

The Evolution of Core Stability in Decentralised Matching Markets

Heinrich H. Nax & H. Peyton Young
October 2011

Extended abstract

Electronic technology has created new forms of markets that involve large numbers of agents who interact in real time at virtually no cost. These markets are larger, more decentralised, and more differentiated than traditional market institutions. Interactions are driven by repeated online participation over extended periods of time without public announcements of bids, offers, or realized prices. Even after many encounters, agents may learn little or nothing about the others' preferences and past actions. Our goal is to construct a dynamic model that incorporates these features, and to explore its convergence and welfare properties. We see this as a first step toward developing a better understanding of how such markets operate, and how they might be more effectively designed.

We shall be particularly interested in bilateral markets where agents on each side of the market submit “demands” and are matched provided that their demands are mutually compatible. Examples include online platforms for matching buyers and sellers of goods, for matching workers and firms, and for matching hotels with hotel clients. These matching markets have traditionally been analyzed using game-theoretic methods (Gale & Shapley [1962], Shapley & Shubik [1972], Roth & Sotomayor [1990]). In much of this literature, however, it is assumed that agents submit preference menus to a central authority, which then employs a suitably designed algorithm to match them. The model we propose here is different in character: agents submit bids that are conditional on the characteristics of those with whom they are matched, and the only role of the central authority (matchmaker) is to create a compatible (not necessarily optimal) set of matches at each point in time. There is no presumption that agents (or the matchmaker) know anything about the preferences of the others, or that they can deduce such information from prior rounds. Instead agents employ a trial-and-error learning model: matched agents occasionally probe higher demands to see if they can “get away with it”, while unmatched agents lower their demands incrementally in the hope of attracting partners. We show that over time this process leads to core outcomes. Moreover, within each matched pair there is a probabilistic bias towards an equitable division of the surplus, that is, the process is more likely to yield outcomes in which the surplus is split as evenly as possible subject to the core constraints than it is to produce highly unequal divisions.

There are three basic elements to our learning procedure. First, behaviour aims to satisfy momentary personal aspirations based on trial and error: agents occasionally explore if alternative actions could lead to higher payoffs and abandon actions that result in worse positions quickly. Second, upward and downward adjustments are made locally and in small increments. Third, demands are sticky: players are more likely to adjust demands

downwards the higher their aspiration.

To the best of our knowledge there is no previous work on dynamic learning models applied to decentralized matching markets. However, there is a sizeable literature on matching algorithms that grows out of the seminal paper by Gale & Shapley [1962]. In this approach agents submit preferences for being matched with agents on the other side of the market, and a central clearing house matches them in a way that yields a core outcome (provided that the reports are truthful).¹ These algorithms have been successfully applied in situations where agents engage in a formal application process, such as students applying for admission to universities, or doctors to hospital residencies.²

In the present paper, by contrast, we consider situations where the market is fluid and decentralized. Agents are matched and rematched over time, and the information they submit takes the form of prices rather than preferences. We shall show that even when agents have minimal amounts of information and use very simple price adjustment rules, the market evolves towards core outcomes.

This result fits into a growing literature showing how cooperative game solutions can be understood as outcomes of a dynamic learning process (Agastya [1997], [1999]; Arnold & Schwalbe [2002], Rozen [2010a], [2010b]; Newton [2010], [2011]). We shall briefly outline Newton's approach here; the others are similar in spirit. In each period a player is selected at random and he or she demands a share of the surplus from some targeted coalition of players. He chooses a demand that amounts to a best reply to the expected demands of the others in the coalition, where his expectations are based on a random sample of the other players' past demands. In fact he chooses a best reply with probability close to one, but with small probability he may make some other demand. This noisy best response process leads to a Markov chain whose ergodic distribution can be characterized using the theory of large deviations. Newton shows that, subject to various regularity conditions, this process converges to a core allocation in games that possess a nonempty interior core. Moreover, under a suitable specification of the error distribution, the stochastically stable outcomes maximize a Rawlsian (maximin) welfare function subject to the core constraints.

How would this type of process operate in the matching markets described earlier? Each agent would know the identities of the agents on the other side of the market and would form expectations about their current demands from a random sample of their previous demands. This presumes much more information than would typically be available in an online matching market. Moreover, to obtain core convergence Newton needs to assume that the cooperative game has a nonempty interior core. Unfortunately this condition does not hold for matching games, because the value of the grand coalition is simply the disjoint sum of the values of an optimal set of matched pairs, hence the core constraints will hold as equalities for some of the subcoalitions as well as for the grand coalition.

The approach we take requires much less information on the part of the agents, and it

¹See Demange & Gale [1985], Demange, Gale & Sotomayor [1986] for examples of such clearing mechanisms. See Shimer [2005], Elliott [2010], [2011] for models with costly search.

²See, for example, Roth [1984] for a discussion of the medical resident market in the US and the National Residency Matching Program.

employs a naïve adjustment rule that is akin to reinforcement learning (Bush & Mosteller [1955]). Specifically we shall assume that: (1) players do not know anything about the identities or past behaviour of other participants in the market, and (2) their learning procedures are completely uncoupled, that is, they are a function only of their own realized payoffs. In particular, players do not attempt to choose best replies to the others' strategies; they simply experiment to see whether they might be able to do better. Rules of this type have a long history in the psychology literature (Thorndike [1898], Hoppe [1931], Estes [1950], Herrnstein [1961]). Furthermore it has recently been shown that there are families of such rules that lead to equilibrium behaviour in generic noncooperative games (Karandikar, Mookherjee, Ray & Vega-Redondo [1998], Foster & Young [2006], Germano & Lugosi [2007], Marden, Young, Arslan & Shamma [2009], Young [2009], Pradelski & Young [2010]).

This framework has not previously been used to study learning dynamics in cooperative games.³ It seems especially well-suited to modelling behaviour in large decentralized markets, where agents have little information about the overall game and the identity of the other market participants. Here we shall restrict our attention to the analysis of learning dynamics in matching (assignment) games, which constitute a particularly important class in practice.

References

- M. Agastya, "Adaptive Play in Multiplayer Bargaining Situations," *Review of Economic Studies* 64, 411-26, 1997.
- M. Agastya, "Perturbed Adaptive Dynamics in Coalition Form Games," *Journal of Economic Theory* 89, 207-233, 1999.
- T. Arnold & U. Schwalbe, "Dynamic coalition formation and the core," *Journal of Economic Behavior and Organization* 49, 363-380, 2002.
- R. Bush & F. Mosteller, *Stochastic Models of Learning*, Wiley, 1955.
- G. Demange & D. Gale, "The strategy of two-sided matching markets," *Econometrica* 53, 873-988, 1985.
- G. Demange, D. Gale & M. Sotomayor, "Multi-item auctions," *Journal of Political Economics* 94, 863-872, 1986.
- M. L. Elliott, "Inefficiencies in networked markets," *working paper*, Stanford University, 2010.
- M. L. Elliott, "Search with multilateral bargaining," *working paper*, Stanford University, 2011.
- W. Estes, "Towards a statistical theory of learning," *Psychological Review* 57, 94-107, 1950.

³Sandholm [2008] reviews many of the previous applications.

- D. Foster & H. P. Young, "Regret testing: Learning to play Nash equilibrium without knowing you have an opponent," *Theoretical Economics* 1, 341-367, 2006.
- D. Gale & L. S. Shapley, "College admissions and the stability of marriage," *American Mathematical Monthly* 69, 9-15, 1962.
- F. Germano & G. Lugosi, "Global Nash convergence of Foster and Young's regret testing," *Games and Economic Behavior* 60, 135-154, 2007.
- H. Heckhausen, "Motivationsanalyse der Anspruchsniveau-Setzung," *Psychologische Forschung* 25, 118-154, 1955.
- R. J. Herrnstein, "Relative and absolute strength of response as a function of frequency of reinforcement," *Journal of Experimental Analysis of Behavior* 4, 267-272, 1961.
- F. Hoppe, "Erfolg und Mißerfolg," *Psychologische Forschung* 14, 1-62, 1931.
- R. Karandikar, D. Mookherjee, D. Ray & F. Vega-Redondo, "Evolving Aspirations and Cooperation," *Journal of Economic Theory* 80, 292-331, 1998.
- J. R. Marden, H. P. Young, G. Arslan, J. Shamma, "Payoff-based dynamics for multi-player weakly acyclic games," *SIAM Journal on Control and Optimization* 48, special issue on "Control and Optimization in Cooperative Networks", 373-396, 2009.
- J. Newton, "Non-cooperative convergence to the core in Nash demand games without random errors or convexity assumptions," Ph.D. thesis, University of Cambridge, 2010.
- J. Newton, "Recontracting and stochastic stability in cooperative games," mimeo, University of Cambridge, 2011.
- B. Pradelski & H. P. Young, "Learning Efficient Nash Equilibria in Distributed Systems," *Department of Economics Working Paper 480*, University of Oxford, 2010.
- H. Raiffa, "Arbitration schemes for generalized two-person games," in *Contributions to the Theory of Games*, Vol. 2, H. Kuhn, A. Tucker & M. Dresher (eds.), Princeton University Press, 361-387, 1953.
- A. E. Roth, "The Evolution of the Labor Markets for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy* 92, 991-1016, 1984.
- A. E. Roth, T. Sönmez & U. Ünver, "Pairwise kidney exchange," *Journal of Economic Theory* 125, 151-188, 2005.
- A. E. Roth & M. Sotomayor, *Two-Sided Matching: A Study in Game Theoretic Modeling and Analysis*, Cambridge University Press, 1990.
- A. E. Roth & M. Sotomayor, "Two-sided matching," in *Handbook of Game Theory with Economic Applications*, Volume 1, R. Aumann & S. Hart (eds.), 485-541, 1992.
- K. Rozen, "Conflict Leads to Cooperation in Nash Bargaining," *mimeo*, Yale University, 2010a.
- K. Rozen, "Conflict Leads to Cooperation in Nash Bargaining: Supplemental Result on Evolutionary Dynamics," *web appendix*, Yale University, 2010b.

- T. Sandholm, "Computing in Mechanism Design," *New Palgrave Dictionary of Economics*, 2008.
- H. Sauermann & R. Selten, "Anspruchsanpassungstheorie der Unternehmung," *Zeitschrift für die Gesamte Staatswissenschaft* 118, 577-597, 1962.
- L. S. Shapley & M. Shubik, "The Assignment Game I: The Core," *International Journal of Game Theory* 1, 111-130, 1972.
- M. Sotomayor, "Some further remark on the core structure of the assignment game," *Mathematical Social Sciences* 46, 261-265, 2003.
- R. Tietz, "An Experimental Analysis of Wage Bargaining Behavior," *Zeitschrift für die gesamte Staatswissenschaft* 131, 44-91, 1975.
- R. Tietz & O. Bartos, "Balancing of aspiration levels as fairness principle in negotiations," *Lecture Notes in Economics and Mathematical Systems* 213, R. Tietz (ed.), 52-66, 1983.
- R. Tietz, W. Daus, J. Lautsch & P. Lotz, "Semi-normative properties of bounded rational bargaining theories," *Experimental Economics* 314, R. Tietz, W. Albers & R. Selten (eds.), 142-159, 1988.
- R. Tietz & H. Weber, "On the nature of the bargaining process in the Kresko-game," in *Contributions to experimental economics*, Vol. 3, H. Sauermann (ed.), 305-334, 1972.
- R. Tietz & H. Weber, "Decision behavior in multi-variable negotiations," in *Contributions to experimental economics*, Vol. 7, H. Sauermann (ed.), 60-87, 1978.
- R. Tietz, H. Weber, U. Vidmajer & C. Wentzel, "On Aspiration Forming Behavior in Repetitive Negotiations," in *Contributions to experimental economics*, Vol. 7, H. Sauermann (ed.), 88-102, 1978.
- E. Thorndike, "Animal Intelligence: An Experimental Study of the Associative Processes in Animals," *Psychological Review* 8, 1898.
- H. Weber, "On the theory of Adaptation of Aspiration Levels in a Bilateral Decision Setting," *Zeitschrift für die gesamte Staatswissenschaft* 132, 582-591, 1976.
- H. P. Young, "Learning by trial and error," *Games and Economic Behavior* 65, 626-643, 2009.
- F. Zeuthen, *Problems of Monopoly and Economic Warfare*, Routledge & Kegan Paul, 1930.