
Towards a theoretical framework for algorithms of market dynamics

Broadly speaking, I am interested in studying market dynamics from an algorithmic perspective. This is a problem at the intersection of machine learning and theoretical economics, and has significant potential for cutting-edge research as well as applications and connections to related areas in both economics and computer science.

To state the research problem more precisely, I will need to first develop some background. I will first describe the significant literature there exists on “market algorithms” (algorithms on markets, whether for equilibrium computation or optimization), and more specifically on my problem of interest: modeling “real” market dynamics in terms of these algorithms.

The closest line of research to this problem I am aware of is probably Kanoria et al [1–3], who demonstrate an equivalence between a simple model of market dynamics and a belief propagation algorithm. I am more generally interested in creating a unified theoretical framework for market algorithms, which predicts them as the dynamics on a market given some agent-level assumptions.

The most immediate implications of my research would be in agent-based computational economics; more computer science oriented implications are described in subsequent sections.

Background.

There are several different methodologies/perspectives to studying market dynamics and equilibrium computation. I briefly review them below, and highlight the differences between them and their limitations.

Algorithmic game theory. There are numerous algorithms studied in mechanism design that amount to computing the equilibrium of a game. We refer to [4–6] for reviews of these mechanisms – of particular interest is the case when agents are organized into a graph [7] and algorithms based on double auctions [8, 9], such as on a supply chain [10, 11].

These algorithms are typically considered over the setting of an artificially-designed market, such as an online supply chain, and the dynamics of the game are thus “hard-coded” into the system, as it is an engineering problem.

Multi-agent learning. Of particular interest are repeated games, in which the agents in the system must learn to optimize their actions to their environment. There is extensive research on multi-agent reinforcement learning [12, 13], as well as other algorithms for game dynamics like no-regret learning [14]. These concern themselves with optimization by explicitly strategic agents; the agents’ internal algorithms may have varying degrees of computational efficiency.

Message-passing algorithms. Message-passing algorithms are commonly used in multi-agent coordination; a classic example is Max-Sum [15].

These algorithms are typically for optimization rather than equilibrium computation, and require full collaboration. If we are interested in predicting market dynamics from individual-level assumptions about agents’ interests, then this sort of algorithm may entail less realistic assumptions, or be less promising. Furthermore, Max-Sum is computationally expensive to scale [16], which raises doubts as to how realistic it is as a model of economic computation in real markets.

However, the perspective that market dynamics is fundamentally a matter of information propagation is worth entertaining, and is also a perspective that is shared by Kanoria et al [1–3].

Out-of-equilibrium phenomena. The “economic” perspective on market algorithms is that of agent-based computational economics, which seeks to directly address the problem of realistically modeling market dynamics outside of equilibrium.

Here, market dynamics are derived by making ad-hoc behavioral assumptions about agent behaviour – we defer to [17] for a review of these methods, and to [18–20] for examples of agent-based modeling of supply shocks, which is the context in which these methods are most widely used.

While literature in each of these areas is extensive, a unified picture is lacking. Outside of the mechanisms studied in ACE/Out-of-equilibrium phenomena, there seems to be very little work on explicitly reverse-engineering the behavioral assumptions that completely specify a particular market mechanism, or conversely on deriving the consequences of particular behavioral assumptions.

The problem of reverse-engineering utility functions (up to a certain equivalence class) from the equilibrium has been studied before in the context of inverse game theory [21–25]. However, the utility functions themselves are not sufficient to determine the *dynamics* of the game, only its equilibrium.

There are efforts in mechanism design to discover certain desired properties of market mechanisms (the seminal result on which is the Myerson-Satterthwaite theorem [26]), these do not amount to being able to uniquely determine the dynamics from some set of assumptions on the agents.

The closest vein of research to my plan appears to be the research of Kanoria et al [1–3], who constructed a simple agent-based model of bargaining in exchange networks and demonstrated its equivalence to a belief propagation algorithm, thereby founding the latter in precise behavioral assumptions. However, like in agent-based computational economics, these behavioral assumptions are ad-hoc, and not expressible in a general theoretical framework; their work cannot directly be extended to a different model of agent behavior.

In general, the dynamics of a game cannot be uniquely determined from the agents’ utility functions, or any equivalent formulation of the game set-up; to do so, a new theoretical framework is needed. This is the gap my research plan seeks to address.

Research problems.

I seek to develop a unified framework for algorithms of market dynamics, under which specific market algorithms can be derived as special cases, as the “natural dynamics” that arise under certain behavioral assumptions.

These assumptions cannot be formulated within the framework of the game itself, as the agents’ utility functions only determine the equilibrium of the game, and not the algorithm by which this equilibrium is computed.

Instead, I suggest that the dynamics of an agent should itself be considered a “choice” in some more abstract sense. Then the dynamics of the original game arises as the unique equilibrium of this new game with expanded choice sets.

Research hypothesis 1. For any game G and algorithm p on G , a temporal game \hat{G} can be constructed with an expanded choice space and a unique equilibrium \hat{x} corresponding to p .

Such a \hat{G} can be constructed uniquely for any (G, p) . Each particular p will then correspond to some set of statements about the utility functions of agents in the temporal game. By altering these agent-level assumptions, we will be able to produce different algorithms of out-of-equilibrium behavior.

The intuition here is thus: in some philosophical sense, the market is “always” at equilibrium. Less blasphemously, one may say that an agent acting sub-optimally to its environment can still be considered optimal if we say that any other choice would be impossible or infinitely expensive, simply in the sense that the agent was not able to consider that choice in time. Obviously, this statement is blatantly false under a precise specification of the game, and thus the formalization of this idea involves constructing a new game.

The thesis that market dynamics may be derived from some utility functions on an expanded game depends fundamentally on the notion of agents being *strategic* – this is not the case with message-passing algorithms, which assume full collaboration.

However, I am still interested in including message-passing algorithms in my proposed framework, as I believe the perspective of market dynamics being fundamentally being a matter of information propagation (via the price aggregation mechanism) is rather interesting and natural. To this end, I suggest that one may “force” the agents to collaborate by imposing a *rights structure* on the agents’ choices.

Rights structures are the central of idea of property rights theory, and are mathematically formalized e.g. in [27, 28]. The simplest example of a rights structure is the initial endowment of a Walrasian setting: in general, a rights structure allows an agent to “forbid” certain choices of another agent’s. This connects naturally to our problem, as this allows agents to enforce collaboration by another agent.

In my own previous work [28], I suggested that there may be a correspondence between rights structures and aggregate utility functions (i.e. that each rights structure maximizes a particular utility function at equilibrium), and that this may be worth further study. This is a problem of direct interest to us, as we are interested in finding the rights structure implied by a particular message-passing algorithm (i.e. by a particular function being optimized).

Research problem 2. Determine the precise correspondence between rights structures and aggregate utility functions, and use this to formulate message-passing algorithms as equilibrium computation algorithms on games with rights structures.

Applications and significance.

The most immediate motivation for RH 1 – i.e. for developing a general theoretical framework for market algorithms – lies in agent-based computational economics and the study of out-of-equilibrium phenomena.

Existing work in this field essentially makes certain ad-hoc assumptions about agent behavior [17] and subsequently derives a model of market dynamics from these assumptions. A general theoretical framework will enable a wider and more systematic study of possible market dynamics, and will allow us to draw general results connecting specific behavioral assumptions to specific emergent behaviors.

In particular, my work will have implications in microfoundational work. Microfoundational economics refers to the attempt to model macroeconomic concepts as emergent phenomena arising from underlying microeconomic laws: for example, previous economics research has attempted to create reductionist models of economic growth [29], business cycles [30] and agent-based simulations of specific observed market behaviors such as loyalty [31] (the last of these isn't technically "macro", but belongs to a similar philosophical paradigm).

My proposed theoretical framework will bring greater scientific clarity to such phenomenological work, as it enables us to make precise scientific predictions, allowing the falsification of models of agent behavior and bounded rationality.

Furthermore, RP 2 directly pertains to the theoretical foundations of welfare economics, and could in similar spirit pave the way for research into the effects of redistribution, or on the relative inequalities in preference fulfilment resulting from a particular rights structure.

Conversely, my work will also have implications in computer science and machine learning algorithms.

There is a popular quote attributed to Christos Papadimitriou in agent-based computational economics: "If your laptop can't find it, then neither can the market". The contrapositive to this is that a realistic model of information propagation in the market will allow for computationally efficient multi-agent co-ordination mechanisms.

The development of a "natural" model of market dynamics from realistic assumptions about the temporal game will contribute directly to this goal. This serves to address the known computational inadequacies of existing message-passing algorithms like Max-Sum [16].

Furthermore, many classic machine learning/optimization algorithms may be mathematically equivalent to multi-agent co-ordination algorithms – this is shown by [1–3] for belief propagation, and I suspect that a differential version of Max-Sum may look a lot like backpropagation when implemented on acyclic graphs (there is a continuous version of Max-Sum in the literature [32], but it is not marginal in nature). While this line of research is more speculative, my research agenda in general may have implications for studying (or reverse-engineering) the convergence and complexity properties of traditional learning algorithms.

Tentative research plan.

In the proposal above, I've listed several closely-aligned problems that I am interested in; the proposal is ambitious, and it is not necessarily the case that they can all practically be completed in a PhD timeframe. While my research problem RH 1 is essential to my research agenda as I currently envision it, the relevance and meaningfulness of RP 2 depends largely on the progress that I will make in my first year of research, and the applications described are essentially auxiliary.

Below is a coarse tentative research plan/outline of the research phase of my PhD:

Background reading. The central theme of my research agenda is the explicit specification of axiomatic behavioral assumptions and deriving market mechanisms emergent from them; one of the purposes of doing so is to derive natural or realistic models of market dynamics.

This is closely related to the goals of agent-based computational economics (ACE), which [5] describes as “growing economies from the ground up”. In order to perform productive research towards this goal, it is essential that I review the existing literature in this area; a comprehensive review work is [17].

Apart from ACE, the most interesting consequences of my research might be for multi-agent learning with strategic agents, an area I know comparatively little of; references include [12–14]. Other relevant references to my work include [1–4, 27]; I will also briefly review the microfoundations literature (including [29, 30]).

I intend to start this background reading before starting my PhD and end it two months into the program at the very latest.

Months 1-8. In the first eight months of my PhD, I will work on “reverse-engineering” the assumptions behind some market algorithms of interest.

At this point, I do not necessarily intend to have rigorously formulated the theoretical framework that I aim to eventually state these assumptions in; I will see if the results I obtain are most naturally interpreted in such a framework, and will keep an open mind to other possible frameworks.

As doing this with message-passing algorithms in this way would require progress on RP 2, I intend to focus mostly on mechanisms like auctions in this period. Depending on what results I get and the natural flow of my work, I might choose to incorporate repeated games or message-passing algorithms (in the latter case, I would have to begin work on RP 2 early).

Based on the results and time constraints, I could then work on applications and analogies to classic optimization algorithms (including addressing RP 2b), which I believe would have the potential to lead to papers in the standard top-tier machine learning conferences.

Months 9-16. I will then make an effort to produce the desired theoretical framework, thus addressing RH 1. I expect coming up with this framework itself to not take too long given the insights from the previous phase, but I could foresee many interesting theoretical problems arising (one that I can already anticipate is the proof of the uniqueness of the temporal game equilibrium).

This is tentatively the most central and important part of the project (as I currently have it planned), and has the potential to lead to several papers e.g. in conferences such as LOFT and EC, and journals like Games and Economic Behavior.

Months 17-24. Depending on the intuition I have developed by this stage for my work on other mechanisms, I may decide to work either on RH 2 (therefore allowing us to incorporate message-passing algorithms into our framework) or may decide to delve deeper into the applications of my research. In the latter case, I would focus primarily on proposing laws for the “natural dynamics” of markets in the temporal game framework.

Months 25-36. In the final year of my PhD, I will focus on exploring the applications I’ve previously described, as I believe my work will open many exciting new directions of future research beyond my PhD. The economic applications of my work are those that currently hold my interest, although this is subject to change based on what I discover in the process.

It is likely that I will only be able to do exploratory research on the applications, and that they could each be the subject of interesting long-term research programs.

I am very excited to begin work on my research plan, which I believe addresses interesting and impactful problems that are particularly significant to the department’s strong program in computational social science and mechanism design.

I have previous research experience in theoretical economics, specifically property rights theory: [28]. This relates directly to my interest in RP 2, but also trained me in the relevant mathematical formalisms for RH 1, as I used games with expanded choice sets to formulate the mathematical notion of rights in this paper. I believe this makes me well-qualified to work on this research plan.

References

- [1] Yashodhan Kanoria et al. *Bargaining dynamics in exchange networks*. 2011. arXiv: 1004.2079 [cs.GT].
- [2] Yashodhan Kanoria et al. “Fast Convergence of Natural Bargaining Dynamics in Exchange Networks”. In: *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. San Francisco, California: Society for Industrial and Applied Mathematics, 2011, pp. 1518–1537.
- [3] Yashodhan Kanoria et al. *A Natural Dynamics for Bargaining on Exchange Networks*. 2010. arXiv: 0911.1767 [cs.GT].
- [4] Vijay V Vazirani. “Combinatorial algorithms for market equilibria”. In: *Algorithmic Game Theory*. New York: Cambridge University Press, 2007. Chap. 5, pp. 22–26, 103–134.
- [5] Leigh Tesfatsion. *Agent-Based Computational Economics*. Working Paper 1. Iowa State University, 2003. URL: <http://www2.econ.iastate.edu/tesfatsi/acewp1.pdf>.
- [6] Dhananjay K. Gode and Shyam Sunder. “Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality”. In: *Journal of Political Economy* 101.1 (1993), pp. 119–137. doi: 10.1086/261868.
- [7] Sham M. Kakade, Michael Kearns, and Luis E. Ortiz. “Graphical Economics”. In: *Learning Theory*. Ed. by John Shawe-Taylor and Yoram Singer. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 17–32. ISBN: 978-3-540-27819-1.
- [8] Rustam Tagiew. *Towards barter double auction as model for bilateral social cooperations*. 2009. arXiv: 0905.3709 [cs.GT].
- [9] R.Preston McAfee. “A dominant strategy double auction”. In: *Journal of Economic Theory* 56.2 (1992), pp. 434–450. ISSN: 0022-0531. doi: 10.1016/0022-0531(92)90091-U.
- [10] M. Babaioff and N. Nisan. “Concurrent Auctions Across The Supply Chain”. In: *Journal of Artificial Intelligence Research* 21 (May 2004), pp. 595–629. ISSN: 1076-9757. doi: 10.1613/jair.1316.
- [11] Moshe Babaioff and William E. Walsh. “Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation”. In: *Decision Support Systems* 39.1 (2005). The Fourth ACM Conference on Electronic Commerce, pp. 123–149. ISSN: 0167-9236. doi: 10.1016/j.dss.2004.08.008.
- [12] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. *Multi-agent reinforcement learning: a selective overview of theories and algorithms*. 2021. arXiv: 1911.10635 [cs.LG].
- [13] Lorenzo Canese et al. “Multi-agent reinforcement learning: a review of challenges and applications”. In: *Applied Sciences* 11.11 (2021). ISSN: 2076-3417. doi: 10.3390/app11114948.
- [14] Jan-P. Calliess and Geoffrey J. Gordon. “No-Regret learning and a mechanism for distributed multiagent planning”. In: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*. Estoril, Portugal: International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 509–516. ISBN: 9780981738109.
- [15] A. Rogers et al. “Bounded approximate decentralised coordination via the max-sum algorithm”. In: *Artificial Intelligence* 175.2 (2011), pp. 730–759. ISSN: 0004-3702. doi: 10.1016/j.artint.2010.11.001.
- [16] Md. Mosaddek Khan, Long Tran-Thanh, and Nicholas R. Jennings. “A generic domain pruning technique for GDL-based DCOP algorithms in cooperative multi-agent systems”. In: *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*. Stockholm, Sweden, 2018, pp. 1595–1603.
- [17] William Brian Arthur. “Out-of-equilibrium economics and agent-based modeling”. In: *Handbook of Computational Economics*. Vol. 2. Amsterdam: Elsevier, 2006. Chap. 32, pp. 1551–1564.

- [18] Alexander Outkin et al. “An agent-based modeling approach to non-equilibrium dynamics of natural gas supply shock propagation”. In: *International Association of Energy Economists 2014 Conference*. 2014.
- [19] Mauro Napoletano, Jean-Luc Gaffard, and Zakaria Babutsidze. *Agent-based models: a new tool for economic and policy analysis*. Tech. rep. hal-01070338f. 2012.
- [20] Dhruv Sharma. *Macroeconomic agent-based models: a statistical physics perspective*. Tech. rep. tel-03199888f. 2020.
- [21] Volodymyr Kuleshov and Okke Schrijvers. “Inverse Game Theory: Learning utilities in succinct games”. In: *Web and Internet Economics*. Ed. by Evangelos Markakis and Guido Schäfer. Berlin, Heidelberg: Springer, 2015, pp. 413–427. ISBN: 978-3-662-48995-6.
- [22] Eddie Dekel and Marciano Siniscalchi. “Epistemic Game Theory”. In: ed. by H. Peyton Young and Shmuel Zamir. Vol. 4. *Handbook of Game Theory with Economic Applications*. Elsevier, 2015. Chap. 12, pp. 619–702. DOI: 10.1016/B978-0-444-53766-9.00012-4.
- [23] Robert J. Aumann. “Backward induction and common knowledge of rationality”. In: *Games and Economic Behavior* 8.1 (1995), pp. 6–19. ISSN: 0899-8256. DOI: 10.1016/S0899-8256(05)80015-6.
- [24] David Abel et al. *On the Expressivity of Markov Reward*. 2021. arXiv: 2111.00876 [cs.LG].
- [25] Nicholas Bishop, Le Cong Dinh, and Long Tran-Thanh. “How to Guide a Non-Cooperative Learner to Cooperate: Exploiting No-Regret Algorithms in System Design”. In: *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2021, pp. 1464–1466. ISBN: 9781450383073.
- [26] Roger B Myerson and Mark A Satterthwaite. “Efficient mechanisms for bilateral trading”. In: *Journal of Economic Theory* 29.2 (1983), pp. 265–281. ISSN: 0022-0531. DOI: 10.1016/0022-0531(83)90048-0.
- [27] Peter Gärdenfors. “Rights, games and social choice”. In: *Noûs* 15.3 (1981), pp. 341–356. ISSN: 00294624, 14680068. DOI: doi.org/10.2307/2215437.
- [28] Abhimanyu Pallavi Sudhir. *A mathematical definition of property rights in a Debreu economy*. 2021. arXiv: 2107.09651 [econ.TH].
- [29] Gunnar Eliasson et al. *Microfoundations of Economic Growth: A Schumpeterian Perspective*. Selected papers from the International Schumpeter Society, Biennial Meeting (6 : 1996 : Stockholm). University of Michigan Press, 1998. ISBN: 9780472109043. URL: <https://books.google.co.uk/books?id=ggfsAAAAMAAJ>.
- [30] Giovanni Dosi, Giorgio Fagiolo, and Andrea Roventini. “The microfoundations of business cycles: an evolutionary, multi-agent model”. In: *Schumpeterian Perspectives on Innovation, Competition and Growth*. Springer, 2009, pp. 161–180.
- [31] Aleksandra Alorić et al. “Emergence of Cooperative Long-Term Market Loyalty in Double Auction Markets”. In: *PLoS ONE* 11 (2016). DOI: 10.1371/journal.pone.0154606.
- [32] Thomas Voice et al. “A hybrid continuous Max-Sum algorithm for decentralised coordination”. In: *19th European Conference on Artificial Intelligence*. 2010, pp. 61–66. URL: <https://eprints.soton.ac.uk/271239/>.