The Pareto frontier for random mechanisms

Mennle, Timo ; Seuken, Sven

Abstract: In many situations, a group of individuals (called agents) must collectively decide on one of several alternatives, e.g., elect the next president. Ordinal mechanisms are systematic procedures to make such decisions based on the agents' preference orders over the alternatives. A mechanism is strategyproof if it makes truthful reporting of preferences a dominant strategy. Strategyproofness is therefore the “gold standard” among the incentive concepts. However, the seminal impossibility result of Gibbard [1977] showed that strategyproofness also greatly restricts the design space of ordinal mechanisms even if they can use randomization. In particular, it is incompatible with many other common desiderata, such as Condorcet consistency, stability, or egalitarian fairness. Thus, trade-offs between strategyproofness and other desiderata are necessary. In this paper, we study these trade-offs. We use approximate strategyproofness to define manipulability, a measure to quantify the incentive properties of non-strategyproof mechanisms, and we introduce deficit, a measure to quantify the performance of mechanisms with respect to another desideratum. A mechanism that minimizes the deficit subject to a particular bound on manipulability is called optimal at this bound; and the mechanisms that are optimal at some bound form the Pareto frontier. Our main contribution is a structural characterization of this Pareto frontier: we show that there exists a finite set of supporting manipulability bounds, such that it suffices to identify optimal mechanisms at each of them. Other mechanisms along the Pareto frontier can then be constructed as hybrids (i.e., convex combinations) of these optimal mechanisms. This allows a concise representation of the Pareto frontier in terms of a finite number of optimal mechanisms and their hybrids. In combination with linear programming, we can exploit this characterization to compute the whole Pareto frontier algorithmically. To illustrate its shape, we apply our results to determine the Pareto frontier for two different desiderata, namely Plurality and Veto scoring, in settings with 3 alternatives and up to 18 agents.

DOI: <https://doi.org/10.1145/2940716.2940786>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-130376>

Conference or Workshop Item

Accepted Version

Originally published at:

Mennle, Timo; Seuken, Sven (2016). The Pareto frontier for random mechanisms. In: The 2016 ACM Conference on Economics and Computation, Maastricht, The Netherlands, 24 August 2016 - 28 August 2016. ACM Press, 769.

DOI: <https://doi.org/10.1145/2940716.2940786>

The Pareto Frontier for Random Mechanisms [Extended Abstract]

TIMO MENNLE, University of Zurich
SVEN SEUKEN, University of Zurich

In many situations, a group of individuals (called *agents*) must collectively decide on one of several alternatives, e.g., elect the next president. *Ordinal mechanisms* are systematic procedures to make such decisions based on the agents' preference orders over the alternatives. A mechanism is *strategyproof* if it makes truthful reporting of preferences a dominant strategy. Strategyproofness is therefore the "gold standard" among the incentive concepts. However, the seminal impossibility result of Gibbard [1977] showed that strategyproofness also greatly restricts the design space of ordinal mechanisms even if they can use randomization. In particular, it is incompatible with many other common desiderata, such as Condorcet consistency, stability, or egalitarian fairness. Thus, trade-offs between strategyproofness and other desiderata are necessary.

In this paper, we study these trade-offs. We use approximate strategyproofness to define *manipulability*, a measure to quantify the incentive properties of non-strategyproof mechanisms, and we introduce *deficit*, a measure to quantify the performance of mechanisms with respect to another desideratum. A mechanism that minimizes the deficit subject to a particular bound on manipulability is called *optimal* at this bound; and the mechanisms that are optimal at some bound form the *Pareto frontier*. Our main contribution is a structural characterization of this Pareto frontier: we show that there exists a finite set of *supporting manipulability bounds*, such that it suffices to identify optimal mechanisms at each of them. Other mechanisms along the Pareto frontier can then be constructed as *hybrids* (i.e., convex combinations) of these optimal mechanisms.

This allows a concise representation of the Pareto frontier in terms of a finite number of optimal mechanisms and their hybrids. In combination with linear programming, we can exploit this characterization to compute the whole Pareto frontier algorithmically. To illustrate its shape, we apply our results to determine the Pareto frontier for two different desiderata, namely *Plurality* and *Veto* scoring, in settings with 3 alternatives and up to 18 agents.

References

Allan Gibbard. 1977. Manipulation of Schemes That Mix Voting with Chance. *Econometrica* 45, 3 (1977), 665–81.

The full version of this paper is available at <http://www.ifi.uzh.ch/ce/publications/EFF.pdf>

We would like to thank (in alphabetical order) Daniel Abaecherli, Haris Aziz, Gianluca Brero, Albin Erlanson, Bettina Klaus, Christian Kroer, Dmitry Moor, Ariel Procaccia, Tuomas Sandholm, Arunava Sen, Steffen Schuldenzucker, William Thomson, and Utku Ünver for helpful comments on this work. We also thank the participants of the Meeting of the COST Action on Computational Social Choice (Istanbul, 2015) and eight anonymous reviewers from EC'15 and EC'16 for their valuable comments. Any errors remain our own.

Part of this research was supported by the SNSF (Swiss National Science Foundation) under grant #156836. Author's addresses: Timo Mennle, Sven Seuken: Department of Informatics, University of Zurich, Switzerland, Email: {mennle,seuken}@ifi.uzh.ch.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

EC'16, July 24–28, 2016, Maastricht, The Netherlands.

ACM 978-1-4503-3936-0/16/07.

<http://dx.doi.org/10.1145/2940716.2940786>