

Research Statement

November 2022

Fedor Sandomirskiy (Caltech)

Linde Postdoctoral Fellow in Economics, PIMCO Fellow in Data Science
links: [homepage](#), [CV](#), [RS](#), [TS](#), [JMP](#), [job market poster](#), [Google Scholar](#)

1. HIGH-LEVEL OVERVIEW

I am a microeconomic theorist specializing in *information economics* and *economic design*. My research brings new methodological tools to understand interactions among strategic agents, and improve the outcomes of these interactions by proposing new mechanisms. It blends microeconomic insights with ideas from algorithmic game theory and relies on the interplay of probability, convexity, and functional analysis. I am especially interested in the following topics.

(A) Information economics

Information economics studies how agents' behavior is affected by the information they have access to. My key interest is in methods for studying problems where different agents acquire information from *different sources*.

My main contribution to this field [1] was published in the *Journal of Political Economy*. The paper describes agents' belief distributions that can emerge when there are multiple sources of information. All my papers on the topic [1, 6, 8, 9, 12, 18, 19] are discussed in Section 2.

(B) Mechanisms and algorithms for fair and efficient resource allocation

The classical theory of economic design suffers from impossibility results: mild normatively appealing conditions on a mechanism are often incompatible. My key interest is in methods escaping these impossibilities. Such methods are at the heart of the new paradigm of economic design which stems from collaborations between economists and computer scientists.

My main contribution is to the theory of fair allocation of resources. The literature has focused on allocation of goods (scarce, valuable resources). My research has highlighted a surprising dissimilarity between this classical setting and problems with *bad*s (chores, waste, or liabilities) and inspired economic and computer science literature on fair division of bads. My papers describe the pseudo-market approach for bads and mixed bads/goods problems [2], give the first algorithmic results for bads [11], develop a novel practical approach to fair division based on sharing-minimization [3], and provide the first results on the robust use of statistical data in fair division [4].

These papers appeared in *Econometrica*, *Mathematics of Operations Research (R&R)*, *Operations Research*, and *Management Science*. They are discussed in Section 3 together with other related papers of mine [2, 17, 4, 11, 3, 14, 15, 16, 10].

(C) Common patterns behind (A), (B), and beyond

Seemingly unrelated problems of economic theory turn out to be connected to majorization, optimal transportation, maximal flows, mathematical tomography, and, more generally, measures with given marginals and their extreme points. Many of my papers uncover and exploit these connections. One of my current goals is to understand the connections themselves.

My working papers [7] and [8] show that multi-item multi-bidder auctions and multi-agent persuasion are, in fact, transportation problems. Ongoing projects explore consumer preference aggregation from this perspective [5] and establish a direct connection between information and Bayesian mechanism design inspired by their mutual link to majorization [13]. See Section 4 for details.

Below I describe my trajectory and how it has shaped my views and skills. I then turn to the detailed discussion of each of the directions (A), (B), and (C).

My background and trajectory. I started my career as a mathematician, defending my Master’s thesis in mathematical physics on disordered quantum systems [20] and completing my Ph.D. in mathematical game theory on¹ “The Value of Information in Repeated Games.” In parallel to my graduate studies, I worked as a researcher at the Math department in the group of Fields medalist Stanislav Smirnov. Later, I worked at the Game Theory Lab at HSE Univeristy. After getting interested in economic theory, I joined the Technion as a postdoc and then became a postdoc at Caltech.

2. INFORMATION ECONOMICS

Information about the environment in which we make decisions can affect our actions as much as monetary incentives and individual preferences. With the current rise of the digital economy, information has been playing an increasingly important role as it has become easy to collect and supply selectively.

2.1. Multi-agent information design. Information design studies what information to supply to agents (receivers) in order to incentivize their behavior in the directions desired. While modern recommendation systems can tailor information specially for each of the agents, the theory of information design has mainly focused on the case with just one receiver, or when all receivers get the same information publicly. In a series of papers [1, 6, 8, 9], we develop tools to handle multi-agent information design problems. A key ingredient in understanding problems of information design is the distribution of beliefs that the agents end up having.

We tackle this problem in a paper that appeared in the *Journal of Political Economy* [1]. It deals with a basic question that is a prerequisite for approaching information design: what belief distributions can agents end up with when they learn from different sources? If there is just one receiver, or the signal is public, the answer is given by Aumann’s Splitting lemma, the key technical tool both for information design and for the theory of repeated games with incomplete information discussed below. We find an 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 the agreement theorem, no-trade theorem, measures with given marginals, and Hilbert space geometry. These insights are applied to a multi-agent version of Bayesian persuasion, an information design framework developed by [21] and dominating the literature on information design. We show that optimal persuasion of two receivers may require an infinite number of signals in contrast to the single-receiver case where two signals are known to be enough.

In a working paper [6], we introduce a privacy concern to multi-agent information design: the information obtained by each of the agents must tell her nothing about the information received by others. We characterize ways to supply information such that it is impossible to give more information to any of the agents without violating privacy. The characterization relies on unexpected connections to the theories of majorization and mathematical tomography (a technique to uncover the geometry of high-dimensional sets by their lower-dimensional projections). Our results are closely related to the design of fair recommendation systems in the presence of protected attributes, and give an explicit recipe for achieving maximal informativeness under the fairness constraint; see also Section 3.3.

Papers [1] and [6] imply that measures with given marginals underlie multi-agent information design. The working paper [8] makes this connection explicit by showing that one

¹The results were published in [18] and [19] and are briefly discussed in Section 2.

can reduce multi-agent Bayesian persuasion to an auxiliary problem of optimal transportation (maximization of a linear objective over a set of measures with given marginals). We show how this reduction can be used to find explicit solutions for particular objectives, and demonstrate that it implies the dual representation for the optimal value, thus extending the $\text{cav}[u]$ -theorem of [21] and the duality by [22] to the case of multiple receivers.

A working paper [9] develops a model of multi-agent information design capturing persuasion in hierarchical organizations: the informed party and decision makers are separated by a sequence of mediators who, pursuing their own interests, may garble the information along the way. For one mediator, we give a geometric characterization of the optimal persuasion similar to the $\text{cav}[u]$ -theorem for non-mediated communication. In particular, two signals are enough for optimal persuasion as if there were no mediator. Surprisingly, adding more mediators complicates the problem dramatically, and two signals are no longer enough to persuade.

2.2. Observational learning. One of the descriptive goals of information economics is to understand how individual decisions shape the way information propagates through society as a whole. In [12], we study the ability of society to aggregate dispersed information, and papers [18, 19] explore how information asymmetries dissipate in long strategic interactions.

It is well-known that even a society of perfectly rational Bayesian agents may be prone to information cascades; whether such cascades emerge depends on the network of social ties. In [12] (*R&R in Journal of Economic Theory*), we study information aggregation and cascade behavior for casual decisions, such as buying cheap stuff. The order in which agents face such decisions is independent of the underlying network. This independence is captured by the assumption that the order in which agents choose is random. By contrast, most of the literature has assumed that the order of actions is correlated with the network structure. We observe that randomness leads to a wealth of new phenomena: agents have bounded radius of influence, global information cascades become unlikely, networks prone to cascades split into a multitude of local ones, and, by observing these local cascades, an outside observer such as social media platform can discover the ground truth.

Papers [18, 19] deal with information flow from the informed part of society to the uninformed one through a multistage strategic interaction. These papers are based on my Ph.D. thesis. The interaction is modeled as a repeated zero-sum game, where one player is informed of a payoff-relevant state and her opponent is not. By observing the actions of the informed player, the uninformed player can try guessing the state. So the information leaks through actions and the informed side needs to balance between exploiting her information now and keeping the information advantage for future rounds. Information dissipation can be captured through the value that the informed side can extract from her information as a function of the game length. In [18], I characterize those games where the value of information remains bounded. The paper [19] explores the opposite extreme and quantifies how fast the value of information can grow via a family of entropy-like functionals of the prior. Both papers rely on a connection between such repeated games and martingale-optimization problems representing the optimal information-revelation process. These optimization problems, in their turn, are connected to optimal transportation.

3. MECHANISMS AND ALGORITHMS FOR FAIR AND EFFICIENT RESOURCE ALLOCATION

Classical economic design is prone to impossibility results, the most famous of which are the Gibbard-Satterthwaite and Arrow's theorems. Ideal mechanisms satisfying modest requirements of efficiency, fairness, and non-manipulability usually fail to exist and the classical theory provides no recipes 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. My papers [2, 17, 4, 11, 3, 14, 15, 16, 10] reflect the new paradigm of economic design based on the interplay between ideas from economics and algorithmic game theory.² These ideas include the importance of small message spaces (concise bidding languages) and algorithmic complexity of a mechanism, robustness with respect to modeling assumptions, quantitative relaxation of requirements via approximation ratios, and a focus on the behavior of a mechanism on typical inputs instead of the worst-case ones.

3.1. Non-similarity of allocation problems with goods and bads. In everyday life, we often decide how to allocate bads, not goods. Think of distributing administrative chores in an academic department, or house chores among family members. However, in contrast to the well-studied case of goods, problems with bads have not received much attention. Apparently, this has occurred due to the wrong intuition that a simple change of the sign in one formula is enough to translate results from goods to bads.

Our paper [2] published in *Econometrica* initiated a systematic study of fair division of bads, and of models with a mixture of goods and bads. In this paper, we assume general homothetic preferences; a paper [17] published in *Social Choice and Welfare* refines the results in the domain of linear preferences that is most important for practice, e.g., the domain that is relevant to Spliddit.org.

We find that problems with goods are not equivalent to those with bads. This structural difference leads to new obstacles and is reflected in new impossibility results which are specific to the model with bads. Moreover, arguably the best approach to allocating goods without money — the pseudo-market approach, or the competitive equilibrium with equal incomes (CEEI) — is not straightforward to extend from goods to bads. In the model with goods, CEEI is a single-valued rule which is envy-free, efficient, and solves a convex optimization problem of Nash Social Welfare maximization. This is no longer true of the model with bads.

We extend CEEI to economies with a mixture of goods and bads. The outcome remains envy-free and efficient; however, the set of competitive equilibria becomes disconnected with a potentially exponential number of connected components. The components 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. The non-convexity of this problem requires new algorithmic insights discussed next.

3.2. Fair division algorithms. To use a mechanism in practice, its outcome must be efficiently computable. This concern has inspired my works on algorithmic questions [11, 3, 14].

In [11] (*R&R in Mathematics of Operations Research*), we propose the first algorithm for finding CEEI for bads. It runs in polynomial time if either the number of agents or of bads is fixed. The algorithm is based on the novel idea that the set of all Pareto-optimal demand structures (the so-called consumption graphs) has a polynomial size for almost all preference profiles, and can be efficiently enumerated. A similar idea is applicable to economies with goods, where it gives an algorithm substantially simpler in implementation than other known exact algorithms.

The technique of Pareto-frontier enumeration from [11] proves to be useful in other problems. It is the key technical tool in our paper [3] (*Operations Research*) proposing a new practical approach to fair division of valuable items. Think of siblings dividing inheritance including expensive real estate. The standard economic approach is to assume that items are divisible or to make them divisible via randomization. For example, a sibling may receive a house with a probability of 50%. Computer science literature usually focuses on indivisible

²The impact of ideas from computer science is reflected in my economic design lecture notes ([link](#)).

items and unavoidably replaces exact fairness requirements with their approximate versions to guarantee the existence of a fair allocation. For example, a sibling may receive an allocation that is envy-free up to one house. None of these two extremes are practical. We find a middle ground by treating the requirements of exact fairness and Pareto optimality as constraints and minimizing the number of fractionally allocated items. Using the technique of Pareto-frontier enumeration, we demonstrate that sharing minimization is computationally tractable except for a zero-measure set of degenerate profiles. The advantages of this approach are demonstrated on synthetic and real-life data from [Spliddit.org](https://spliddit.org).

If 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 computer scientists, e.g., envy-freeness and proportionality up to one item. The paper [14] published in *Operations Research Letters* develops an algorithm for computing approximately-proportional Pareto optimal allocations for a mixture of indivisible goods and bads. Even the existence of such allocations was not known. The algorithm relies on trading cycles for finding fractional Pareto improvements (a side result in our paper [3]) and an extension of the Barman-Krishnamurthy rounding [23].

One methodological corollary of my algorithmic results is that the concept of fractional Pareto optimality introduced in [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].

3.3. Fairness and efficiency in other economic design settings. Allocation and other collective decisions are often made with only partial information about agents' preferences or even based solely on agents' types. My papers [4, 10, 16, 15] explore fairness and efficiency issues in such settings.

A paper [4] published in *Management Science* offers the first robust mechanism design results in fair division. We consider a designer who has only partial statistical information about agents' preferences and show how she can achieve fairness and maximize welfare no matter what the exact distribution is. In the case of goods, we get a one-parametric family of undominated mechanisms, while for bads, such mechanism is unique (a repercussion of goods/bads non-similarity as in Section 3.1). The analysis uncovers connections with information design and optimal transportation.

Agents may have richer preferences in the background than a mechanism allows them to report. If true preferences are cardinal and reported, ordinal, we end up with notions of ex-post and ex-ante efficiency. The first one simply ignores the cardinal counterpart and the second one relies on cardinal preferences but in a robust way, namely, by looking at all possible cardinal preferences compatible with the ordinal report. In the working paper [10], we explore those domains of ordinal preferences where the two notions coincide. This project was motivated by the abundance of unresolved conjectures regarding the best ordinal mechanisms; see Section 5.5 of my [lecture notes](#) for one of such conjectures.

Fairness issues similar to those in allocation problems arise in the context of classification and recommendation systems. There is growing evidence that such systems can discriminate users by protected attributes such as gender or race. This may happen as a result of biased training data or insufficient data on particular groups. The common wisdom in artificial intelligence suggests that fairness has to be made a separate design concern. In a paper [16] published in proceedings of *AAAI*, we consider a fairness constraint imposed on a self-interested party, e.g., a revenue-maximizing bank predicting the reliability of a borrower. We show that popular fairness notions may harm the disadvantaged protected group since it may be profitable for the decision-maker to reallocate the "price of fairness" to this group. We propose a concept of welfare-equalizing fairness which is free of this flaw. Our results bridge

the literature on fair classification with the egalitarian approach to fairness from welfare economics. The paper [16] is close in spirit to applications of the information design paper [6] discussed in Section 2.1.

Our paper [15] (*Journal of Artificial Intelligence Research*) deals with fairness and efficiency issues arising in the context of direct democracy. Citizens, or members of a large group (such as a blockchain ecosystem), need to decide directly on a sequence of relevant policy issues. Running a referendum on each of the issues is good at representing the opinion of each agent (fairness) but is a heavy burden on the population if the number of issues is large (inefficiency). Within a broad family of alternative mechanisms, we show that the sortition — sampling a uniformly random committee and letting members decide on each issue via the majority vote — achieves the best fairness/efficiency compromise. These results are related to modern ideas of liquid democracy and voting in metric spaces.

4. COMMON PATTERNS IN MECHANISM AND INFORMATION DESIGN

A recent methodological breakthrough in mechanism and information design is understanding the importance and the unifying role played by the majorization theory [26, 27, 28, 29].

My papers discussed above suggest that majorization is just one out of a list of related patterns. Majorization arises naturally in [1, 6, 9, 18] as well as connections to optimal transportation [8, 4, 18], mathematical tomography [6], and maximal flows [1, 11]. All these patterns can be seen as manifestations of a more fundamental connection with the structure of measures with given marginals.

Exploring this connection and its implications is a key part of my current agenda. The paper [8] that reduces multi-agent persuasion to transportation is a recent example.

A working paper [7] approaches auction design from the same perspective. It demonstrates that one can think of revenue-maximization in multi-item multi-bidder auctions as a continuous transportation problem studied by Beckmann [30]. This result has structural and algorithmic implications, and extends the celebrated connection between the multi-item monopolist problem and the Monge-Kantorovich optimal transportation discovered by [31] in the multi-bidder case.

The ongoing project [13] is inspired by the observation that feasibility questions in multi-agent information design under privacy constraints [6] and feasibility questions for reduced-form mechanisms in Bayesian mechanism design lead to similar majorization conditions. I show that this is not a coincidence and that there is a direct way to translate feasibility results from mechanism to information design and back. Roughly, one can think of a mechanism as the joint distribution of agents' types and an outcome-decision, and then interpret types as signals about the decision. As a result, agents' posterior beliefs coincide with their reduced-form mechanisms. This construction implies novel versions of Border's theorem in mechanism design, including those that were thought to be intractable [32]. It also provides alternative proofs for some of our information design results from [1] and [6].

Another ongoing project [5] offers a novel perspective on the classical question of consumer demand aggregation uncovering structures common to aggregation, information, and mechanism design. We show that aggregation of homothetic preferences boils down to taking convex combinations in the space of logarithmic expenditure functions. This observation relates aggregation to convexification in functional spaces, similar to convexification in Bayesian persuasion. Using this link, we characterize domains invariant with respect to aggregation and explore the structure of indecomposable preferences corresponding to extreme points. The results are applied to robust welfare analysis, the pseudo-market approach in fair division, and to identification questions.

REFERENCES

Selected Publications

- [1] Itai Arieli, Yakov Babichenko, Fedor Sandomirskiy, and Omer Tamuz “Feasible joint posterior beliefs” *Journal of Political Economy*, 2021, 129(9)
EC’20 best paper award ([link](#))
preprint: [arXiv:2002.11362](#),
[popular version](#) published in **SIGecom Exchanges**
- [2] 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](#)
- [3] Fedor Sandomirskiy, Erel Segal-Halevi “Fair division with minimal sharing”
Operations Research, 2022, 70(3):1293-1952,
preprint: [arXiv:1908.01669](#)
- [4] Anna Bogomolnaia, Herve Moulin, and Fedor Sandomirskiy “A simple Online Fair Division problem”
Management Science, 2022, 68(2):1174-1194
preprint [arXiv:1903.10361](#)

Working papers

- [5] Fedor Sandomirskiy and Philip Ushchev “The geometry of consumer preference aggregation”
Job Market Paper: [link](#)
- [6] Kevin He, Fedor Sandomirskiy, and Omer Tamuz “Private private information”
preprint: [arXiv:2112.14356](#)
extended abstract published at **EC’22**
- [7] Alexander Kolesnikov, Fedor Sandomirskiy, Aleh Tsyvinski, and Alexander P. Zimin “Beckmann’s approach to multi-item multi-bidder auctions”
preprint: [arXiv:2203.06837](#)
- [8] Itai Arieli, Yakov Babichenko, and Fedor Sandomirskiy “Persuasion as transportation”
preprint: [link](#)
extended abstract published at **EC’22**
- [9] Itai Arieli, Yakov Babichenko, and Fedor Sandomirskiy “Bayesian persuasion with mediators”
preprint: [arXiv:2203.04285](#)
- [10] Federico Echenique, Joseph Root, Fedor Sandomirskiy “Efficiency in random resource allocation and social choice”
preprint: [arXiv:2203.06353](#)
- [11] Simina Brânzei, Fedor Sandomirskiy “Algorithms for Competitive Division of Chores”
R&R in Mathematics of Operations Research
preprint: [arXiv:1907.01766](#)
- [12] Itai Arieli, Fedor Sandomirskiy, and Rann Smorodinsky “On social networks that support learning”
R&R in Journal of Economic Theory
preprint: [arXiv:2011.05255](#)
extended abstract published at **EC’21**
- [13] Fedor Sandomirskiy “From information design to mechanism design and back”
on arXiv soon

Other publications

- [14] Haris Aziz, Herve Moulin, Fedor Sandomirskiy “A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation”
Operations Research Letters, 2020, 48(5):573-578
preprint: [arXiv:1909.00740](#)
- [15] Reshef Meir, Fedor Sandomirskiy, and Moshe Tennenholtz “Representative Committees of Peers”
Journal of Artificial Intelligence Research, 2021, 71:401-429
preprint: [arXiv:2006.07837](#)
- [16] Omer Ben-Porat, Fedor Sandomirskiy, Moshe Tennenholtz “Protecting the Protected Group: Circumventing Harmful Fairness”
proceedings of AAAI-21 (The Thirty-Fifth AAAI Conference on Artificial Intelligence), 2021
preprint: [arXiv:1905.10546](#)

- [17] Anna Bogomolnaia, Herve Moulin, Fedor Sandomirskiy, Elena Yanovskaya “Dividing goods or bads under additive utilities”
Social Choice and Welfare, 2019, 52(3):395-417
preprint: [arXiv:1608.01540](https://arxiv.org/abs/1608.01540)
- [18] Fedor Sandomirskiy “On repeated zero-sum games with incomplete information and asymptotically bounded values”
Dynamic Games and Applications, 2018, 8(1):180-198
preprint: [arXiv:1509.01727](https://arxiv.org/abs/1509.01727)
- [19] Fedor Sandomirskiy “Repeated games of incomplete information with large sets of states”
International Journal of Game Theory, 2014, 43(4):767-789
preprint: [arXiv:1205.6791](https://arxiv.org/abs/1205.6791)
- [20] Alexander Fedotov, Fedor Sandomirskiy “An exact renormalization formula for the Maryland model”
Communications in Mathematical Physics, 2015, 334(2):1083-1099
preprint: [arXiv:1311.6930](https://arxiv.org/abs/1311.6930)

Other references

- [21] E. Kamenica and M. Gentzkow “Bayesian persuasion” *American Economic Review*, 2011, 101(6):2590-2615
- [22] P. Dworzak and A. Kolotilin “The persuasion duality” *arXiv preprint arXiv:1910.11392*, 2019
- [23] S. Barman, SK. Krishnamurthy “On the proximity of markets with integral equilibria” *In Proceedings of the AAAI Conference on Artificial Intelligence 2019*, 33:1748-1755
- [24] S. Barman, SK. Krishnamurthy, R. Vaish “Finding fair and efficient allocations” *In Proceedings of the 2018 ACM Conference on Economics and Computation*, 2018, 557-574
- [25] B. de Keijzer, S. Bouveret, T. Klos, Y. Zhang “On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences” *In International Conference on Algorithmic Decision Theory*, 2009, 98-110
- [26] A. Kleiner, B. Moldovanu, and P. Strack “Extreme points and majorization: Economic applications” *Econometrica*, 2021, 89(4):1557-1593
- [27] A. Nikzad “Constrained majorization: Applications in mechanism design” *In Proceedings of the 23rd ACM Conference on Economics and Computation*, 2022, 330-331
- [28] O. Candogan, P. Strack “Optimal disclosure of information to a privately informed receiver” *arXiv preprint arXiv:2101.10431*, 2021
- [29] I. Arieli, Y. Babichenko R. Smorodinsky, T. Yamashita “Optimal persuasion via bi-pooling” *SSRN 3511516*, 2019
- [30] M. Beckmann “A continuous model of transportation. *Econometrica*, 1952, 643-660
- [31] C. Daskalakis, A. Deckelbaum, and C. Tzamos “Strong duality for a multiple-good monopolist” *Econometrica*, 2017, 85(3):735-767
- [32] P. Gopalan, N. Nisan, T. Roughgarden “Public projects, boolean functions, and the borders of Border’s theorem” *ACM Transactions on Economics and Computation (TEAC)*, 2018, 6(3-4):1-21.