November 2019

Research Statement: game theory & economic design & algorithms

Fedor Sandomirskiy, Ph.D. (Candidate of Sciences),
postdoc at Technion, Game Theory group,
senior research fellow at GT Lab HSE St. Petersburg
homepage: https://www.fedors.info/

1. General overview

I have a passion for applying sophisticated mathematical tools for understanding interactions in multi-agent systems and for finding ways to improve the outcomes of these interactions by selecting clever rules.

My research is on *game theory* and *economic design* in a broad sense. I am especially interested in *fair and efficient resource allocation mechanisms* (see Section 2) and *strategic use of Information* (Section 3), but my research is not limited to these topics (Section 4).

Methods and problems that I like. In my research on *economic design*, I combine the approach of *microeconomics* with the modern insights from *algorithmic game theory* to get *positive results*, concrete recommendations for practitioners instead of impossibilities immanent in classical literature.

In *game theory*, I am intrigued by the problems, where equilibrium behavior is too complicated to be described explicitly. For example, this may occur due to information effects or the multistage nature of the interaction. Analyzing the qualitative properties of such games, I combine *strategic insights* with the "magic of mathematics," which allows investigating properties of the strategic behavior without describing the behavior itself explicitly.

In *Mathematics*, I like the interplay of *Convex Analysis*, *Probability*, and *Functional analysis*. I am interested in the behavior of large random structures and concentration effects, e.g., typical and asymptotic properties of large markets, large games, etc. In seemingly different contexts of resource-allocation and information design, I found powerful connections with structured optimization problems such as optimal transportation and maximal flows and their duals. These tools seem to be still underestimated in microeconomics.

My background and trajectory. I hold a master degree in Mathematical Physics from St. Petersburg State University. This experience gave me a strong math background (especially, in functional analysis, probability, linear models, and dynamical systems), which I applied later in my game theory research. The results of my master thesis (2011) about disordered quantum systems were published in *Communications in Mathematical Physics*, one of the top journals [6] in the field.

I became interested in game theory at the end of my academic studies at the Physics department and decided to make it my main research direction. As a part of my Ph.D. studies at the Central Econ-Math Institute RAS, I learned economics, and acquired additional training in Computer Science. My thesis on "Extreme asymptotic regimes in repeated games with incomplete information" was defended in two and a half years (2014), which is rather fast even for the Russian academic tradition. The results were considered outstanding by the committee; they are briefly discussed in Section 3. These results were published by *International Journal of Game Theory* [5] and *Dynamic Games and Applications* [4]. In parallel to my graduate studies, I worked as a researcher at the Math department in the group of Stanislav Smirnov, a Fields medalist.

In 2015, my colleagues and I established a game theory research group at HSE St. Petersburg. Herve Moulin and Anna Bogomolnaia from Glasgow University became its academic supervisors. Meeting Herve and Anna, who introduced me to the world of economic design and became my long-term collaborators, opened a new page of my research career and led to publications in *Econometrica* [1], *Social Choice and Welfare* [2], and a recent R&R in *Management Science* [3].

In 2015-2016, I was involved in various administrative and organizational activities, teaching, and advising at HSE. In 2017 I headed the group. However, later, I decided to focus more on my research than on administrative issues.

In 2018, I joined Technion as a postdoc, where I deepened my understanding of interconnectedness of economic design, game theory, and artificial intelligence. Enjoying many exciting collaborations at Technion, I keep the professional connection with the St. Petersburg group.

2. Economic design: mechanisms and algorithms beyond worst-case analysis

My research and the new paradigm of economic design. Classic economic design is full of impossibility results (even minimal reasonable assumptions of non-manipulability, efficiency, and fairness are usually incompatible), e.g., the famous Gibbard–Satterthwaite theorem or Zhou's impossibilities [14, 15]. These impossibilities demonstrate that there is no "ideal" mechanism satisfying all the requirements; however, classical results are unable to provide any recipe for reasonable compromises so important for practice.

Motivated by this concern, I got interested in ways to avoid impossibilities and make the theoretical economic design closer to practice. The classical worst-case paradigm labels a mechanism as bad if it fails to satisfy the desired property on one preference profile. In my papers, I use alternative approaches to assessing mechanism's performance: evaluating trade-offs between quantitatively relaxed design objectives rather than proving their incompatibility and analyzing average-case behavior of a mechanism on a random preference profile or on all profiles except a small subset of degenerate ones. Recent results in the field of algorithmic game theory confirm the power of these ideas: a moderate relaxation of design objectives is usually enough to escape impossibilities [16]; "bad" instances are rare (e.g., [17], [18], [19], [20], [23]); and mechanisms demonstrate much better performance on the real-world or random preference profiles compared to worst-case predictions. Other ideas making my research closer to practice is the importance of small message spaces (aka concise bidding languages), algorithmic complexity issues, and robustness with respect to modeling assumptions.

2.1. **Competitive equilibrium in fair division of bads and other constrained economies.**

In everyday life, we often decide how to allocate bads, not goods. Think of distributing paperwork at the department or house chores. However, in contrast to the well-studied case of goods, problems with bads have not been receiving much attention. Apparently, this occurred due to the wrong intuition that a simple change of sign in one formula is enough to translate results from goods to bads.

2.1.1. *Surprising non-similarity of economies with goods and with bads.*

Anna Bogomolnaia, Herve Moulin, Elena Yanovskaya, and I initiated the systematic study of fair division of divisible bads.

In our paper [1], published in *Econometrica*, we assume general homothetic preferences, and [2] (*Social Choice and Welfare*) refines the results in the special case of additive utilities important for practice, e.g., used on Spliddit.org.

We found that problems with goods and those with bads are not equivalent, and this structural non-similarity leads to new obstacles. For example, no single-valued envy-free and efficient division rule for bads is robust to small perturbations of preference profile. For goods, the competitive equilibrium from equal incomes (CEEI) in the associated exchange economy is known to have all these properties; in addition, it is unique and is efficiently computable as a maximizer of the Nash Social Welfare (Eisenberg-Gale convex program).

We extended the notion of competitive equilibrium to economies with a mixture of goods and bads. Equilibria remain envy-free and efficient; however, the set of competitive equilibria becomes disconnected with a potentially exponential number of connected components. They correspond to critical points (local minima, maxima, and saddle points) of the Nash Social Welfare on the Pareto frontier, while none of the global extrema is fair and efficient.

2.1.2. *Bads as a constrained economy and open questions.* The multiplicity of competitive equilibria has a deep origin: the economy with bads can be reduced to an economy with goods but with additional constraints on consumption. Constrained economies are known for their non-monotonicity and multiplicity phenomena.

This observation leads to many intriguing questions. Do the same phenomena hold for other constrained fair division problems, e.g., the assignment economy of Hylland and Zeckhauser [21]? How to pick a natural single-valued selector? Preliminary results show that multiplicity disappears in large economies. How to compute a competitive equilibrium? Existing methods use the convexity of the problem, which disappears in constrained economies. We discuss our algorithmic results in the next subsection.

2.2. **Fair division algorithms.** To use a mechanism in practice, its outcome must be efficiently computable. Motivated by this concern, I started working on algorithmic questions, which form the core of algorithmic game theory.

2.2.1. *Computing competitive equilibria in constrained economies.* The results on fair division of bads described above raise the question of computing competitive equilibria in economies with divisible bads or, more generally, in constrained economies. Existing algorithms rely on the convexity of the problem and are non-applicable to non-convex economies, where the set of competitive equilibria becomes disconnected. The only exception is the recent paper [22], which computes an equilibrium in Hylland-Zeckhauser model, but relies on a "black-box" of cell-enumeration technique from computational algebraic geometry, more a theoretical construction than a practical solution.

In a working paper [8] with Simina Branzei submitted to *Mathematics of Operations Research*, we suggest an algorithm for computing all equilibria of economies with bads under additive utilities. It runs in polynomial time if either the number of agents or bads is fixed. The algorithm is based on the novel idea that the set of all Pareto-optimal demand structures (aka consumption graphs) has a polynomial size for almost all preference profiles. We enumerate all such structures and then check them for compatibility with equilibrium conditions, which boils down to solving a maximum-flow problem.

Our approach is much more elementary and easier to implement than that based on cell-enumeration. We believe that it also extends to other constrained economies and, in particular, provides a much simpler algorithm for Hylland-Zeckhauser assignment economies.

2.2.2. *Sharing-minimization as a new approach to fair division with valuable items.* In a follow-up paper [9] with Erel Segal-Halevi, we demonstrated the power of Pareto-frontier enumeration in a different problem.

The modern literature on fair division falls into one of the two extremes: economists focus on divisible items, while CS people consider discrete objects and are then forced to use approximate fairness notions to guarantee existence. Practical problems are somewhere in between: while objects are usually indivisible (e.g., apartments and cars), a divorcing couple would prefer keeping a small number of objects shared to an allocation "envy-free up to one apartment."

We proposed an alternative approach of sharing minimization, bridging the two streams of literature. The idea is to keep fairness and efficiency as constraints and then minimize the number of shared items. Using the Pareto-frontier enumeration approach from my paper [8] with Simina, we show that this minimization is algorithmically tractable for preference profiles that are not too degenerate.

One methodological corollary of our analysis confirms that the concept of fractional Pareto optimality from [24] has better algorithmic properties than widely-used discrete Pareto-optimality. This opens a possibility to refine many negative results about algorithmic hardness in economies with indivisibilities, e.g., [25].

2.2.3. *Existence of approximately-fair allocations.* When sharing is impossible due to institutional constraints or agents' views, fair allocation may fail to exist. To circumvent this non-existence, several notions of approximate fairness were suggested by the CS community such as envy-freeness up to one item, proportionality up to one item, or max-min share.

Our recent note [10] with Haris Aziz and Herve Moulin, submitted to *Operations Research Letters*, provides a simple algorithm for computing approximately-proportional fractionally Pareto-optimal allocations for a mixture of indivisible goods and bads. Even the existence of such allocations was not known. The simple algorithm relies on two ingredients: a trading-cycle algorithm for finding divisible fractional-Pareto-improvements (a side result in our paper [9]) and an extension of Barman-Krishnamurthy rounding [23].

2.3. **Dynamic resource-allocation and matching markets.** Many resource allocation problems are dynamic (aka "online") by their nature, for example, distributing the incoming flow of tasks among workers. In such problems, allocation decisions are made in an intertemporal context and are agnostic to the exact future of the process. The importance of making timely decisions becomes especially stark if agents have high waiting costs as in kidney-exchange market.

The online fair division recently attracted the attention of the CS and Economic communities (e.g., [26]); however, the known results achieve fairness without any efficiency guarantees.

2.3.1. *Combining fairness and efficiency in a simple model.* Our working paper [3] with Anna Bogomolnaia and Herve Moulin (R&R in *Management Science*) is the first attempt to achieve both goals. We consider a simple Bayesian prior-independent setting with a fixed set of agents, where objects are i.i.d. across periods with some unknown distribution (possibly selected in an adversarial way). In this setting, we identify ex-ante fair rules that are undominated in terms of expected welfare. In the case of goods, we get a one-parametric family of undominated rules, while for bads such rules are unique (a repercussion of goods/bads non-similarity as in Subsection 2.1).

Values of Competitive Ratios and Price of Fairness accompanied by numerical examples demonstrate that collecting additional statistical information about prior leads to a tiny gain. This justifies using our simple and robust prior-independent rules, even if some statistical information is available.

These results are based on novel links between prior-independent guarantees and convexity, geometric insights about the structure of worst-case instances, and probabilistic decoupling arguments.

2.3.2. *Extensions.* Motivated by these positive results in a simple model, I am now working on finding fair and efficient online rules in more-complicated markets, where agents may have non-trivial intertemporal preferences as in the realistic foodbank-problem [27] and may also arrive and depart in the online fashion as in kidney-exchange markets.

Preliminary results suggest that building such an online rule can be reduced to finding a competitive equilibrium in an auxiliary constrained static economy. However, questions of its existence and computability remain widely open.

2.4. **Fairness and Incentives in Artificial Intelligence.** Being a member of the research group of Moshe Tennenholtz, I learned about multiple (yet undeveloped!) connections between familiar problems of economic design and fairness & incentive issues in algorithmic classification and recommendation systems.

2.4.1. *Fairness in classification.* algorithmic decisions often inherit human biases. For example, Apple's new credit card is not free of gender bias [28], and the automated recruiting tool used by Amazon was favoring men [29]. To circumvent such unfairness, the academic community came up with quite a few design constraints that a "fair" algorithm must satisfy.

In our recent working paper [11] with Omer Ben-Porat and Moshe Tennenholtz, we bridge the literature on fair classification with the egalitarian approach to fairness from welfare-economics. We focus on scenarios where fairness constraint is imposed on a self-interested party (e.g., a revenue-maximizing bank predicting reliability of a borrower and making lending decisions).

We uncover a serious gap in the preceding analysis of fair classification. Popular fairness constraints may harm the disadvantaged protected group of borrowers (defined by race or gender) even more because, for a selfish decision-maker, it is profitable to reallocate the "price of fairness" to this group. We also introduce the concept of welfare-equalizing fairness that allows avoiding such paradoxes.

2.4.2. *Incentivizing high-quality content by clever ranking algorithms.* Recommendation systems and search engines not only help users to enjoy the desired content but also implicitly determine who of the content providers receives the revenue generated by a particular user. Therefore, in a long run, the ranking algorithm affects which businesses grow and flourish and which disappear from the market.

In an ongoing project with Moshe Tennenholtz, our goal is to determine those ranking mechanisms that incentivize content providers to create better content in equilibrium. This direction turns out to be close to designing all-pay auctions with non-standard revenue-like objective as in [30].

2.4.3. *Voting in the policy space.* In an ongoing project with Reshef Meir and Moshe Tennenholtz, we aim to improve the quality of elections by allowing voters to submit ballots containing their preferences on possible policies (perhaps, not represented by any of the candidates) instead of their preferences on candidates. This line of research is related to ideas of liquid democracy [31] and to voting in metric spaces [32].

3. Strategic use of information and information design

My first encounter with game theory was in my last year at the department of physics. From Viktor Domansky, who became my Ph.D. advisor in a few months, I learned about repeated games with incomplete information: a classic model introduced by Robert Aumann and Michael Maschler. I was impressed both by the soundness of the model and by the elegance of mathematics (martingale-based arguments and Blackwell's approachability) and decided to make game theory my field of study.

I still find intriguing the role that information plays in strategic interactions and how rigorous analysis of these interactions spans over various branches of mathematics.

3.1. **Multistage interactions with asymmetric information.**

In a repeated game with incomplete information, an agent informed of some payoff-relevant state interacts with an uninformed opponent. However, the interaction has many rounds, and by observing the actions of the informed player, the uninformed one can try guessing the state. So the information leaks through actions.

I was intrigued by the question what the value of information in long multistage interactions is? Does this value reflect any information-theoretic "quantity" of information like Shannon's entropy? How it depends on the strategic properties of a one-stage game? These questions inspired my papers [5] in *International Journal of Game Theory* and [4] in *Dynamic Games and Applications*, which contain essential results from my Ph.D. thesis.

3.1.1. *Games with anomalously high value of information.* The publication [5] is devoted to games with large sets of states. I showed that the value of information in such games can grow anomalously fast with the number of repetitions and the speed of growth is controlled by entropy-like family of functionals.

3.1.2. *When the value of information remains bounded?* The paper [4] focuses on the opposite extreme. It characterizes the whole class of games where the value of information remains bounded. An easy-checkable property of a one-stage game turns out to be responsible for this phenomenon. Examples of such games were known in the context of insider trading on financial markets, see [33] and [34]; however, the analysis there was based on ad hoc arguments using the structure of a particular example.

My result relies on a reduction of the game to a martingale-optimization problem and on ideas from theories of optimal-transportation and monotone operators.

A conjecture that puzzles links boundedness of the value of information and existence of the optimal policy of informed player that reveals the whole information in finite time.

3.2. **Information design (=Bayesian persuasion).**

The most fundamental question of strategic use of information is what beliefs can a more informed agent (sender) induce among less-informed agents (receivers) by sending them a certain signal? This signal can be either a certain message as in the theory of Bayesian persuasion or an action as in repeated games discussed above.

If there is just one receiver or the signal is public, the answer is given by the celebrated Splitting lemma of Robert Aumann: any posterior distribution of beliefs satisfying the so-called martingale condition (the average posterior equals prior), can be induced by a certain signaling policy. This result is a key technical tool both in the theory of repeated games with incomplete information and in the theory of Bayesian persuasion. However, with private signals, no extension of the Splitting lemma to several receivers was known.

In an ongoing project [12] (soon on arXiv) with Itai Arieli, Yakov Babichenko, and Omer Tamuz we find the analog of the Splitting lemma for two or more receivers. The result turns out to be surprisingly non-trivial and sheds light on the connection of the problem with maximal flows and optimal transportation. We apply these insights to problems of Bayesian persuasion.

3.3. **Social learning.**

Models of repeated games with incomplete information and Bayesian persuasion describe information-transmission on the micro-level. Theory of social learning deals with macro-level counterpart. Understanding how information propagates through society and how it depends on ties in a social network is crucial for designing marketing campaigns. Such campaigns usually aim to identify the most influencing agents in a social network and make them early adopters of a certain new product in a hope that this adoption will spread over society. However, recent results [35] challenge the effectiveness of targeting influencers and show that using random agents leads to almost the same result.

In a project with Itai Arieli and Rann Smorodinsky [13] (on arXiv in December 2019), we go further and challenge the very existence of a small group of agents critical for information aggregation. We consider a social network with fully-rational Bayesian agents and random arrival order. Each agent gets a bounded signal about an unobserved payoff-relevant and makes an action after observing actions chosen by his neighbors that arrived earlier. We say that a sequence of networks supports learning if the fraction of agents making payoff-maximizing action tends to one along this sequence.

First, we show that learning is robust to the elimination of large random groups of agents: if a sequence of network exhibits learning, then after deleting randomly 99% of the population in each network, learning persists with high probability. However, this does not exclude the possibility that a certain small minority of agents is critical for learning. Then we demonstrate that, surprisingly, learning can also be possible in totally-egalitarian networks, where all agents are symmetric and thus ex-ante all of them are equally influential. We construct such networks using novel insights from the theory of expanders and demonstrate that even adversarial elimination of 99% of agents cannot spoil learning outcomes.

4. Some other research interests

My interests are not limited to the directions mentioned above. Some topics exited my curiosity because of involvement in applied projects, IT-industry consulting, and student supervision at HSE St. Petersburg. To name a few directions: designing payment-mechanisms for crowd-sourcing platforms with heterogeneous workers, detection of collusion and corruption in online procurement auctions using machine-learning methods, incentives in blockchain, exploring the typical structure of reported preferences of real users by the data from Spliddit.org.

Apart from these topics, I am interested in large random games (concentration effects for equilibrium payoffs and strategies); in typical properties of extensive-form games (e.g., is playing White is advantageous on average if the game of chess starts from a random position, similarly to Fischer random chess?); in using the topological properties of Nash-correspondence to design smooth reforms (how to reach good equilibrium from a bad equilibrium by smoothly changing the payoffs in an optimal way?); and in some others.

References

My publications

[1] Anna Bogomolnaia, Herve Moulin, Fedor Sandomirskiy, Elena Yanovskaya Competitive division of a mixed manna
Econometrica, 2017, vol.85:6, p.1847-1871
preprint: arXiv:1702.00616, slides: link

[2] Anna Bogomolnaia, Herve Moulin, Fedor Sandomirskiy, Elena Yanovskaya Dividing goods or bads under additive utilities
Social Choice and Welfare, 2019, vol.52:3, p.395-417.
preprint: arXiv:1608.01540

[3] Anna Bogomolnaia, Herve Moulin, and Fedor Sandomirskiy A simple Online Fair Division problem
Management Science (R&R October 2019)
arXiv 2019: arXiv:1903.10361, slides: link

[4] Fedor Sandomirskiy On repeated zero-sum games with incomplete information and asymptotically bounded values *Dynamic Games and Applications*, 2017, doi:10.1007/s13235-017-0217-7
preprint: arXiv:1509.01727

[5] Fedor Sandomirskiy Repeated games of incomplete information with large sets of states *International Journal of Game Theory*, 2014, vol.43:4, p.767-789
preprint: arXiv:1205.6791

[6] Alexander Fedotov, Fedor Sandomirskiy An exact renormalization formula for the Maryland model
Communications in Mathematical Physics, 2015, vol.334:2, p.1083-1099
preprint: arXiv:1311.6930

[7] Fedor Sandomirskiy Variation of measure-valued martingales and repeated games with incomplete information
Doklady Mathematics (English translations of Proceedings of the Russian Academy of Sciences), 2012, vol.86 p.796-798

My working papers

[8] Simina Brânzei, Fedor Sandomirskiy Algorithms for Competitive Division of Chores
submitted to *Mathematics of Operations Research*
arXiv 2019: arXiv:1907.01766, slides: link

[9] Fedor Sandomirskiy, Erel Segal-Halevi Fair division with minimal sharing
arXiv 2019: arXiv:1908.01669, slides: link

[10] Haris Aziz, Herve Moulin, Fedor Sandomirskiy A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation
submitted to *Operations Research Letters*
arXiv 2019: arXiv:1909.00740

[11] Omer Ben-Porat, Fedor Sandomirskiy, Moshe Tennenholtz Protecting the Protected Group: Circumventing Harmful Fairness
arXiv 2019: arXiv:1905.10546

[12] Itai Arieli, Yakov Babichenko, Fedor Sandomirskiy, and Omer Tamuz Feasible joint posterior beliefs
soon on arXiv.org

[13] Itai Arieli, Fedor Sandomirskiy, Rann Smorodinsky Who matters for social learning?
on arXiv.org in December 2019

Other references

[14] Zhou L. On a conjecture by Gale about one-sided matching problems. Journal of Economic Theory. 1990 Oct 1;52(1):123-35.

[15] Zhou L. Impossibility of strategy-proof mechanisms in economies with pure public goods. The Review of Economic Studies. 1991 Jan 1;58(1):107-19.

[16] Procaccia A., Tennenholtz M. Approximate mechanism design without money. InProceedings of the 10th ACM conference on Electronic commerce 2009 Jul 6 (pp. 177-186). ACM.

[17] D. Kurokawa, A. Procaccia, J. Wang. Fair Enough: Guaranteeing Approximate Maximin Shares. *Journal of the ACM* (forthcoming)

[18] Dickerson JP, Goldman J, Karp J, Procaccia AD, Sandholm T. The computational rise and fall of fairness. InTwenty-Eighth AAAI Conference on Artificial Intelligence 2014 Jun 21.

[19] B. Plaut, T. Roughgarden Almost envy-freeness with general valuations. 2017. arXiv preprint arXiv:1707.04769

[20] M. Babaioff, N. Nisan, I. Talgam-Cohen Competitive equilibria with indivisible goods and generic budgets. 2017 arXiv preprint arXiv:1703.08150.

[21] Hylland A, Zeckhauser R. The efficient allocation of individuals to positions. Journal of Political economy. 1979 Apr 1;87(2):293-314.

[22] Alaei S, Jalaly Khalilabadi P, Tardos E. Computing Equilibrium in Matching Markets. In Proceedings of the 2017 ACM Conference on Economics and Computation 2017 Jun 20 (pp. 245-261). ACM.

[23] Barman S, Krishnamurthy SK. On the proximity of markets with integral equilibria. InProceedings of the AAAI Conference on Artificial Intelligence 2019 Jul 17 (Vol. 33, pp. 1748-1755).

[24] Barman S, Krishnamurthy SK, Vaish R. Finding fair and efficient allocations. InProceedings of the 2018 ACM Conference on Economics and Computation 2018 Jun 11 (pp. 557-574). ACM.

[25] de Keijzer B, Bouveret S, Klos T, Zhang Y. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. InInternational Conference on Algorithmic Decision-Theory 2009 Oct 20 (pp. 98-110). Springer, Berlin, Heidelberg.

[26] Benade G, Kazachkov AM, Procaccia AD, Psomas CA. How to make envy vanish over time. In Proceedings of the 2018 ACM Conference on Economics and Computation 2018 Jun 11 (pp. 593-610). ACM.

[27] Lee MK, Kusbit D, Kahng A, Kim JT, Yuan X, Chan A, Noothigattu R, See D, Lee S, Psomas CA, Procaccia AD. WeBuildAI: Participatory Framework for Fair and Efficient Algorithmic Governance. Preprint. 2018.

[28] Taylor T. Apple Card algorithm sparks gender bias allegations against Goldman Sachs, The Washington Post, November 11, 2019, `https://www.washingtonpost.com/business/2019/11/11/apple-card-algorithm-sparks-gender-bias-allegations-against-goldman-sachs/`

[29] Dastin J. Amazon scraps secret AI recruiting tool that showed bias against women. San Fransico, CA: Reuters. Retrieved on October. 2018 Oct;9:2018. `https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G`

[30] Chawla S, Hartline JD, Sivan B. Optimal crowdsourcing contests. Games and Economic Behavior. 2019 Jan 1;113:80-96.

[31] Cohensius G, Mannor S, Meir R, Meirom E, Orda A. Proxy voting for better outcomes. InProceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems 2017 May 8 (pp. 858-866). International Foundation for Autonomous Agents and Multiagent Systems.

[32] Anshelevich E, Bhardwaj O, Elkind E, Postl J, Skowron P. Approximating optimal social choice under metric preferences. Artificial Intelligence. 2018 Nov 1;264:27-51.

[33] Domansky V. Repeated games with asymmetric information and random price fluctuations at finance markets. International Journal of Game Theory. 2007 Oct 1;36(2):241-57.

[34] De Meyer B, Marino A. Continuous versus discrete market game. Cowles Foundation Discussion Paper. 2005(1535).

[35] Akbarpour M, Malladi S, Saberi A. Just a few seeds more: value of network information for diffusion. Available at SSRN 3062830. 2018 Apr 1.