Elliot Anshelevich and Shreyas Sekar. Computing Stable Coalitions: Approximation Algorithms for Reward Sharing. *Proc of 11th Conference on Web and Internet Economics (WINE 2015).*

[25] Elliot Anshelevich and Shreyas Sekar. Approximate Equilibrium and Incentivizing Social Coordination. *Proc. of 28th Conference on Artificial Intelligence (AAAI 2014).*

[26] Elliot Anshelevich, Bruce Shepherd, and Gordon Wilfong. Strategic Network Formation through Peering and Service Agreements. *Games and Economic Behavior,* Volume 73, Issue 1, September 2011, Pages 17-38. Conference version appeared in *Proc. 47th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2006).

[27] Elliot Anshelevich and Gordon Wilfong. Network Formation and Routing by Strategic Agents using Local Contracts. In *Proc. 4th International Workshop On Internet And Network Economics* (WINE 2008).

[28] Onkar Bhardwaj, Elliot Anshelevich, and Koushik Kar. Coalitionally Stable Pricing Schemes for Inter-domain Forwarding. *Computer Networks,* Volume 97, Issue C (March 2016), Pages 128–146.

[29] Elliot Anshelevich, Ameya Hate, and Koushik Kar. Strategic Pricing in Next-hop Routing with Elastic Demands. *Theory of Computing Systems,* Volume 54, Issue 3 (2014), pages 407-430. Conference version appeared in *SAGT 2011.*

[30] Elliot Anshelevich, Onkar Bhardwaj, and Koushik Kar. Strategic Network Formation through Intermediaries. *Proc of 24th International Joint Conference on Artificial Intelligence (IJCAI 2015).*

[31] Elliot Anshelevich, Koushik Kar, and Shreyas Sekar. Envy-Free Pricing in Large Markets: Approximating Revenue and Welfare. *Proc of 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015).*

[32] Elliot Anshelevich, Koushik Kar, and Shreyas Sekar. Pricing to Maximize Revenue and Welfare Simultaneously in Large Markets. *Proc of 12th Conference on Web and Internet Economics (WINE 2016).*

[33] Elliot Anshelevich and Shreyas Sekar. Price Doubling and Item Halving: Robust Revenue Guarantees for Item Pricing. *Proc of 18th ACM Conference on Economics and Computation (EC 2017).*

[34] Elliot Anshelevich and Shreyas Sekar. Price Competition in Networked Markets: How do monopolies impact social welfare? *Proc of 11th Conference on Web and Internet Economics (WINE 2015).*