

# The Pareto Frontier for Random Mechanisms\*

Timo Mennle                      Sven Seuken  
University of Zurich          University of Zurich

First version: February 13, 2015

This version: January 10, 2017

## Abstract

We study the trade-offs between strategyproofness and other desiderata, such as efficiency or fairness, that often arise in the design of random ordinal mechanisms. We use  $\varepsilon$ -approximate strategyproofness to define *manipulability*, a measure to quantify the incentive properties of non-strategyproof mechanisms, and we introduce the *deficit*, a measure to quantify the performance of mechanisms with respect to another desideratum. When this desideratum is incompatible with strategyproofness, mechanisms that trade off manipulability and deficit optimally form the *Pareto frontier*. Our main contribution is a structural characterization of this Pareto frontier, and we present algorithms that exploit this structure to compute it. To illustrate its shape, we apply our results for two different desiderata, namely Plurality and Veto scoring, in settings with 3 alternatives and up to 18 agents.

**Keywords:** Random Mechanisms, Pareto Frontier, Ordinal Mechanisms, Strategyproofness, Approximate Strategyproofness, Positional Scoring, Linear Programming

**JEL:** D71, D82

---

\* Department of Informatics, University of Zurich, Switzerland, Email: {mennle, seuken}@ifi.uzh.ch. Updates available from <https://www.ifi.uzh.ch/ce/publications/EFF.pdf>. We would like to thank (in alphabetical order) Daniel Abächerli, Haris Aziz, Gianluca Brero, Albin Erlanson, Bettina Klaus, Christian Kroer, Dmitry Moor, Ariel Procaccia, Tuomas Sandholm, Arunava Sen, Steffen Schudenzucker, William Thomson, and Utku Ünver for helpful comments on this work. Furthermore, we are thankful for the feedback we received from participants at COST COM-SOC Meeting 2016 (Istanbul, Turkey) and multiple anonymous referees at EC'15 and EC'16. Any errors remain our own. A 1-page abstract based on this paper has been published in the conference proceedings of EC'16 (Mennle and Seuken, 2016). Part of this research was supported by the SNSF (Swiss National Science Foundation) under grant #156836.

# 1. Introduction

In many situations, a group of individuals has to make a collective decision by selecting an alternative: *who should be the next president? who gets a seat at which public school? or where to build a new stadium?* *Mechanisms* are systematic procedures to make such decisions. Formally, a mechanism collects the individuals' (or *agents*') preferences and selects an alternative based on this input. While the goal is to select an alternative that is desirable for society as a whole, it may be possible for individual agents to manipulate the outcome to their own advantage by being insincere about their preferences. However, if these agents are lying about their preferences, the mechanism will have difficulty in determining an outcome that is desirable with respect to the true preferences. Therefore, incentives for truth-telling are a major concern in mechanism design. In this paper, we consider *ordinal* mechanisms with *random outcomes*. Such mechanisms collect preference orders and select lotteries over alternatives. We formulate our results for the full domain where agents can have any weak or strict preferences, but our results continue to hold on many domain restrictions, including the domain of strict preferences, the assignment domain, and the two-sided matching domain.

## 1.1. The Curse of Strategyproofness

*Strategyproof* mechanisms make truthful reporting a dominant strategy for all agents and it is therefore the “gold standard” among the incentive concepts. However, the seminal impossibility results by [Gibbard \(1973\)](#) and [Satterthwaite \(1975\)](#) established that if there are at least three alternatives and all strict preference orders are possible, then the only unanimous and deterministic mechanisms that are strategyproof are *dictatorships*. [Gibbard \(1977\)](#) extended these insights to mechanisms that involve randomization and showed that all strategyproof random mechanisms are probability mixtures of strategyproof unilateral and strategyproof duple mechanisms. These results greatly restrict the design space of strategyproof mechanisms. In particular, many common desiderata are incompatible with strategyproofness, such as Condorcet consistency, stability of matchings, or egalitarian fairness.

When a desideratum is incompatible with strategyproofness, designing “ideal” mechanisms is impossible. For example, suppose that our goal is to select an alternative that is the first choice of as many agents as possible. The *Plurality* mechanism selects an

alternative that is the first choice for a maximal number of agents. Thus, it achieves our goal perfectly. At the same time, Plurality and any other mechanism that achieves this goal perfectly must be manipulable. In contrast, the *Random Dictatorship* mechanism selects the first choice of a random agent. It is strategyproof but there is a non-trivial probability that it selects “wrong” alternatives. If Plurality is “too manipulable” and Random Dictatorship is “wrong too often,” then trade-offs are necessary. In this paper, we study mechanisms that make *optimal trade-offs* between (non-)manipulability and another desideratum. Such mechanisms form the *Pareto frontier* as they achieve the desideratum as well as possible, subject to a given limit on manipulability.

## 1.2. Measuring Manipulability and Deficit

In order to understand these trade-offs formally, we need measures for the performance of mechanisms in terms of incentives and in terms of the desideratum.

**Approximate Strategyproofness and Manipulability** A strategyproof mechanism does not allow agents to obtain a strictly positive gain from misreporting. *Approximate strategyproofness* is a relaxation of strategyproofness: instead of requiring the gain from misreporting to be non-positive,  $\varepsilon$ -approximate strategyproofness imposes a small (albeit positive) upper bound  $\varepsilon$  on this gain. The economic intuition is that if the potential gain is small, then agents might stick to truthful reporting (e.g., due to cognitive cost). To obtain a notion of approximate strategyproofness for ordinal mechanisms, we follow earlier work and consider agents with von Neumann-Morgenstern utility functions that are bounded between 0 and 1 (Birrell and Pass, 2011; Carroll, 2013). This allows the formulation of a parametric measure for incentives: the *manipulability*  $\varepsilon(\varphi)$  of a mechanism  $\varphi$  is the smallest bound  $\varepsilon$  for which  $\varphi$  is  $\varepsilon$ -approximately strategyproof.

**Desideratum Functions and Deficit** By allowing higher manipulability, the design space of mechanisms grows. This raises the question how the new freedom can be harnessed to improve the performance of mechanisms with respect to other desiderata. To measure this performance we introduce the *deficit*.

We first illustrate its construction via an example: consider the desideratum to select an alternative that is the first choice of as many agents as possible. For any alternative  $a$

it is natural to think of “the share of agents for whom  $a$  is the first choice” as the *value* that society derives from selecting  $a$ . For some alternative ( $b$ , say) this value is maximal. Then the *deficit of  $a$*  is the difference between the share of agents with first choice  $b$  and the share of agents with first choice  $a$ . A Dictatorship mechanism may select alternatives with strictly positive deficit. We define the *deficit* of this mechanism to be the highest deficit of any alternatives that it selects across all possible preference profiles.

In general, any notion of deficit is constructed in three steps: first, we express the desideratum via a *desideratum function*, which specifies the value to society from selecting a given alternative at a given preference profile. Since we consider random mechanisms, we extend desideratum functions to random outcomes by taking expectations.<sup>1</sup> Second, we define the deficit of outcomes based on the desideratum function. Third, we define the deficit of mechanisms based on the deficit of outcomes.<sup>2</sup> For any mechanism  $\varphi$ , we denote its deficit by  $\delta(\varphi)$ .

### 1.3. Optimal Mechanisms and the Pareto Frontier

Together, manipulability  $\varepsilon(\varphi)$  and deficit  $\delta(\varphi)$  yield a way to compare different mechanisms in terms of incentives and their performance with respect to some desideratum. Specifically, the *signature* of a mechanism  $\varphi$  is the point  $(\varepsilon(\varphi), \delta(\varphi))$  in the unit square  $[0, 1] \times [0, 1]$ . Figure 1 illustrates this comparison. Ideally, a mechanism would be strategyproof *and* would always select the most desirable alternatives. This corresponds to a signature in the origin  $(0, 0)$ . However, for desiderata that are incompatible with strategyproofness, designing ideal mechanisms is not possible. Instead, there exist strategyproof mechanisms which have a non-trivial deficit, such as Random Dictatorship; and there exist value maximizing mechanisms which have non-trivial manipulability, such as

---

<sup>1</sup>Many important desiderata can be expressed via desideratum functions, including binary desiderata such as *unanimity*, *Condorcet consistency*, *egalitarian fairness*, *Pareto optimality*, *v-rank efficiency* of assignments, *stability* of matchings, any desideratum specified via a *target mechanism* (or a *target correspondence*), or any *logical combination* of these. Moreover, it is possible to express quantitative desiderata, such as *maximizing positional score* in voting, *maximizing v-rank value* of assignments, or *minimizing the number of blocking pairs* in matching. We discuss the generality and limitations of desideratum functions in Section E.

<sup>2</sup>In the example we considered *absolute* differences to construct the deficit of outcomes and we considered the *worst-case* deficit to construct the deficit of mechanisms. Alternatively, we can consider *relative* differences and we can also define an *ex-ante* deficit of mechanisms. Many meaningful notions of deficit can be constructed in this way (see Appendix A); the results in this paper hold for any of them.

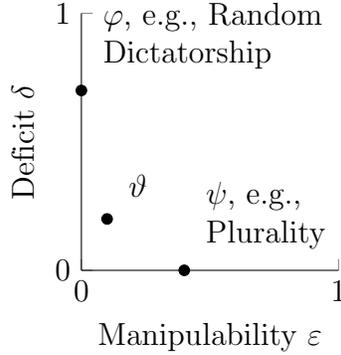


Figure 1: Example signatures of mechanisms in a manipulability-deficit-plot.

Plurality (if the goal is to select an alternative that is the first choice of as many agents as possible). Finally, there may exist mechanisms that have intermediate signatures, such as  $\vartheta$  in Figure 1. Choosing between these mechanisms means to make trade-offs.

**Finding Optimal Mechanisms** Naturally, we want to make *optimal* trade-offs. A mechanism is *optimal at manipulability bound*  $\varepsilon$  if it has manipulability of at most  $\varepsilon$  and it has the lowest deficit among all such mechanisms. For a given  $\varepsilon$  we denote by  $\text{OPT}(\varepsilon)$  the set of all mechanisms that are optimal at  $\varepsilon$ . Given an optimal mechanism, it is not possible to reduce the deficit without increasing manipulability at the same time. In this sense, the set of all optimal mechanisms constitutes the *Pareto frontier*. Our first result yields a finite set of linear constraints that is equivalent to  $\varepsilon$ -approximate strategyproofness. This equivalence allows us to formulate the linear program  $\text{FINDOPT}$ , whose solutions uniquely identify the optimal mechanisms at  $\varepsilon$ .

**Trade-offs via Hybrid Mechanisms** Given two mechanisms, *mixing* them suggests itself as a natural approach to create new mechanisms with intermediate signatures. Formally, the  $\beta$ -*hybrid* of two mechanisms  $\varphi$  and  $\psi$  is the  $\beta$ -convex combination of the two mechanisms. Such a hybrid can be implemented by first collecting the agents' preference reports, then randomly deciding to use  $\psi$  or  $\varphi$  with probabilities  $\beta$  and  $1 - \beta$ , respectively. If  $\varphi$  has lower manipulability and  $\psi$  has lower deficit, then one would expect their hybrid to inherit a share of both properties. Our second result in this paper formalizes this intuition: we prove that the signature of a  $\beta$ -hybrid is always *weakly*

*preferable* on both dimensions to the  $\beta$ -convex combination of the signatures of the two original mechanisms. This insight teaches us that interesting intermediate mechanisms can indeed be obtained via mixing.

Our result has important consequences for the Pareto frontier: it implies that the first unit of manipulability that we sacrifice yields the greatest return in terms of a reduction of deficit. Furthermore, the marginal return on further relaxing incentive requirements decreases as the mechanisms become more manipulable. This is good news for mechanism designers because it means that the most substantial improvements already arise by relaxing strategyproofness just “a little bit.”

**Structural Characterization of the Pareto Frontier** To fully understand the possible and necessary trade-offs between manipulability and deficit, we need to identify the whole Pareto frontier across *all* manipulability bounds. Our main result in this paper is a structural characterization of this Pareto frontier. We show that there exists a finite set  $\varepsilon_0 < \dots < \varepsilon_K$  of *supporting manipulability bounds*, such that between any two of them ( $\varepsilon_{k-1}$  and  $\varepsilon_k$ , say) the Pareto frontier consists precisely of the hybrids of two mechanisms that are optimal at  $\varepsilon_{k-1}$  and  $\varepsilon_k$ , respectively. Consequently, the two building blocks of the Pareto frontier are, first, the optimal mechanisms at the supporting manipulability bounds  $\varepsilon_k$  and, second, the hybrids of optimal mechanisms at adjacent supporting manipulability bounds for any intermediate  $\varepsilon \neq \varepsilon_k$ . Thus, the Pareto frontier can be represented concisely in terms of a finite number of optimal mechanisms and their hybrids. In combination with the linear program FINDOPT, we can exploit this characterization to compute the whole Pareto frontier algorithmically.

In summary, we provide a novel perspective on the possible and necessary trade-offs between incentives and other desiderata. Our results unlock the Pareto frontier of random mechanisms to analytic, axiomatic, and algorithmic exploration.

## 2. Related Work

Severe impossibility results restrict the design of strategyproof ordinal mechanisms. The seminal Gibbard-Satterthwaite Theorem (Gibbard, 1973; Satterthwaite, 1975) established that if all strict preferences over at least 3 alternatives are possible, then the only unanimous, strategyproof, and deterministic mechanisms are dictatorial, and Gibbard

(1977) extended this result to random mechanisms. Thus, many important desiderata are incompatible with strategyproofness, such as selecting a Condorcet winner or maximizing Borda count (Pacuit, 2012). Similar restrictions persist in other domains: in the random assignment problem, strategyproofness is incompatible with rank efficiency (Featherstone, 2011), and in the two-sided matching problem, strategyproofness is incompatible with stability (Roth, 1982).

Many research efforts have been made to circumvent these impossibility results and to obtain better performance on other dimensions. One way to reconcile strategyproofness with other desiderata is to consider restricted domains. Moulin (1980) showed that in the single-peaked domain, all strategyproof, anonymous, and efficient mechanisms are variants of the Median mechanism with additional virtual agents, and Ehlers, Peters and Storcken (2002) extended this result to random mechanisms. Chatterji, Sanver and Sen (2013) showed that a *semi-single-peaked* structure is *almost* the defining characteristic of domains that admit the design of strategyproof deterministic mechanisms with appealing properties; an analogous result for random mechanisms is outstanding.

An alternative way to circumvent impossibility results is to continue working in full domains but to relax the strategyproofness requirement “a little bit.” This can enable the design of mechanisms that come closer to achieving a given desideratum but still have appealing (albeit imperfect) incentive properties. Mennle and Seuken (2017b) introduced *partial strategyproofness*, a relaxation of strategyproofness that has particular appeal in the assignment domain. Azevedo and Budish (2015) proposed *strategyproofness in the large*, which requires that the incentives for any individual agent to misreport should vanish in large markets. However, strategyproofness in the large is unsuited for the exact analysis of finite settings which we perform in this paper. Instead, we follow Birrell and Pass (2011) and Carroll (2013), who used *approximate strategyproofness for agents with bounded vNM utility functions* to quantify manipulability of non-strategyproof ordinal mechanisms and derived asymptotic results. We also use approximate strategyproofness but we give exact results for finite settings.

Some prior work has considered trade-offs explicitly. Using efficiency notions based on dominance relations, Aziz, Brandt and Brill (2013) and Aziz, Brandl and Brandt (2014) proved compatibility and incompatibility of various combinations of incentive and efficiency requirements. Procaccia (2010) considered an *approximation ratio* based on positional scoring and gave bounds on how well strategyproof random mechanisms

can approximate optimal positional scores as markets get large. While he found most of these to be inapproximable by strategyproof mechanisms, [Birrell and Pass \(2011\)](#) obtained positive limit results for the approximation of deterministic target mechanisms via *approximately* strategyproof random mechanisms. In [\(Mennle and Seuken, 2017a\)](#) we studied how hybrid mechanisms make trade-offs between strategyproofness (in terms of the *degree of strategyproofness*) and efficiency (in terms of *dominance*) in random assignment. While hybrid mechanisms also play a central role in our present paper, we consider general ordinal mechanisms and different measures (i.e., *manipulability* and *deficit*).

### 3. Formal Model

Let  $N$  be a set of  $n$  agents and  $M$  be a set of  $m$  alternatives, where the tuple  $(N, M)$  is called a *setting*. Each agent  $i \in N$  has a *preference order*  $P_i$  over alternatives, where  $P_i : a \geq b$ ,  $P_i : a > b$ , and  $P_i : a \sim b$  denote *weak* preference, *strict* preference, and *indifference*, respectively, and  $\mathcal{P}$  denotes the set of all preference orders. For agent  $i$ 's preference order  $P_i$ , the *rank* of alternative  $j$  under  $P_i$  is the number of alternatives that  $i$  strictly prefers to  $j$  plus 1, denoted  $\text{rank}_{P_i}(j)$ .<sup>3</sup> A *preference profile*  $\mathbf{P} = (P_i, P_{-i})$  is a collection of preference orders from all agents, and  $P_{-i}$  are the preference orders of all other agents, except  $i$ . A (*random*) *mechanism* is a mapping  $\varphi : \mathcal{P}^N \rightarrow \Delta(M)$ . Here  $\Delta(M)$  is the space of lotteries over alternatives, and any  $x \in \Delta(M)$  is called an *outcome*.

We extend agents' preferences over alternatives to preferences over lotteries via von Neumann-Morgenstern utility functions: each agent  $i \in N$  has a *utility function*  $u_i : M \rightarrow [0, 1]$  that represents their preference order, i.e.,  $u_i(a) \geq u_i(b)$  holds whenever  $P_i : a \geq b$ . Note that utility functions are bounded between 0 and 1, so that the model admits a non-degenerate notion of approximate strategyproofness (see Remark 2). We denote the set of all utility functions that represent the preference order  $P_i$  by  $U_{P_i}$ .

*Remark 1.* We formulate our results for the full domain but they naturally extend to a variety of domains, including the domain of strict preferences, the assignment domain, and the two-sided matching domain (see Appendix E).

---

<sup>3</sup>1 is added to ensure that first choice alternatives have rank 1, not 0.

## 4. Approximate Strategyproofness and Manipulability

Our goal in this paper is to study mechanisms that trade off manipulability and other desiderata optimally. For this purpose we need measures for the performance of different mechanisms with respect to the two dimensions of this trade-off. In this section, we review approximate strategyproofness, derive a measure for the incentive properties of non-strategyproof mechanisms, and present our first main result.

### 4.1. Strategyproofness and Approximate Strategyproofness

The most demanding incentive concept is strategyproofness. It requires that truthful reporting is a dominant strategy for all agents. For random mechanisms, this means that truthful reporting always maximizes any agent's *expected utility*.

**Definition 1** (Strategyproofness). Given a setting  $(N, M)$ , a mechanism  $\varphi$  is *strategyproof* if for any agent  $i \in N$ , any preference profile  $(P_i, P_{-i}) \in \mathcal{P}^N$ , any utility  $u_i \in U_{P_i}$  that represents  $P_i$ , and any misreport  $P'_i \in \mathcal{P}$ , we have

$$\sum_{j \in M} u_i(j) \cdot (\varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i})) \leq 0. \quad (1)$$

The left side of (1) is the change in its own expected utility that  $i$  can affect by falsely reporting  $P'_i$  instead of reporting  $P_i$ . For later use, we denote this difference by

$$\varepsilon(u_i, (P_i, P_{-i}), P'_i, \varphi) = \sum_{j \in M} u_i(j) \cdot (\varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i})). \quad (2)$$

The fact that  $\varepsilon(u_i, (P_i, P_{-i}), P'_i, \varphi)$  is upper bounded by 0 for strategyproof mechanisms means that deviating from the true preference report weakly decreases expected utility for any agent in any situation, independent of the other agents' reports. Conversely, if a mechanism  $\varphi$  is *not* strategyproof, there necessarily exists at least one situation in which  $\varepsilon(u_i, (P_i, P_{-i}), P'_i, \varphi)$  is strictly positive. Imposing a different bound leads to the notion of approximate strategyproofness (Birrell and Pass, 2011; Carroll, 2013).

**Definition 2** ( $\varepsilon$ -Approximate Strategyproofness). Given a setting  $(N, M)$  and a bound  $\varepsilon \in [0, 1]$ , a mechanism  $\varphi$  is  $\varepsilon$ -*approximately strategyproof* if for any agent  $i \in N$ ,

any preference profile  $(P_i, P_{-i}) \in \mathcal{P}^N$ , any utility  $u_i \in U_{P_i}$  that represents  $P_i$ , and any misreport  $P'_i \in \mathcal{P}$ , we have

$$\varepsilon(u_i, (P_i, P_{-i}), P'_i, \varphi) = \sum_{j \in M} u_i(j) \cdot (\varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i})) \leq \varepsilon. \quad (3)$$

This definition is analogous to Definition 1 of strategyproofness, except that the upper bound in (3) is  $\varepsilon$  instead of 0. Thus, 0-approximate strategyproofness coincides with strategyproofness. Furthermore, the gain never exceeds 1, which makes 1-approximate strategyproofness a void constraint that is trivially satisfied by any mechanism.

The interpretation of intermediate values of  $\varepsilon \in (0, 1)$  is more challenging. Unlike utilities in quasi-linear domains, vNM utilities are not comparable across agents. Thus, we cannot simply think of  $\varepsilon$  as the “value” (e.g., in dollars) that an agent can gain by misreporting. Instead,  $\varepsilon$  is a *relative* bound: since  $u_i$  is between 0 and 1, a change of magnitude 1 in expected utility corresponds to the selection of an agent’s first choice instead of that agent’s last choice. Thus, “1” is the *maximal gain* from misreporting that *any* agent could obtain under an *arbitrary* mechanism. The bound  $\varepsilon$  is the *share of this maximal gain* by which any agent can at most improve its expected utility under an  $\varepsilon$ -approximately strategyproof mechanism.

*Remark 2.* The bounds on utilities are essential for  $\varepsilon$ -approximate strategyproofness to be a useful relaxation of strategyproofness for ordinal mechanisms. Suppose that a non-strategyproof mechanism  $\varphi$  allows a gain of  $\varepsilon(u_i, (P_i, P_{-i}), P'_i, \varphi) > 0$ . Then scaling the utility function  $u_i$  by a factor  $\alpha > 1$  results in a linear increase of this gain. Thus,  $\varepsilon$ -approximate strategyproofness for unbounded utilities would imply strategyproofness.

## 4.2. Manipulability

If  $\varphi$  is  $\varepsilon$ -approximately strategyproof, then it is also  $\varepsilon'$ -approximately strategyproof for any  $\varepsilon' \geq \varepsilon$ . Thus, lower values of  $\varepsilon$  correspond to stronger incentives. With this in mind, we define the manipulability of a mechanism.

**Definition 3** (Manipulability). Given a setting  $(N, M)$  and mechanism  $\varphi$ , the *manipulability of  $\varphi$  (in the setting  $(N, M)$ )* is given by

$$\varepsilon(\varphi) = \min\{\varepsilon' \in [0, 1] : \varphi \text{ is } \varepsilon'\text{-approximately strategyproof in } (N, M)\}. \quad (4)$$

Intuitively,  $\varepsilon(\varphi)$  is the lowest bound  $\varepsilon'$  for which  $\varphi$  is  $\varepsilon'$ -approximately strategyproof. This minimum is in fact attained because all inequalities from (3) are weak. Note that in a different setting  $(N', M')$ , the manipulability may vary. However, for all statements in this paper, a setting is held fixed and the value  $\varepsilon(\varphi)$  should be understood as the manipulability of the mechanism  $\varphi$  in the fixed setting from the respective context.

### 4.3. An Equivalent Condition for Approximate Strategyproofness

Recall that the definition of  $\varepsilon$ -approximate strategyproofness imposes an upper bound on the gain that agents can obtain by misreporting. In particular, inequality (3) must hold for all utility functions that represent the agent's preference order. Since there are infinitely many such utility functions, a naïve approach to verifying  $\varepsilon$ -approximate strategyproofness of a given mechanism would involve checking an infinite number of constraints. This is somewhat unattractive from an axiomatic perspective and even prohibitive from an algorithmic perspective. Fortunately, we can bypass this issue, as the next Theorem 1 shows.

**Theorem 1.** *Given a setting  $(N, M)$ , a bound  $\varepsilon \in [0, 1]$ , and a mechanism  $\varphi$ , the following are equivalent:*

1.  $\varphi$  is  $\varepsilon$ -approximately strategyproof in  $(N, M)$ .
2. For any agent  $i \in N$ , any preference profile  $(P_i, P_{-i}) \in \mathcal{P}^N$ , any misreport  $P'_i \in \mathcal{P}$ , and any rank  $r \in \{1, \dots, m\}$ , we have

$$\sum_{j \in M: \text{rank}_{P'_i}(j) \leq r} \varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i}) \leq \varepsilon. \quad (5)$$

*Proof Outline (formal proof in Appendix F.1).* The key idea is to represent any utility function as an element of the convex hull of a certain set of *extreme utility functions*. For any combination  $(i, (P_i, P_{-i}), P'_i, k)$  the inequality in statement (2) is just the  $\varepsilon$ -approximate strategyproofness constraints for one extreme utility function.  $\square$

Theorem 1 yields that  $\varepsilon$ -approximate strategyproofness can be equivalently expressed as a *finite* set of *weak, linear* inequalities. This has far-reaching consequences. In general, it unlocks approximate strategyproofness for use under the automated mechanism

design paradigm (Sandholm, 2003). Specifically, it enables our identification of optimal mechanisms as solutions to a particular linear program (Section 6).

## 5. Desideratum Functions and Deficit

While it is important to elicit truthful preferences, good incentives alone do not make a mechanism attractive. Instead, it should ultimately select *desirable* alternatives, where desirability depends on the agents' preferences. In this section, we introduce a formal method to quantify the ability of mechanisms to achieve a given desideratum.

### 5.1. Desideratum Functions

To express a desideratum formally, we define *desideratum functions*. These reflect the value that society derives from selecting a particular alternative when the agents have a particular preference profile.

**Definition 4** (Desideratum Function). A *desideratum function* is a mapping  $d : M \times \mathcal{P}^N \rightarrow [0, 1]$ , where  $d(j, \mathbf{P})$  is the *d-value* associated with selecting alternative  $j$  when the agents have preferences  $\mathbf{P}$ .

We illustrate how desideratum functions represent desiderata with two examples.

**Example 1.** Suppose that our goal is to select an alternative that is the first choice of as many agents as possible. We can define the corresponding desideratum function by setting  $d^{\text{PLU}}(j, \mathbf{P}) = n_j^1/n$ , where  $n_j^1$  is the number of agents whose first choice under  $\mathbf{P}$  is  $j$ . Note that  $d^{\text{PLU}}(j, \mathbf{P})$  is proportional to the *Plurality score* of  $j$  under  $\mathbf{P}$ .

**Example 2.** Alternatively, we can consider a binary desideratum. An alternative  $j$  is a *Condorcet winner at  $\mathbf{P}$*  if it dominates all other alternatives in a pairwise majority comparison. A mechanism is *Condorcet consistent* if it selects Condorcet winners whenever they exist. We can reflect this desideratum by setting  $d^{\text{CON}}(j, \mathbf{P}) = 1$  for any  $j$  that is a Condorcet winner at  $\mathbf{P}$ , and  $d^{\text{CON}}(j, \mathbf{P}) = 0$  otherwise.

Desideratum functions are extended to (random) outcomes by taking expectations.

**Definition 5** (Expected  $d$ -value). Given a desideratum function  $d$ , a preference profile  $\mathbf{P} \in \mathcal{P}^N$ , and an outcome  $x \in \Delta(M)$ , the (*expected*)  $d$ -value of  $x$  at  $\mathbf{P}$  is given by

$$d(x, \mathbf{P}) = \sum_{j \in M} x_j \cdot d(j, \mathbf{P}). \quad (6)$$

The interpretation of  $d(x, \mathbf{P})$  is straightforward: if  $d$  *quantifies* the value of alternatives (as in Example 1), then  $d(x, \mathbf{P})$  is the *expectation* of the societal value from selecting an alternative according to  $x$ . If  $d$  reflects a *binary desideratum* (as in Example 2), then  $d(x, \mathbf{P})$  is the *probability* of selecting an alternative with the desirable property.

*Remark 3.* By taking expectations, the  $d$ -value of random outcomes is fully determined by the  $d$ -values of the alternatives. This linear structure is a key ingredient to our results. In Appendix E, we show that many (but not all) popular desiderata admit such a representation, and we also discuss the limitations.

Ideally, mechanisms would always select outcomes that maximize the  $d$ -value.

**Definition 6** ( $d$ -maximizing). Given a desideratum function  $d$  and a preference profile  $\mathbf{P} \in \mathcal{P}^N$ , an outcome  $x \in \Delta(M)$  is  *$d$ -maximizing at  $\mathbf{P}$*  if  $d(x, \mathbf{P}) = \max_{j \in M} d(j, \mathbf{P})$ . A mechanism  $\varphi$  is  *$d$ -maximizing* if, for any  $\mathbf{P} \in \mathcal{P}^N$ ,  $\varphi(\mathbf{P})$  is  $d$ -maximizing at  $\mathbf{P}$ .

Note that for any preference profile  $\mathbf{P}$  there always exists at least one alternative that is  $d$ -maximizing at  $\mathbf{P}$  by construction. Furthermore, any  $d$ -maximizing random outcome must be a lottery over alternatives that are all  $d$ -maximizing at  $\mathbf{P}$ .

**Example 1, continued.** Recall the desideratum function  $d^{\text{PLU}}(j, \mathbf{P}) = n_j^1/n$ , where  $n_j^1$  is the number of agents who ranked  $j$  first under  $\mathbf{P}$ . A mechanism is  $d^{\text{PLU}}$ -maximizing if and only if it is a Plurality mechanism (i.e., a mechanism that selects only alternatives which are the first choice of a maximum number of agents).

**Example 2, continued.** Recall that  $d^{\text{CON}}(j, \mathbf{P}) = \mathbb{1}_{\{j \text{ Condorcet winners at } \mathbf{P}\}}$  expresses the desideratum to select Condorcet winners when they exist. Indeed, any  $d^{\text{CON}}$ -maximizing mechanism is Condorcet consistent. Moreover, at any preference profile  $\mathbf{P}$  where no Condorcet winner exists, the minimal and maximal achievable  $d$ -values are both zero. At these  $\mathbf{P}$ , the  $d^{\text{CON}}$ -maximizing mechanisms are therefore free to choose any outcome. Consequently, maximizing  $d^{\text{CON}}$ -value is *equivalent* to Condorcet consistency.

## 5.2. Deficit of Mechanisms

Intuitively, the deficit of an outcome is the loss that society incurs from choosing that outcome instead of a  $d$ -maximizing outcome.

**Definition 7** (Deficit of Outcomes). Given a desideratum function  $d$ , a preference profile  $\mathbf{P} \in \mathcal{P}^N$ , and an outcome  $x \in \Delta(M)$ , the  $d$ -deficit of  $x$  at  $\mathbf{P}$  is

$$\delta_d(x, \mathbf{P}) = \max_{j \in M} d(j, \mathbf{P}) - d(x, \mathbf{P}). \quad (7)$$

*Remark 4* (Relative Deficit). The difference in (7) is *absolute*; however, in some situations, it may be more natural to consider *relative* differences, such as the ratio between the achieved and the maximal achievable  $d$ -value. As we show in Appendix A.1, it is without loss of generality that we restrict our attention to absolute differences.

Equipped with the deficit of outcomes, we define the deficit of mechanisms. This measure is the “desideratum counterpart” to the measure  $\varepsilon(\varphi)$  for incentive properties.

**Definition 8** (Deficit of Mechanisms). Given a setting  $(N, M)$ , a desideratum function  $d$ , and a mechanism  $\varphi$ , the (*worst-case*)  $d$ -deficit of  $\varphi$  (in  $(N, M)$ ) is the highest  $d$ -deficit incurred by  $\varphi$  across all preference profiles; formally,

$$\delta_d^{\max}(\varphi) = \max_{\mathbf{P} \in \mathcal{P}^N} \delta_d(\varphi(\mathbf{P}), \mathbf{P}). \quad (8)$$

Intuitively, the deficit of the mechanism is determined by the most severe violation of the desideratum across all preference profiles. Thus, a mechanism with low deficit violates the desideratum only “a little bit” across all possible preference profiles.

*Remark 5.* Another way to define the deficit of mechanisms arises when agents’ preference profiles are drawn from a known distribution. In Appendix A.2 we define this *ex-ante deficit* formally. All results in this paper hold for both notions of deficit.

For the remainder of this paper, we fix an arbitrary setting  $(N, M)$ , a desideratum expressed via a desideratum function  $d$ , and a notion of deficit  $\delta$  derived from  $d$ . The triple  $(N, M, \delta)$  is called a *problem*. Unless explicitly stated otherwise, our results are understood to hold for any fixed problem  $(N, M, \delta)$  separately.

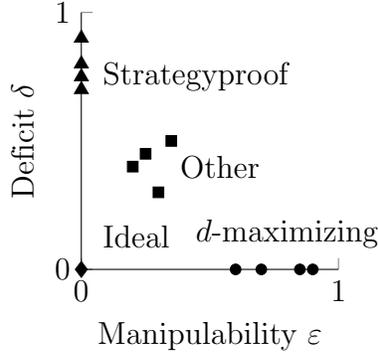


Figure 2: Manipulability-deficit-plot with mechanism signatures: strategyproof (triangles),  $d$ -maximizing (circles), ideal (diamond), others (squares).

## 6. Optimal Mechanisms

We are now in a position to formalize and study optimal mechanisms, which are mechanisms that trade off manipulability and deficit optimally.

### 6.1. Signatures of Mechanisms

To compare different mechanisms, we introduce signatures.

**Definition 9** (Signature). Given a problem  $(N, M, \delta)$  and a mechanism  $\varphi$ , the tuple  $(\varepsilon(\varphi), \delta(\varphi)) \in [0, 1] \times [0, 1]$  is called the *signature of  $\varphi$  (in the problem  $(N, M, \delta)$ )*.<sup>4</sup>

Signatures allow a convenient graphical representation of the performance of any mechanism in terms of manipulability and deficit. Figure 2 gives examples of such signatures: since 0-approximate strategyproofness is equivalent to strategyproofness, the signature of any strategyproof mechanism have an  $\varepsilon$ -component of 0. On the other hand, any  $d$ -maximizing mechanism have a signature with a  $\delta$ -component of 0. If an ideal mechanism exists, it has a signature in the origin  $(0, 0)$ . Mechanisms that are neither strategyproof nor  $d$ -maximizing have signatures in the half-open unit square  $(0, 1] \times (0, 1]$ .

<sup>4</sup>Since we fix a problem  $(N, M, \delta)$ , we refer to the tuple  $(\varepsilon(\varphi), \delta(\varphi))$  simply as the *signature of  $\varphi$* , keeping in mind that a mechanism's signature may be different for different problems.

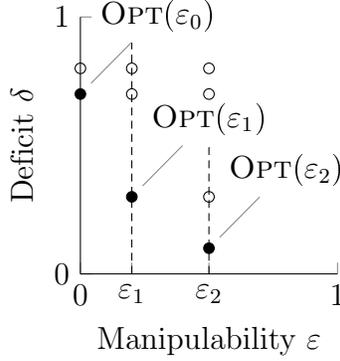


Figure 3: Example signatures of optimal mechanisms (filled circles) and non-optimal mechanisms (empty circles) at manipulability bounds  $\varepsilon_0, \varepsilon_1, \varepsilon_2$ .

## 6.2. Definition and Existence of Optimal Mechanisms

When impossibility results prohibit the design of ideal mechanisms, the decision in favor of any mechanism necessarily involves a trade-off between manipulability and deficit. To make an informed decision about this trade-off, a mechanism designer must be aware of the different design options. One straightforward approach to this problem is to decide on a maximal acceptable manipulability  $\varepsilon \in [0, 1]$  up front and use a mechanism that minimizes the deficit among all  $\varepsilon$ -approximately strategyproof mechanisms. We define  $\delta(\varepsilon) = \min \{\delta(\varphi) \mid \varphi \text{ } \varepsilon\text{-approximately SP}\}$  to be the lowest deficit that is achievable by any  $\varepsilon$ -approximately strategyproof mechanism.

**Definition 10** (Optimal Mechanism). Given a problem  $(N, M, \delta)$  and a bound  $\varepsilon \in [0, 1]$ , a mechanism  $\varphi$  is *optimal at  $\varepsilon$*  if  $\varphi$  is  $\varepsilon$ -approximately strategyproof and  $\delta(\varphi) = \delta(\varepsilon)$ . Denote by  $\text{OPT}(\varepsilon)$  the set of all mechanisms that are optimal at  $\varepsilon$ . Any optimal mechanism  $\varphi \in \text{OPT}(\varepsilon)$  is called a *representative of  $\text{OPT}(\varepsilon)$* .

Proposition 1 shows that optimal mechanisms always exist.

**Proposition 1.** *Given a problem  $(N, M, \delta)$  and a manipulability bound  $\varepsilon \in [0, 1]$ , there exists at least one mechanism that is optimal at  $\varepsilon$ .*

Existence follows via a compactness argument (see Appendix F.2 for a proof).

Proposition 1 justifies the use of the minimum (rather than the infimum) in the definition of  $\delta(\varepsilon)$ , since the deficit  $\delta(\varepsilon) = \delta(\varphi)$  is in fact attained by some mechanism. Figure 3 illustrates signatures of optimal and non-optimal mechanisms. On the vertical lines at each of the manipulability bounds  $\varepsilon_0 = 0, \varepsilon_1, \varepsilon_2$ , the circles (empty circles) correspond to signatures of non-optimal mechanisms. The signatures of optimal mechanisms (filled circles) from  $\text{OPT}(\varepsilon_k)$ ,  $k = 0, 1, 2$  take the lowest positions on the vertical lines.

### 6.3. Identifying Optimal Mechanisms

The existence proof for optimal mechanisms is implicit and does not provide a way of actually determining them. Our next result establishes a correspondence between the set  $\text{OPT}(\varepsilon)$  and the set of solutions to a linear program. We can solve this program algorithmically to find representatives of  $\text{OPT}(\varepsilon)$  and to compute  $\delta(\varepsilon)$ .

**Linear Program 1** (FINDOPT).

$$\begin{aligned}
& \text{minimize} && \mathbf{d} && && \text{(Objective)} \\
& \text{subject to} && \sum_{j \in M: \text{rank}_{P_i}(j) \leq k} \mathbf{f}_j(P'_i, P_{-i}) - \mathbf{f}_j(P_i, P_{-i}) \leq \varepsilon, && && (\varepsilon\text{-approximate SP}) \\
& && \forall i \in N, (P_i, P_{-i}) \in \mathcal{P}^N, P'_i \in \mathcal{P}, k \in \{1, \dots, m\} \\
& && \mathbf{d} \geq \max_{j \in M} d(j, \mathbf{P}) - \sum_{j \in M} \mathbf{f}_j(\mathbf{P}) \cdot d(j, \mathbf{P}), \quad \forall \mathbf{P} \in \mathcal{P}^N && && \text{(Deficit)} \\
& && \sum_{j \in M} \mathbf{f}_j(\mathbf{P}) = 1, \quad \forall \mathbf{P} \in \mathcal{P}^N && && \text{(Probability)} \\
& && \mathbf{f}_j(\mathbf{P}) \in [0, 1], \quad \forall \mathbf{P} \in \mathcal{P}^N, j \in M && && \text{(Outcome variables)} \\
& && \mathbf{d} \in [0, 1], && && \text{(Deficit variable)}
\end{aligned}$$

Each variable  $\mathbf{f}_j(\mathbf{P})$  corresponds to the probability with which the mechanism  $\varphi$  selects alternative  $j$  if the agents report preference profile  $\mathbf{P}$ . Consequently, any assignment of the variables  $\{\mathbf{f}_j(\mathbf{P}) : j \in M, \mathbf{P} \in \mathcal{P}^N\}$  corresponds to a unique mapping  $\varphi : M \times \mathcal{P}^N \rightarrow \mathbb{R}^M$ . The two constraints labeled *(Probability)* and *(Outcome variables)* ensure that the variable assignment in fact corresponds to a *mechanism* (rather than just a mapping from preference profiles to  $\mathbb{R}^M$ ). The additional variable  $\mathbf{d}$  represents this mechanism's deficit and the objective is to minimize its value.

The constraints labeled ( *$\varepsilon$ -approximate SP*) reflect the equivalent condition for  $\varepsilon$ -approximate strategyproofness that we obtained from Theorem 1. In combination, these constraints ensure that the mechanisms corresponding to the feasible variable assignments of FINDOPT are *exactly* the  $\varepsilon$ -approximately strategyproof mechanisms. The constraint labeled (*Deficit*) makes  $d$  an upper bound for the deficit of  $\varphi$ .<sup>5</sup>

The following Proposition 2 formalizes the correspondence between optimal mechanisms and solutions of the linear program FINDOPT.

**Proposition 2.** *Given a problem  $(N, M, \delta)$  and a bound  $\varepsilon \in [0, 1]$ , a variable assignment  $\{\mathbf{f}_j(\mathbf{P}) : j \in M, \mathbf{P} \in \mathcal{P}^N\}$  is a solution of FINDOPT at  $\varepsilon$  if and only if the mechanism  $\varphi$  defined by  $\varphi_j(\mathbf{P}) = \mathbf{f}_j(\mathbf{P})$  for all  $j \in M, \mathbf{P} \in \mathcal{P}^N$  is optimal at  $\varepsilon$ .*

The proof follows directly from the discussion above. One important consequence of Proposition 2 is that we can compute optimal mechanisms for any given problem  $(N, M, \delta)$  and any manipulability bound  $\varepsilon \in [0, 1]$ . Going back to the mechanism designer’s problem of trading off manipulability and deficit, we now have a way of determining optimal mechanisms for particular manipulability bounds  $\varepsilon$ . With FINDOPT we can evaluate algorithmically what deficit we must accept when manipulability must not exceed some fixed limit  $\varepsilon$ .

Shifting the burden of design to a computer by encoding good mechanisms in optimization problems is the central idea of *automated mechanism design* (Sandholm, 2003). A common challenge with this approach is that the optimization problem can become large and difficult to solve; and naïve implementations of FINDOPT will face this issue as well. Substantial run-time improvements are possible, e.g., by exploiting additional axioms such as anonymity and neutrality (Mennle, Abaecherli and Seuken, 2015). Nonetheless, determining optimal mechanisms remains a computationally expensive operation.

Computational considerations aside, Proposition 2 provides a new understanding of optimal mechanisms: since  $\text{OPT}(\varepsilon)$  corresponds to the solutions of the linear program FINDOPT, it can be interpreted as a convex polytope. In Section 8 we use methods from convex geometry to derive our structural characterization of the Pareto frontier. The representation of optimal mechanisms as solutions to FINDOPT constitutes the first building block of this characterization.

---

<sup>5</sup>Alternatively, we can implement the *ex-ante deficit* by replacing the (*Deficit*)-constraint (see Appendix A.2).

## 7. Hybrid Mechanisms

In this section, we introduce *hybrid mechanisms*, which are convex combinations of two *component mechanisms*. Intuitively, by mixing one mechanism with low manipulability and another mechanism with low deficit, we may hope to obtain new mechanisms with intermediate signatures. Initially, the construction of hybrids is independent of the study of optimal mechanisms. However, in Section 8, they will constitute the second building block for our structural characterization of the Pareto frontier.

**Definition 11** (Hybrid). For  $\beta \in [0, 1]$  and mechanisms  $\varphi$  and  $\psi$ , the  $\beta$ -hybrid  $h_\beta$  is given by  $h_\beta(\mathbf{P}) = (1 - \beta)\varphi(\mathbf{P}) + \beta\psi(\mathbf{P})$  for any preference profile  $\mathbf{P} \in \mathcal{P}^N$ .

In practice, “running” a hybrid mechanism is straightforward: first, collect the preference reports. Second, toss a  $\beta$ -coin to determine whether to use  $\psi$  (probability  $\beta$ ) or  $\varphi$  (probability  $1 - \beta$ ). Third, apply this mechanism to the reported preference profile. Our next result formalizes the intuition that hybrids have at least intermediate signatures.

**Theorem 2.** *Given a problem  $(N, M, \delta)$ , mechanisms  $\varphi, \psi$ , and  $\beta \in [0, 1]$ , we have*

$$\varepsilon(h_\beta) \leq (1 - \beta)\varepsilon(\varphi) + \beta\varepsilon(\psi), \quad (9)$$

$$\delta(h_\beta) \leq (1 - \beta)\delta(\varphi) + \beta\delta(\psi). \quad (10)$$

*Proof Outline (formal proof in Appendix F.3).* We write out the definitions of  $\varepsilon(h_\beta)$  and  $\delta(h_\beta)$ , each of which may involve taking a maximum. The two inequalities are then obtained with the help of the triangle inequality.  $\square$

In words, the signatures of  $\beta$ -hybrids are always weakly better than the  $\beta$ -convex combination of the signatures of the two component mechanisms.

There are two important takeaways from Theorem 2. First, it yields a strong argument in favor of randomization: given two mechanisms with attractive manipulability and deficit, randomizing between the two always yields a mechanism with a signature that is at least as attractive as the  $\beta$ -convex combination of the signatures of the two mechanisms. As Example 3 in Appendix B shows, randomizing in this way can yield *strictly* preferable mechanisms that have *strictly* lower manipulability and *strictly* lower deficit than either of the component mechanisms.

The second takeaway is that the common fairness requirement of *anonymity* comes “for free” in terms of manipulability and deficit (provided that the deficit measure  $\delta$  is itself anonymous): given any mechanism  $\varphi$ , an anonymous mechanism can be constructed by randomly assigning the agents to new roles. This yields a hybrid mechanism with many components, each of which corresponds to the original mechanism with agents assigned to different roles. Under an anonymous deficit notion, every component will have the same signature as  $\varphi$ . It follows from Theorem 2 that this new anonymous mechanism has a weakly better signature than  $\varphi$ . Similarly, we can impose *neutrality* without having to accept higher manipulability or more deficit (if  $\delta$  is itself neutral). We formalize these insights in Appendix C.

## 8. The Pareto Frontier

Recall that optimal mechanisms are those mechanisms that trade off manipulability and deficit optimally. They form the Pareto frontier because we cannot achieve a strictly lower deficit without accepting strictly higher manipulability.

**Definition 12** (Pareto Frontier). Given a problem  $(N, M, \delta)$ , let  $\bar{\varepsilon}$  be the smallest manipulability bound that is compatible with  $d$ -maximization; formally,

$$\bar{\varepsilon} = \min\{\varepsilon \in [0, 1] \mid \exists \varphi : \varphi \text{ } d\text{-maximizing \& } \varepsilon\text{-approximately SP}\}. \quad (11)$$

The *Pareto frontier* is the union of all mechanisms that are optimal for some manipulability bound  $\varepsilon \in [0, \bar{\varepsilon}]$ ; formally,

$$\text{PF} = \bigcup_{\varepsilon \in [0, \bar{\varepsilon}]} \text{OPT}(\varepsilon). \quad (12)$$

The special manipulability bound  $\bar{\varepsilon}$  is chosen such that maximal  $d$ -value can be achieved with an  $\bar{\varepsilon}$ -approximately strategyproof mechanisms ( $\bar{\varphi}$ , say) but not with any mechanism that has strictly lower manipulability. Since  $\bar{\varphi}$  has deficit 0, any mechanism  $\varphi$  that is optimal at some larger manipulability bound  $\varepsilon > \bar{\varepsilon}$  may be more manipulable than  $\bar{\varphi}$  but  $\varphi$  cannot have a strictly lower deficit. Thus, we can restrict attention to manipulability bounds between 0 and  $\bar{\varepsilon}$  (instead of 0 and 1). From the perspective of the mechanism designer, mechanisms on the Pareto frontier are the only mechanisms

that should be considered; if a mechanism is *not* on the Pareto frontier, we can find another mechanism that is a Pareto-improvement in the sense that it has strictly lower manipulability, strictly lower deficit, or both.

## 8.1. Monotonicity and Convexity

Recall that we have defined  $\delta(\varepsilon)$  as the smallest deficit that can be attained by any  $\varepsilon$ -approximately strategyproof mechanism. Thus, the signatures of mechanisms on the Pareto frontier are described by the mapping  $\varepsilon \mapsto \delta(\varepsilon)$  that associates each manipulability bound  $\varepsilon \in [0, \bar{\varepsilon}]$  with the deficit  $\delta(\varepsilon)$  of optimal mechanisms at this manipulability bound. Based on our results so far, we can already make interesting observations about the Pareto frontier by analyzing this mapping.

**Corollary 1.** *Given a problem  $(N, M, \delta)$ , the mapping  $\varepsilon \mapsto \delta(\varepsilon)$  is monotonically decreasing and convex.*

Monotonicity follows from the definition of optimal mechanisms, and convexity is a consequence of Theorem 2 (see Appendix F.4 for a formal proof). While monotonicity is merely reassuring, convexity is non-trivial and very important. It means that when we relax strategyproofness, the first unit of manipulability that we give up allows the largest reduction of deficit. For any additional unit of manipulability that we sacrifice, the deficit will be reduced at a lower rate, which means *decreasing marginal returns*. Thus, we can expect to capture most gains in  $d$ -value from relaxing strategyproofness early on. Moreover, convexity and monotonicity together imply *continuity*. This means that trade-offs along the Pareto frontier are *smooth* in the sense that a tiny reduction of the manipulability bound  $\varepsilon$  does not require accepting a vastly higher deficit.

For mechanism designers, these observations deliver an important lesson: the most substantial gains (per unit of manipulability) arise from relaxing strategyproofness just “a little bit.” This provides encouragement to investigate the gains from accepting even small amounts of manipulability. On the other hand, if the initial gains are not worth the sacrifice, then gains from accepting more manipulability will not be “surprisingly” attractive either.

## 8.2. A Structural Characterization of the Pareto Frontier

In Section 6, we have shown that we can identify optimal mechanisms for individual manipulability bounds by solving the linear program FINDOPT. In Section 7, we have introduced hybrids, and we have shown how mixing two mechanisms results in new mechanisms with intermediate or even superior signatures. We now give our main result, a structural characterization of the Pareto frontier in terms of these two building blocks, namely *optimal mechanisms* and *hybrids*.

**Theorem 3.** *Given a problem  $(N, M, \delta)$ , there exists a finite set of supporting manipulability bounds*

$$\varepsilon_0 = 0 < \varepsilon_1 < \dots < \varepsilon_{K-1} < \varepsilon_K = \bar{\varepsilon}, \quad (13)$$

*such that for any  $k \in \{1, \dots, K\}$  and any  $\varepsilon \in [\varepsilon_{k-1}, \varepsilon_k]$  with  $\beta = \frac{\varepsilon - \varepsilon_{k-1}}{\varepsilon_k - \varepsilon_{k-1}}$  we have that*

$$\text{OPT}(\varepsilon) = (1 - \beta)\text{OPT}(\varepsilon_{k-1}) + \beta\text{OPT}(\varepsilon_k), \quad (14)$$

$$\delta(\varepsilon) = (1 - \beta)\delta(\varepsilon_{k-1}) + \beta\delta(\varepsilon_k). \quad (15)$$

*Proof Outline (formal proof in Appendix F.5).* Our proof exploits that  $\text{OPT}(\varepsilon)$  corresponds to the solutions of the linear program FINDOPT (Section 6.3) with feasible set  $F_\varepsilon = \{x \mid Dx \leq d, Ax \leq \varepsilon\}$ , where neither  $D$ , nor  $d$ , nor  $A$  depend on  $\varepsilon$ . First, we show that if a set of constraints is binding for  $F_\varepsilon$ , then it is binding for  $F_{\varepsilon'}$  for  $\varepsilon'$  within a compact interval  $[\varepsilon^-, \varepsilon^+]$  that contains  $\varepsilon$  and not binding for any  $\varepsilon'' \notin [\varepsilon^-, \varepsilon^+]$ . With finiteness of the number of constraints of the LP, this yields a finite segmentation of  $[0, \bar{\varepsilon}]$ . The vertex-representations (Grünbaum, 2003) can then be used to show that on each segment  $[\varepsilon_{k-1}, \varepsilon_k]$ , the *solution sets*  $S_\varepsilon = \text{argmin}_{F_\varepsilon} d$  are exactly the  $\beta$ -convex combinations of  $S_{\varepsilon_{k-1}}$  and  $S_{\varepsilon_k}$  with  $\beta = \frac{\varepsilon - \varepsilon_{k-1}}{\varepsilon_k - \varepsilon_{k-1}}$ .  $\square$

It would be particularly simple if the optimal mechanisms at some manipulability bound  $\varepsilon$  were just the  $\beta$ -hybrids of optimal mechanisms at 0 and  $\bar{\varepsilon}$  with  $\beta = \varepsilon/\bar{\varepsilon}$ . While this is not true in general, Theorem 3 shows that the Pareto frontier has a linear structure over each interval  $[\varepsilon_{k-1}, \varepsilon_k]$ . Thus, it is completely identified by two building blocks: (1) the *sets of optimal mechanisms*  $\text{OPT}(\varepsilon_k)$  for finitely many  $\varepsilon_k, k = 0, \dots, K$ , and (2) *hybrid mechanisms*, which provide the missing link for  $\varepsilon \neq \varepsilon_k$ . Representatives of  $\text{OPT}(\varepsilon_k)$  can be obtained by solving the linear program FINDOPT at the supporting manipulability

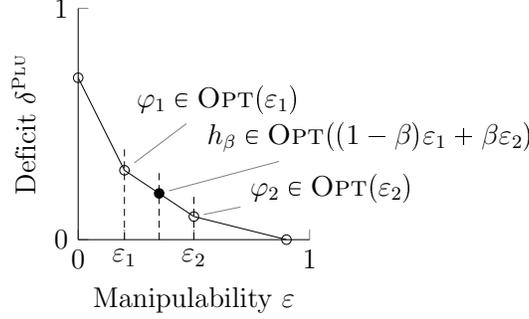


Figure 4: Illustrative example of the signatures of the Pareto frontier.

bounds  $\varepsilon_k$ ; and for intermediate bounds  $\varepsilon = (1 - \beta)\varepsilon_{k-1} + \beta\varepsilon_k$ , a mechanism  $\varphi$  is optimal at  $\varepsilon$  *if and only if* it is a  $\beta$ -hybrid of two mechanisms  $\varphi_{k-1} \in \text{OPT}(\varepsilon_{k-1})$  and  $\varphi_k \in \text{OPT}(\varepsilon_k)$ .

Theorem 3 allows an additional insight about the mapping  $\varepsilon \mapsto \delta(\varepsilon)$  (where  $\delta(\varepsilon) = \min\{\delta(\varphi) \mid \varphi \text{ } \varepsilon\text{-approximately SP}\}$ ). We have already observed that this mapping is monotonic, decreasing, convex, and continuous (see Corollary 1). In addition, we obtain *piecewise linearity*.

**Corollary 2.** *Given a problem  $(N, M, \delta)$ ,  $\varepsilon \mapsto \delta(\varepsilon)$  is piecewise linear.*

Figure 4 illustrates the results of Theorem 3 and Corollary 2: first, the mapping  $\varepsilon \mapsto \delta(\varepsilon)$  is monotonic, decreasing, convex, continuous, and piecewise linear, and, second, optimal mechanisms with intermediate manipulability can be obtained by mixing two mechanisms that are optimal at the two adjacent supporting manipulability bounds, respectively.

## 9. Computing the Pareto Frontier: Algorithms and Examples

In this section we discuss an algorithm that exploits the structural characterization from Theorem 3 to compute all supporting manipulability bounds. Then we derive the Pareto frontier for two example problems.

## 9.1. Algorithm FindBounds

Recall that the linear program `FINDOPT` can be used to determine a representative of the set of optimal mechanisms  $\text{OPT}(\varepsilon)$ . One naïve approach to identifying the Pareto frontier would be to run `FINDOPT` for many different manipulability bounds to obtain optimal mechanisms at each of these bounds, and then consider these mechanisms and their hybrids. However, this method has two drawbacks: first, and most importantly, it would not yield the correct Pareto frontier. The result can, at best, be viewed as a *conservative estimate*, since choosing fixed manipulability bounds is not guaranteed to identify any actual supporting manipulability bounds exactly. Second, from a computational perspective, executing `FINDOPT` is expensive, which is why we would like to keep the number of executions as low as possible.

Knowing the structure of the Pareto frontier, we can use the information obtained from iterated applications of `FINDOPT` to interpolate the most promising candidates for supporting manipulability bounds. In Appendix D we provide the algorithms `FINDBOUNDS` (Algorithm 1) and `FINDLOWER` (Algorithm 2) that do this. Proposition 3 summarizes their properties.

**Proposition 3.** *Given a problem  $(N, M, \delta)$ , the algorithms `FINDLOWER` and `FINDBOUNDS` require at most  $4K + \log_2(1/\varepsilon_1) - 1$  executions of the linear program `FINDOPT` to determine all supporting manipulability bounds of the Pareto frontier, where  $K$  is the number of supporting manipulability bounds and  $\varepsilon_1$  is the smallest non-trivial supporting manipulability bound.*

Due to space constraints we delegate the detailed discussion of the algorithms and the proof of Proposition 3 to Appendix D.

## 9.2. Examples: Plurality Scoring and Veto Scoring

In this section, we consider two concrete problems and derive the respective Pareto frontiers. The two examples highlight the different shapes that the Pareto frontier can take. Both settings contain 3 alternatives and 3 agents with strict preferences over the alternatives, but they differ in the desideratum.

**Plurality and Random Dictatorships** Consider a setting with agents  $N = \{1, 2, 3\}$ , alternatives  $M = \{a, b, c\}$ , and where agents' preferences are strict. Suppose further that the desideratum is to select alternatives that are the first choice of as many agents as possible. From Example 1 we know that the corresponding desideratum function is  $d^{\text{PLU}}(j, \mathbf{P}) = n_j^1/n$ , where  $n_j^1$  is the number of agents whose first choice under  $\mathbf{P}$  is  $j$ . Finally, suppose that the deficit  $\delta^{\text{PLU}}$  is the worst-case  $d^{\text{PLU}}$ -deficit; formally,

$$\delta^{\text{PLU}}(\varphi) = \max_{\mathbf{P} \in \mathcal{P}^N} \left( \max_{j \in M} (d^{\text{PLU}}(j, \mathbf{P})) - d^{\text{PLU}}(\varphi(\mathbf{P}), \mathbf{P}) \right). \quad (16)$$

Using `FINDBOUNDS`, we find that in this problem, the Pareto frontier has only two supporting manipulability bounds  $\varepsilon_0 = 0$  and  $\bar{\varepsilon} = \varepsilon_1 = 1/3$ . Thus, it consists precisely of the hybrids of mechanisms that are optimal at  $\varepsilon_0$  and  $\varepsilon_1$  by Theorem 3. Moreover, we can use `FINDOPT` to obtain representatives of  $\text{OPT}(0)$  and  $\text{OPT}(1/3)$ . Interestingly, the representatives determined by the algorithm have a familiar structure: first, the representative of  $\text{OPT}(0)$  corresponds to Random Dictatorship with a uniform choice of the dictator. Thus, in this problem, no strategyproof mechanism has lower deficit than Random Dictatorship. Second, the representative of  $\text{OPT}(1/3)$  corresponds to *Uniform Plurality*, a mechanism that determines the set of  $d^{\text{PLU}}$ -maximizing alternatives and selects one of these uniformly at random. Thus, in this problem, no  $d$ -maximizing mechanism has lower manipulability than Uniform Plurality. Figure 5 (a) depicts the signatures of the Pareto frontier.

In addition, we prove these insights analytically.

**Proposition 4.** *In a problem  $(N, M, \delta^{\text{PLU}})$  with  $n = 3$  agents,  $m = 3$  alternatives, strict preferences, and where  $\delta^{\text{PLU}}$  is the worst-case  $d^{\text{PLU}}$ -deficit, the following hold:*

1. *The Pareto frontier has two supporting manipulability bounds 0 and 1/3.*
2. *Random Dictatorship is a representative of  $\text{OPT}(0)$ .*
3. *Uniform Plurality is a representative of  $\text{OPT}(1/3)$ .*

The formal proof is given in Appendix F.6.

**Veto and Random Duple** The second problem again involves a setting with three agents, three alternatives, and agents' preferences are strict. The difference to the previous

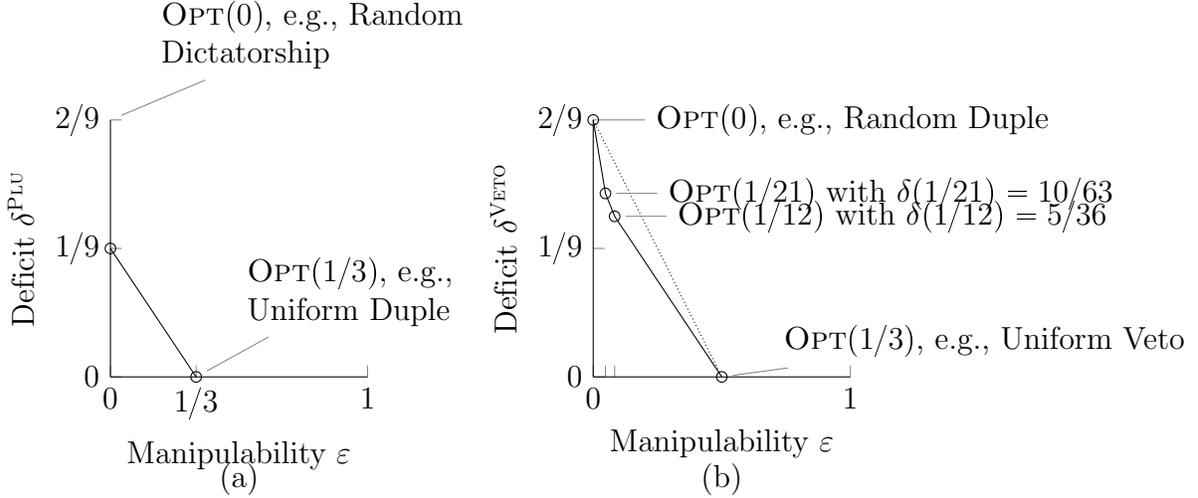


Figure 5: Plot of signatures of the Pareto frontier for  $n = 3$  agents,  $m = 3$  alternatives, strict preferences, (a) worst-case  $d^{\text{PLU}}$ -deficit  $\delta^{\text{PLU}}$  and (b) worst-case  $d^{\text{VETO}}$ -deficit  $\delta^{\text{VETO}}$ .

problem is the different desideratum: this time, our goal is to select an alternative that is the *last* choice of as *few* agents as possible. This desideratum is reflected by the desideratum function  $d^{\text{VETO}}(j, \mathbf{P}) = (n - n_j^m)/n$ , where  $n_j^m$  is the number of agents whose last choice under  $\mathbf{P}$  was  $j$ . The worst-case  $d^{\text{VETO}}$ -deficit  $\delta^{\text{VETO}}$  is given by

$$\delta^{\text{VETO}}(\varphi) = \max_{\mathbf{P} \in \mathcal{P}^N} \left( \max_{j \in M} (d^{\text{VETO}}(j, \mathbf{P})) - d^{\text{VETO}}(\varphi(\mathbf{P}), \mathbf{P}) \right). \quad (17)$$

Again, we use the algorithm `FINDBOUNDS` to determine the supporting manipulability bounds of the Pareto frontier in this problem. These are  $\varepsilon_0 = 0$ ,  $\varepsilon_1 = 1/21$ ,  $\varepsilon_2 = 1/12$ ,  $\bar{\varepsilon} = \varepsilon_3 = 1/2$ . As in the previous example, we can compute representatives of the optimal mechanisms at each of these bounds. For the extreme bounds  $\varepsilon_0 = 0$  and  $\bar{\varepsilon} = 1/2$ , these representatives again have a familiar structure: first, the representative of  $\text{OPT}(0)$  corresponds to *Random Duple*, a mechanism that picks two alternatives uniformly at random and then selects the one that is preferred to the other by a majority of the agents (breaking ties randomly). Second, the representative of  $\text{OPT}(1/2)$  corresponds to *Uniform Veto*, a mechanism that finds all alternatives that are the last choice of a minimal number of agents and selects one of these alternatives uniformly at random.

Figure 5 (b) depicts the signatures of the Pareto frontier in this problem.

We also created representatives of  $\text{OPT}(1/21)$  and  $\text{OPT}(1/12)$  via `FINDOPT` (they are given in Appendix F.7 in their numerical form). While the interpretation of these two mechanisms is not straightforward, it is clear that neither of them is a hybrid Random Duple and Uniform Veto. Indeed, in order to generate mechanisms on the Pareto frontier in this problem, we cannot simply consider hybrids of optimal mechanisms from the extreme supporting manipulability bounds. Instead, we can (and have to) exploit the additional design freedom in the particular problem by separately designing optimal mechanisms for the two intermediate supporting manipulability bounds  $\varepsilon_1 = 1/21$  and  $\varepsilon_2 = 1/12$  specifically.

As in the previous example, we convince ourselves of the correctness of these assertions by proving them analytically.

**Proposition 5.** *In a problem  $(N, M, \delta^{\text{VETO}})$  with  $n = 3$  agents,  $m = 3$  alternatives, strict preferences, and where  $\delta^{\text{VETO}}$  is the worst-case  $d^{\text{Veto}}$ -deficit, the following hold:*

1. *The Pareto frontier has four supporting manipulability bounds  $0, 1/21, 1/12, 1/2$ .*
2. *Random Duple is a representative of  $\text{OPT}(0)$ .*
3. *Uniform Veto is a representative of  $\text{OPT}(1/2)$ .*
4. *Hybrids of Random Duple and Uniform Veto not optimal for  $\beta \in (0, 1/2)$ .*

The formal proof is given in Appendix F.7.

*Remark 6.* The fact that we obtain Random Dictatorship and Random Duple as the optimal strategyproof mechanisms in the two problems, respectively, highlights a connection between our results and the characterization of strategyproof random mechanisms by Gibbard (1977). We discuss this connection in Appendix C and we present a new symmetric decomposition result that may be of independent interest.

**Pareto Frontiers in Larger Settings** For larger settings analytic insights about the Pareto frontier are presently not available. However, we can apply the algorithm `FINDBOUNDS` to identify the Pareto frontier algorithmically.<sup>6</sup> Figure 6 shows signatures of Pareto frontiers for  $m = 3$  alternatives, up to  $n = 18$  agents, and the worst-case  $d^{\text{PLU}}$ -deficit (a) and  $d^{\text{VETO}}$ -deficit (b), respectively.

---

<sup>6</sup>This can be achieved by solving reduced linear programs, see (Mennle, Abaecherli and Seuken, 2015).

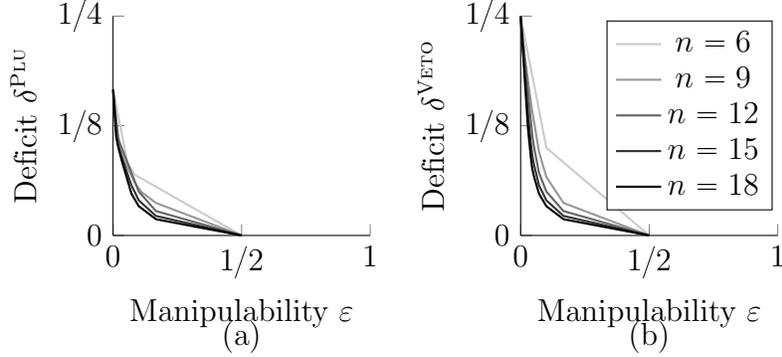


Figure 6: Plots of signatures of the Pareto frontier for  $n \in \{6, 9, 12, 15, 18\}$  agents,  $m = 3$  alternatives, strict preferences, and worst-case  $d^{\text{PLU}}$ -deficit (left) and worst-case  $d^{\text{VETO}}$ -deficit (right).

The plots suggest that the Pareto frontier converges to an L-shape as the number of agents grows. Based on this observation, we conjecture that for these two desiderata near-ideal mechanisms can be found in large markets. If true, we could find mechanisms that are *almost* strategyproof and *almost* achieve the desideratum perfectly. Proving these limit results will be an interesting topic for future research.

## 10. Conclusion

In this paper, we have presented a structural characterization of the Pareto frontier for random ordinal mechanisms. Loosely speaking, the *Pareto frontier* consists of those mechanisms that make optimal trade-offs between incentive properties (measured by manipulability  $\varepsilon(\varphi)$ ) and other desiderata (measured by deficit  $\delta(\varphi)$ ), such as Condorcet consistency or the goal of choosing an alternative that is the first choice of as many agents as possible.

We have achieved our main result in three distinct steps: first, we have shown that  $\varepsilon$ -approximate strategyproofness can be equivalently expressed by a finite set of linear constraints. This has enabled us to define the linear program FINDOPT to identify the set of all mechanisms that have minimal deficit but satisfy  $\varepsilon$ -approximate strategyproofness for a given manipulability bound  $\varepsilon$ . Second, we have shown how hybrid mechanisms, which are convex combinations of two component mechanisms, trade off manipulability

and deficit. In particular, we have given a guarantee that the signature of a  $\beta$ -hybrid  $h_\beta$  is always at least as good as the  $\beta$ -convex combination of the signatures of the two component mechanisms. Third, we have shown that the Pareto frontier consists of two building blocks: (1) there exists a finite set of *supporting manipulability bounds*  $\varepsilon_0, \dots, \varepsilon_K$  such that we can characterize the set of optimal mechanisms at each of the bounds  $\varepsilon_k$  as the set of solutions to the linear program FINDOPT at  $\varepsilon_k$ , and (2) for any intermediate manipulability bound  $\varepsilon = (1 - \beta)\varepsilon_{k-1} + \beta\varepsilon_k$ , the set of optimal mechanisms at  $\varepsilon$  is precisely the set of  $\beta$ -hybrids of optimal mechanisms at each of the two adjacent supporting manipulability bounds  $\varepsilon_{k-1}$  and  $\varepsilon_k$ .

Our results have a number of interesting consequences (beyond their relevance in this paper): first, Theorem 1 gives a finite set of linear constraints that is equivalent to  $\varepsilon$ -approximate strategyproofness. This makes  $\varepsilon$ -approximate strategyproofness accessible to algorithmic analysis. In particular, it enables the use of this incentive requirement under the automated mechanism design paradigm.

Second, the performance guarantees for hybrid mechanisms from Theorem 2 yield convincing arguments in favor of randomization. In particular, we learn that the important requirements of anonymity and neutrality come “for free;” mechanism designers do not have to accept a less desirable signature when imposing either or both (provided that the deficit measure is anonymous, neutral, or both).

Third, our main result, Theorem 3, has provided a *structural understanding* of the whole Pareto frontier. Knowledge of the Pareto frontier enables mechanism designers to make a completely informed decision about trade-offs between manipulability and deficit. In particular, we now have a way to determine precisely by how much the performance of mechanisms (with respect to a given desideratum) can be improved when allowing additional manipulability. An important learning is that the mapping  $\varepsilon \mapsto \delta(\varepsilon)$ , which associates each manipulability bound with the lowest achievable deficit at this manipulability bound, is *monotonic, decreasing, convex, continuous, and piecewise linear*. This means that when trading off manipulability and deficit along the Pareto frontier, the trade-offs are smooth, and the earliest sacrifices yield the greatest gains. Thus, it can be worthwhile to consider even small bounds  $\varepsilon > 0$  in order to obtain substantial improvements.

Finally, we have illustrated our results by considering two concrete problems. In both problems, three agents had strict preferences over three alternatives. In the first

problem, the desideratum was to choose an alternative that is the *first choice of as many agents as possible* (i.e., Plurality scoring), and in the second problem, the desideratum was to choose an alternative that is the *last choice of as few agents as possible* (i.e., Veto scoring). In both problems, we have computed the Pareto frontier and verified the resulting structure analytically. The examples have shown that the Pareto frontier may be completely linear (first problem) or truly non-linear (second problem). For the same desiderata and up to  $n = 18$  agents we have determined the Pareto frontier algorithmically and formulated a conjecture about its limit behavior.

In summary, we have given novel insights about the Pareto frontier for random ordinal mechanisms. We have proven our results for the full ordinal domain that includes indifferences, but they continue to hold for many other interesting domains that arise by restricting the space of preference profiles, such as the assignment domain and the two-sided matching domain. When impossibility results restrict the design of strategyproof mechanisms, we have provided a new perspective on the unavoidable trade-off between incentives and other desiderata along this Pareto frontier.

## References

- Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 2003. “School Choice: A Mechanism Design Approach.” *American Economic Review*, 93(3): 729–747.
- Azevedo, Eduardo, and Eric Budish.** 2015. “Strategy-proofness in the Large.” Working Paper.
- Aziz, Haris, Felix Brandt, and Markus Brill.** 2013. “On the Tradeoff Between Economic Efficiency and Strategy Proofness in Randomized Social Choice.” In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*.
- Aziz, Haris, Florian Brandl, and Felix Brandt.** 2014. “On the Incompatibility of Efficiency and Strategyproofness in Randomized Social Choice.” In *Proceedings of the 28th Conference on Artificial Intelligence (AAAI)*.
- Birrell, Eleanor, and Rafael Pass.** 2011. “Approximately Strategy-Proof Voting.” In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*.

- Budish, Eric.** 2011. “The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes.” *Journal of Political Economy*, 119(6): 1061–1103.
- Carroll, Garbiel.** 2013. “A Quantitative Approach to Incentives: Application to Voting Rules.” Working Paper.
- Chatterji, Shurojit, Remzi Sanver, and Arunava Sen.** 2013. “On Domains that Admit Well-behaved Strategy-proof Social Choice Functions.” *Journal of Economic Theory*, 148(3): 1050–1073.
- Ehlers, Lars, Hans Peters, and Ton Storcken.** 2002. “Strategy-Proof Probabilistic Decision Schemes for One-Dimensional Single-Peaked Preferences.” *Journal of Economic Theory*, 105(2): 408–434.
- Featherstone, Clayton.** 2011. “A Rank-based Refinement of Ordinal Efficiency and a new (but Familiar) Class of Ordinal Assignment Mechanisms.” Working Paper.
- Gibbard, Allan.** 1973. “Manipulation of Voting Schemes: a General Result.” *Econometrica*, 41(4): 587–601.
- Gibbard, Allan.** 1977. “Manipulation of Schemes That Mix Voting with Chance.” *Econometrica*, 45(3): 665–81.
- Grünbaum, Branko.** 2003. *Convex Polytopes. Graduate Texts in Mathematics.* 2 ed., Springer.
- Mennle, Timo, and Sven Seuken.** 2016. “The Pareto Frontier for Random Mechanisms [Extended Abstract].” In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*.
- Mennle, Timo, and Sven Seuken.** 2017a. “Hybrid Mechanisms: Trading Off Strategyproofness and Efficiency of Random Assignment Mechanisms.” Working Paper.
- Mennle, Timo, and Sven Seuken.** 2017b. “Partial Strategyproofness: Relaxing Strategyproofness for the Random Assignment Problem.” Working Paper.
- Mennle, Timo, Daniel Abaecherli, and Sven Seuken.** 2015. “Computing Pareto Frontier for Randomized Mechanisms.” Mimeo.
- Moulin, Hervé.** 1980. “On Strategy-Proofness and Single Peakedness.” *Public Choice*, 35(4): 437–455.

- Pacuit, Eric.** 2012. “Voting Methods.” In *The Stanford Encyclopedia of Philosophy*. Winter 2012 ed., ed. Edward Zalta. Stanford University.
- Procaccia, Ariel.** 2010. “Can Approximation Circumvent Gibbard-Satterthwaite?” In *Proceedings of the 24th Conference on Artificial Intelligence (AAAI)*.
- Procaccia, Ariel, and Jeffrey Rosenschein.** 2006. “The Distortion of Cardinal Preferences in Voting.” In *Cooperative Information Agents X*. Vol. 4149 of *Lecture Notes in Computer Science*, ed. Matthias Klusch, Michael Rovatsos and Terry Payne, 317–331. Springer.
- Roth, Alvin.** 1982. “The Economics of Matching: Stability and Incentives.” *Mathematical Operations Research*, 7: 617–628.
- Roth, Alvin.** 1984. “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory.” *Journal of Political Economy*, 92(6): 991–1016.
- Sandholm, Tuomas.** 2003. “Automated Mechanism Design: A New Application Area for Search Algorithms.” In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*.
- Satterthwaite, Mark.** 1975. “Strategy-proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions.” *Journal of Economic Theory*, 10(2): 187–217.

# APPENDIX

## A. Relative and Ex-ante Deficit

We have defined the deficit of outcomes (Definition 7) as the absolute difference between the achievable and the achieved  $d$ -value; and we have defined the deficit of mechanisms (Definition 8) as the worst-case deficit across all possible preference profiles. In this section we present two variations of these definitions, the *relative deficit* of outcomes and the *ex-ante deficit* of mechanisms. Our results in this paper hold for any of these variations.

### A.1. Relative Deficit of Outcomes

In some situations, it may be more natural to consider a *relative* difference, e.g., the ratio between the achieved and the maximal achievable  $d$ -value. As we show next, it is without loss of generality that in this paper we have restricted our attention to absolute differences. This is because the relative  $d$ -deficit can always be expressed as an absolute  $\tilde{d}$ -deficit, where the adjusted desideratum function  $\tilde{d}$  is obtained from  $d$  by scaling. Proposition 6 makes this argument precise.

To state this equivalence formally, we need to define the relative deficit: for any preference profile  $\mathbf{P} \in \mathcal{P}^N$ , the  *$d$ -value margin at  $\mathbf{P}$*  is the difference between the highest and the lowest  $d$ -value achievable by any alternative at  $\mathbf{P}$ . We set  $d^{\max}(\mathbf{P}) = \max_{j \in M} d(j, \mathbf{P})$ ,  $d^{\min}(\mathbf{P}) = \min_{j \in M} d(j, \mathbf{P})$ , and  $d^{\text{margin}}(\mathbf{P}) = d^{\max}(\mathbf{P}) - d^{\min}(\mathbf{P})$ . Note that for the special case where  $d^{\text{margin}}(\mathbf{P}) = 0$ , all alternatives (and therefore all outcomes) have the same  $d$ -value. In this case, *any* alternative is  $d$ -maximizing at  $\mathbf{P}$ . For an outcome  $x \in \Delta(M)$ , the *relative  $d$ -deficit of  $x$  at  $\mathbf{P}$*  is the  $d$ -deficit of  $x$  at  $\mathbf{P}$ , normalized by the  $d$ -value margin at  $\mathbf{P}$ ; formally,

$$\delta_d^{\text{relative}}(x, \mathbf{P}) = \begin{cases} \frac{d^{\max}(\mathbf{P}) - d(x, \mathbf{P})}{d^{\text{margin}}(\mathbf{P})}, & \text{if } d^{\text{margin}}(\mathbf{P}) > 0, \\ 0, & \text{else.} \end{cases} \quad (18)$$

**Proposition 6.** *For any desideratum function  $d$ , there exists a desideratum function  $\tilde{d}$  such that the relative  $d$ -deficit coincides with the absolute  $\tilde{d}$ -deficit, such that for all*

outcomes  $x \in \Delta(M)$  and all preference profiles  $\mathbf{P} \in \mathcal{P}^N$ , we have

$$\delta_d^{relative}(x, \mathbf{P}) = \delta_{\tilde{d}}(x, \mathbf{P}). \quad (19)$$

The proof follows immediately by setting  $\tilde{d}(j, \mathbf{P}) = \frac{d^{\max}(\mathbf{P}) - d(j, \mathbf{P})}{d^{\text{margin}}(\mathbf{P})}$ , whenever  $d^{\text{margin}}(\mathbf{P}) > 0$ , and  $\tilde{d}(j, \mathbf{P}) = 0$  otherwise. Proposition 6 shows that including relative deficit does not enrich the space of possible criteria but that the space of desideratum functions is rich enough to cover relative deficits implicitly. Thus, it is without loss of generality that we have restricted attention to absolute deficits in this paper.

## A.2. Ex-ante Deficit of Mechanisms

Recall that we defined the deficit of mechanisms as the worst-case loss that society incurs from using a mechanism that is not  $d$ -maximizing. This is most meaningful when we have no prior knowledge of the agents' preferences and therefore need to design mechanisms that achieve the desideratum as well as possible across all possible preference profiles.

However, in some situations, we may have probabilistic information about the agents' preferences, which we would like to exploit to design better mechanisms. Suppose that the agents' preference profiles are drawn from a known distribution  $\mathbb{P}$ . In this case, we may prefer a mechanism that induces high *expected*  $d$ -value under  $\mathbb{P}$ .

**Definition 13** (Ex-ante Deficit). Given a setting  $(N, M)$ , a desideratum function  $d$ , a probability distribution  $\mathbb{P}$  over preference profiles, and a mechanism  $\varphi$ , the *ex-ante  $d$ -deficit of  $\varphi$  with respect to  $\mathbb{P}$*  (in  $(N, M)$ ) is

$$\delta_d^{\mathbb{P}}(\varphi) = \sum_{\mathbf{P} \in \mathcal{P}^N} \mathbb{P}[\mathbf{P}] \cdot \delta_d(\varphi(\mathbf{P}), \mathbf{P}). \quad (20)$$

Minimizing  $\delta_d^{\mathbb{P}}(\varphi)$  corresponds to minimizing the expected  $d$ -deficit from applying  $\varphi$  *ex-ante* (i.e., before the agents' preferences are instantiated from  $\mathbb{P}$ ). This approach is attractive in situations where the same mechanism is applied repeatedly for different groups of agents, so that the outcomes are attractive *on average* across all the repetitions.

We can incorporate the ex-ante deficit into the linear program FINDOPT (Linear Program 1) by exchanging the constraint labeled (*Deficit*). Instead of making the variable  $\mathbf{d}$  an upper bound for the deficit of the mechanism across all preference profiles, we

need to make  $d$  an upper bound for the expected deficit of the mechanism, where this expectation is taken with respect to the distribution  $\mathbb{P}$  over preference profiles. We achieve this by including the following constraint.

$$d \geq \sum_{\mathbf{P} \in \mathcal{P}^N} \mathbb{P}[\mathbf{P}] \cdot \left( \max_{j \in M} d(j, \mathbf{P}) - \sum_{j \in M} f_j(\mathbf{P}) \cdot d(j, \mathbf{P}) \right). \quad (\text{Ex-ante deficit})$$

## B. Strict Improvements from Hybrids in Theorem 2

Theorem 2 showed two weak inequalities for hybrid mechanisms. We now give an example that shows that a hybrid of two mechanisms can in fact have a strictly lower manipulability and a strictly lower deficit than both of its component mechanisms.

**Example 3.** Consider a problem with one agent and three alternatives  $a, b, c$ , where  $\delta$  is the worst-case deficit that arises from Plurality scoring. Let  $\varphi$  and  $\psi$  be two mechanisms whose outcomes depend only on the agent's relative ranking of  $b$  and  $c$ .

	$\varphi$			$\psi$		
<b>Report</b>	<b><math>a</math></b>	<b><math>b</math></b>	<b><math>c</math></b>	<b><math>a</math></b>	<b><math>b</math></b>	<b><math>c</math></b>
If $P : b \geq c$	0	2/3	1/3	5/9	1/9	1/3
If $P : c > b$	1/3	1/3	1/3	1/9	5/9	1/3

It is a simple exercise to verify that the mechanisms' signatures are  $(\varepsilon(\varphi), \delta(\varphi)) = (1/3, 1)$  and  $(\varepsilon(\psi), \delta(\psi)) = (4/9, 8/9)$ . Furthermore, for  $\beta = 3/7$ , the hybrid  $h_\beta$  is constant and therefore strategyproof, and it has a deficit of  $\delta(h_\beta) = 16/21$ . Figure 7 illustrates these signatures. The guarantees from Theorem 2 are represented by the shaded area, where the signature of  $h_\beta$  can lie anywhere below and to the left of the  $\beta$ -convex combination of the signatures of  $\varphi$  and  $\psi$ .

## C. Anonymity, Neutrality, Symmetric Decomposition

Two common requirements in the design of ordinal mechanisms are anonymity and neutrality. Anonymity captures the intuition that a mechanism should not discriminate between agents; instead, the influence of any agent on the outcome should be the same

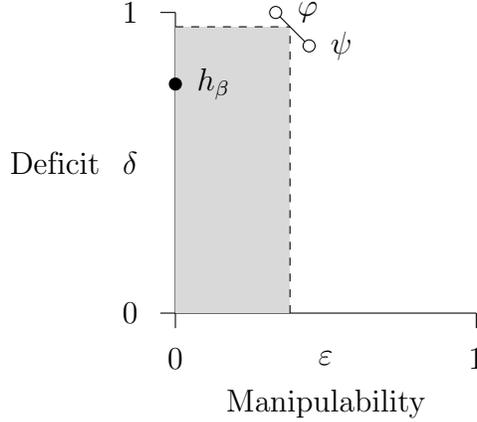


Figure 7: Example signature of hybrid  $h_\beta$  ( $\bullet$ ) must lie in the shaded area, weakly to the left and below the  $\beta$ -convex combinations of the signatures of  $\varphi$  and  $\psi$  ( $\circ$ ).

and independent of the agent's name. Neutrality requires that the mechanism is not biased towards particular alternatives; the decision should depend only on the agents' preferences but not on the names of the alternatives.

In this section we describe the implications of these two requirements for the design of optimal mechanisms: after providing definitions, we show how they can be incorporated in the linear program FINDOPT. Then we formally prove the second takeaway from Theorem 2 that (under certain conditions) both requirements come for free in terms of the signatures of optimal mechanisms. Finally, we prove a new symmetric decomposition result for strategyproof, anonymous, neutral mechanisms, that extends the characterization of strategyproof random mechanisms in (Gibbard, 1977).

### C.1. Definition of Anonymity and Neutrality

First, we define anonymity: for any *renaming of the agents* (i.e., any bijection  $\pi : N \rightarrow N$ ) and any preference profile  $\mathbf{P} = (P_1, \dots, P_n) \in \mathcal{P}^N$ , let  $\mathbf{P}^\pi = (P_{\pi(1)}, \dots, P_{\pi(n)})$  be the preference profile where the agents have exchanged their roles according to  $\pi$ . Agent  $i$  is now reporting the preference order  $P_{\pi(i)}$  that was reported by agent  $\pi(i)$  under the original preference profile  $\mathbf{P}$ . For any mechanism  $\varphi$  let  $\varphi^\pi$  be the mechanism under which the agents trade roles according to  $\pi$ ; formally, let  $\varphi^\pi(\mathbf{P}) = \varphi(\mathbf{P}^\pi)$  for any preference profile  $\mathbf{P} \in \mathcal{P}^N$ .

**Definition 14** (Anonymity). • A desideratum function  $d$  is *anonymous* if for all renamings  $\pi : N \rightarrow N$ , preference profiles  $\mathbf{P} \in \mathcal{P}^N$ , and alternatives  $j \in M$ , we have  $d(j, \mathbf{P}) = d(j, \mathbf{P}^\pi)$  (i.e., the  $d$ -value is independent of the order in which the agents submit their preferences).

- A probability distribution  $\mathbb{P}$  over preference profiles is *anonymous* if for all renamings  $\pi : N \rightarrow N$  and preference profiles  $\mathbf{P} \in \mathcal{P}^N$ , we have  $\mathbb{P}[\mathbf{P}] = \mathbb{P}[\mathbf{P}^\pi]$  (i.e., the probability of a preference profile does not depend on the order in which the different preference orders appear).
- The worst-case deficit  $\delta_d$  is *anonymous* if the underlying desideratum function  $d$  is anonymous. The ex-ante deficit  $\delta_d^{\mathbb{P}}$  is *anonymous* if the underlying desideratum function and the probability distribution  $\mathbb{P}$  are anonymous.
- A mechanism  $\varphi$  is *anonymous* if for all renamings  $\pi : N \rightarrow N$  and preference profiles  $\mathbf{P} \in \mathcal{P}^N$ , we have  $\varphi(\mathbf{P}) = \varphi^\pi(\mathbf{P})$  (i.e., the outcome of the mechanism is independent of the order in which the agents submit their preferences).

Next, we define neutrality: for any *renaming of the alternatives* (i.e., any bijection  $\varpi : M \rightarrow M$ ) and any preference order  $P_i \in \mathcal{P}$ , let  $P_i^\varpi$  be the preference order under which  $P_i^\varpi : \varpi(j) \geq \varpi(j')$  whenever  $P_i : j \geq j'$  for any alternatives  $j, j' \in M$ . This means that  $P_i^\varpi$  corresponds to  $P_i$ , except that the all alternatives have been renamed according to  $\varpi$ . For any preference profile  $\mathbf{P} = (P_1, \dots, P_n) \in \mathcal{P}^N$ , let  $\mathbf{P}^\varpi = (P_1^\varpi, \dots, P_n^\varpi)$  be the preference profile where the alternatives inside the agents' preference orders have been renamed according to  $\varpi$ . For any mechanism  $\varphi$  let  $\varphi^\varpi$  be the mechanism under which the alternatives are renamed according to  $\varpi$ ; formally, let  $\varphi_j^\varpi(\mathbf{P}) = \varphi_{\varpi(j)}(\mathbf{P}^\varpi)$  for all preference profile  $\mathbf{P} \in \mathcal{P}^N$  and alternative  $j \in M$ .

**Definition 15** (Neutrality). We define the following:

- A desideratum function  $d$  is *neutral* if for all renamings  $\varpi : M \rightarrow M$ , preference profiles  $\mathbf{P} \in \mathcal{P}^N$ , and alternatives  $j \in M$ , we have  $d(j, \mathbf{P}) = d(\varpi(j), \mathbf{P}^\varpi)$  (i.e., the  $d$ -value of any alternative is independent of its name).
- A probability distribution  $\mathbb{P}$  over preference profiles is *neutral* if for all renamings  $\varpi : M \rightarrow M$  and preference profiles  $\mathbf{P} \in \mathcal{P}^N$ , we have  $\mathbb{P}[\mathbf{P}] = \mathbb{P}[\mathbf{P}^\varpi]$  (i.e., the probability of a preference profile does not depend on the names of the alternatives).

- The worst-case deficit  $\delta_d$  is *neutral* if the underlying desideratum function  $d$  is neutral. The ex-ante deficit  $\delta_d^{\mathbb{P}}$  is *neutral* if the underlying desideratum function and the probability distribution  $\mathbb{P}$  are neutral.
- A mechanism  $\varphi$  is *neutral* if for all renamings  $\varpi : M \rightarrow M$ , preference profiles  $\mathbf{P} \in \mathcal{P}^N$ , and alternatives  $j \in M$ , we have  $\varphi_j(\mathbf{P}) = \varphi_j^{\varpi}(\mathbf{P}) = \varphi_{\varpi(j)}(\mathbf{P}^{\varpi})$  (i.e., the outcomes of the mechanism are independent of the names of the alternatives).

Incorporating anonymity and neutrality as additional constraints in the linear program FINDOPT is straightforward and can be done as follows.

LINEAR PROGRAM 1 (FINDOPT), EXTENDED.

$$\begin{aligned} & \dots \\ & \mathbf{f}_j(\mathbf{P}) = \mathbf{f}_j(\mathbf{P}^{\pi}), \quad \forall \mathbf{P} \in \mathcal{P}^N, j \in M, \pi : N \rightarrow N \text{ bijection} \quad (\textit{Anonymity}) \\ & \mathbf{f}_j(\mathbf{P}) = \mathbf{f}_{\varpi(j)}(\mathbf{P}^{\varpi}), \quad \forall \mathbf{P} \in \mathcal{P}^N, j \in M, \varpi : M \rightarrow M \text{ bijection} \quad (\textit{Neutrality}) \end{aligned}$$

## C.2. Costlessness of Requiring Anonymity and Neutrality

With these notions of anonymity and neutrality in mind, observe that for any given mechanism it is possible to construct an anonymous mechanism or a neutral mechanism by randomizing over the roles of the agents or the roles of the alternatives in the mechanism, respectively: let

$$\varphi_j^{\text{anon}}(\mathbf{P}) = \frac{1}{n!} \sum_{\substack{\pi : N \rightarrow N \\ \text{bijection}}} \varphi_j^{\pi}(\mathbf{P}) \quad (21)$$

and

$$\varphi_j^{\text{neut}}(\mathbf{P}) = \frac{1}{m!} \sum_{\substack{\varpi : M \rightarrow M \\ \text{bijection}}} \varphi_j^{\varpi}(\mathbf{P}) \quad (22)$$

for all preference profiles  $\mathbf{P} \in \mathcal{P}^N$  and alternatives  $j \in M$ . Observe that  $\varphi^{\text{anon}}$  and  $\varphi^{\text{neut}}$  are simply hybrid mechanisms with many components, where each component is used

with a probability of  $\frac{1}{n!}$  or  $\frac{1}{m!}$ , respectively. With these definitions we can formally derive the second main takeaway from Theorem 2.

**Corollary 3.** *Given a problem  $(N, M, \delta)$ , where the deficit is anonymous/neutral/both, for any mechanism  $\varphi$  there exists a mechanism  $\tilde{\varphi}$  that is anonymous/neutral/both and has a weakly better signature than  $\varphi$ ; formally,  $\varepsilon(\tilde{\varphi}) \leq \varepsilon(\varphi)$  and  $\delta(\tilde{\varphi}) \leq \delta(\varphi)$*

*Proof.* Observe that if the deficit  $\delta$  is anonymous, then  $\varphi$  and  $\varphi^\pi$  have the same signature for any renaming of the agents  $\pi : N \rightarrow N$ . Consequently,  $\varphi^{\text{anon}}$  is a hybrid of  $n!$  mechanisms which all have the same signature, and therefore it has a weakly better signature by Theorem 2. Similarly, if  $\delta$  is neutral, then  $\varphi$  and  $\varphi^\varpi$  have the same signature for any renaming of the alternatives  $\varpi : M \rightarrow M$ , and the result follows analogously.  $\square$

Intuitively, Corollary 3 means that, given the right desideratum, the two requirements anonymity and neutrality are “free” in terms of manipulability and deficit. We do not have to accept a less attractive signatures in order to achieve either property.

### C.3. Non-trivial Signature Cost of Other Desiderata

In contrast to anonymity and neutrality, other common desiderata do not come for free, such as Condorcet consistency, Pareto optimality, or even the rather weak requirement of unanimity.

**Definition 16** (Unanimity, Pareto Optimality, Condorcet Consistency). For a given preference profile  $\mathbf{P} \in \mathcal{P}^N$ , define the following:

- An alternative  $j \in M$  is a *unanimity winner* at  $\mathbf{P}$  if for all agents  $i \in N$  and all other alternatives  $j' \in M$  we have  $P_i : j \geq j'$ . Let  $M^{\text{unan}}(\mathbf{P})$  be the set of all unanimity winners at  $\mathbf{P}$  and  $M^{-\text{unan}}(\mathbf{P}) = M \setminus M^{\text{unan}}(\mathbf{P})$  the set of non-winners.
- An alternative  $j \in M$  *Pareto dominates* another alternative  $j' \in M$  at  $\mathbf{P}$  if for all agents  $i \in N$  we have  $P_i : j \geq j'$  and for some agent  $i'$  we have  $P_{i'} : j > j'$ .  $j$  is *Pareto optimal* at  $\mathbf{P}$  if there exists no other alternative  $j'$  that Pareto dominates  $j$ . Let  $M^{\text{Pareto}}(\mathbf{P})$  be the set of Pareto optimal alternatives at  $\mathbf{P}$ , and let  $M^{-\text{Pareto}}(\mathbf{P}) = M \setminus M^{\text{Pareto}}(\mathbf{P})$  be the set of alternatives that are Pareto dominated by another alternative at  $\mathbf{P}$ .

- For any two alternatives  $a, b \in M$  let  $n_{a>b}(\mathbf{P}) = \#\{i \in N \mid P_i : a > b\}$  be the number of agents who strictly prefer  $a$  to  $b$  at  $\mathbf{P}$ . An alternative  $j \in M$  is a *Condorcet winner at  $\mathbf{P}$*  if for all other alternatives  $j' \in M$  we have  $n_{j>j'}(\mathbf{P}) \geq n_{j'>j}(\mathbf{P})$ .  $M^{\text{Condorcet}}(\mathbf{P})$  is the set of Condorcet winners at  $\mathbf{P}$ , and the set of non-Condorcet winners is  $M^{-\text{Condorcet}}(\mathbf{P}) = M \setminus M^{\text{Condorcet}}(\mathbf{P})$ .

A mechanism  $\varphi$  satisfies unanimity, Pareto optimality, or Condorcet consistency, if it only selects alternatives that have the respective property whenever they exist. One way to incorporate these desiderata into the linear program FINDOPT is to include them in the objective function by using a desideratum function that assigns higher value to alternatives that have the respective property.

However, this is not sufficient if the optimal mechanisms must satisfy the requirement *completely*, independent of the resulting increase in manipulability. For example, the goal could be to find a unanimous, strategyproof mechanism that minimizes the worst-case deficit based on Veto scoring. In this case, it would not suffice to include unanimity in the desideratum function, because the only way to then guarantee unanimity would be to impose full value maximization. Alternatively, we can incorporate the property as additional constraints in the linear program FINDOPT directly. The following linear constraints can be used to require unanimity, Pareto optimality, and Condorcet consistency, respectively.

LINEAR PROGRAM 1 (FINDOPT), EXTENDED.

$$\begin{aligned}
 & \dots \\
 & \mathbf{f}_j(\mathbf{P}) = 0, && \text{(Unanimity)} \\
 & \quad \forall \mathbf{P} \in \mathcal{P}^N \text{ such that } M^{\text{unan}}(\mathbf{P}) \neq \emptyset \text{ and } j \in M^{-\text{unan}}(\mathbf{P}) \\
 & \mathbf{f}_j(\mathbf{P}) = 0, && \text{(Pareto)} \\
 & \quad \forall \mathbf{P} \in \mathcal{P}^N, j \in M^{-\text{Pareto}}(\mathbf{P}) \\
 & \mathbf{f}_j(\mathbf{P}) = 0, && \text{(Condorcet)} \\
 & \quad \forall \mathbf{P} \in \mathcal{P}^N \text{ such that } M^{\text{Condorcet}}(\mathbf{P}) \neq \emptyset \text{ and } j \in M^{-\text{Condorcet}}(\mathbf{P})
 \end{aligned}$$

The next example illustrates that, unlike anonymity and neutrality, unanimity does not come for free in terms of the mechanisms' signatures. Since Pareto optimality and Condorcet consistency imply unanimity, the same is true for both other desiderata.

**Example 4.** Consider the same problem as in Section 9.2, where  $n = m = 3$ , agents have strict preferences over alternatives  $a, b, c$ ,  $\delta$  is the worst-case deficit based on Veto scoring. Let  $\varphi$  be a strategyproof mechanism that is also unanimous.

By the characterization in (Gibbard, 1977),  $\varphi$  must be a hybrid of strategyproof unilateral and strategyproof duple mechanisms. However, for  $m \geq 3$  alternatives, no duple mechanism is unanimous, and unilateral mechanisms are unanimous only if they always select the first choice of the dictating agent. But as soon as a single component of a hybrid is not unanimous, the hybrid is not unanimous either. Consequently,  $\varphi$  must be a hybrid of dictatorships. Since unanimity is an anonymous and neutral constraint (i.e., the constraint is invariant to renamings of agents or alternatives), we obtain from Corollary 3 that the signature of  $\varphi$  is at most as good as the signature of Random Dictatorship, where the dictating agent is chosen *uniformly* at random. However, Random Dictatorship has a deficit of  $\frac{4}{9}$  at the preference profile

$$P_1, P_2 : a > b > c, \tag{23}$$

$$P_3 : c > b > a. \tag{24}$$

But we already observed that Random Duple is strategyproof and optimal (but not unanimous) in this problem with strictly lower deficit of  $\frac{2}{9}$ . This means that requiring unanimity in this problem leads to a strict increase in the lowest deficit that is achievable with strategyproof, optimal mechanisms.

#### C.4. Symmetric Decomposition of Strategyproof, Anonymous, and Neutral Mechanisms

We present a refinement of Gibbard’s strong characterization of strategyproof mechanisms (Gibbard, 1977). Our *symmetric decomposition* characterizes mechanisms that are *strategyproof*, *anonymous*, and *neutral*. We use this result to establish the shapes of Pareto frontiers in Sections 9.2 analytically. Beyond this application, the symmetric decomposition may be of independent interest.

In the full domain of *strict* preferences, Gibbard (1977) showed that any strategyproof mechanism is a hybrid of multiple “simple” mechanisms, namely strategyproof *unilateral* and *duple* mechanisms.

**Definition 17** (Gibbard, 1977). A mechanism *uni* is *unilateral* if the outcome only depends on the report of a single agent; formally, there exists  $i \in N$  such that for all preference profiles  $\mathbf{P}, \mathbf{P}' \in \mathcal{P}^N$  we have that  $P_i = P'_i$  implies  $\text{uni}(\mathbf{P}) = \text{uni}(\mathbf{P}')$ .

**Definition 18** (Gibbard, 1977). A mechanism *dup* is *duple* if only two alternatives are possible; formally, there exist  $a, b \in M$  such that for all preference profiles  $\mathbf{P} \in \mathcal{P}^N$  we have  $\text{dup}_j(\mathbf{P}) = 0$  for all  $j \neq a, b$ .

The strong characterization result is the following.

**Fact 1** (Gibbard, 1977). A mechanism  $\varphi$  is *strategyproof* if and only if it can be written as a hybrid of mechanisms  $\varphi^1, \dots, \varphi^K$ , and each component  $\varphi^k$  is *strategyproof* and either *unilateral* or *duple*.<sup>7</sup>

Obviously, *duple* mechanisms cannot satisfy *neutrality* (unless  $m = 2$ ) and *unilateral* mechanisms cannot satisfy *anonymity* (unless the mechanism is constant or  $n = 1$ ). This means that *anonymity* and *neutrality* of *strategyproof* mechanisms are the result of mixing the *unilateral* and *duple* components “correctly.” This intuition gives rise to the following more refined decomposition of *strategyproof*, *anonymous*, *neutral* mechanism.

**Theorem 4** (Symmetric Decomposition). A mechanism  $\varphi$  is *strategyproof*, *anonymous*, and *neutral* if and only if there exist

1. *strategyproof*, *neutral*, *unilateral* mechanisms  $\text{uni}_k, k \in \{1, \dots, K^{\text{uni}}\}$ ,
2. *strategyproof*, *anonymous*, *duple* mechanisms  $\text{dup}_k, k \in \{K^{\text{uni}} + 1, \dots, K^{\text{uni}} + K^{\text{dup}}\}$ ,
3. coefficients  $\beta_k \geq 0, k \in \{1, \dots, K^{\text{uni}} + K^{\text{dup}}\}$  with  $\sum_{k=1}^{K^{\text{uni}} + K^{\text{dup}}} \beta_k = 1$ ,

such that

$$\varphi = \sum_{\substack{\pi: N \rightarrow N \\ \text{bijection}}} \sum_{k=1}^{K^{\text{uni}}} \left( \frac{\beta_k}{n!} \right) \text{uni}_k^\pi + \sum_{\substack{\varpi: M \rightarrow M \\ \text{bijection}}} \sum_{k=K^{\text{uni}}+1}^{K^{\text{uni}}+K^{\text{dup}}} \left( \frac{\beta_k}{m!} \right) \text{dup}_k^\varpi. \quad (25)$$

*Proof.* By *anonymity* and *neutrality* of  $\varphi$  we get that

$$\varphi = (\varphi^\pi)^\varpi = (\varphi^\varpi)^\pi = \varphi^{\pi, \varpi} \quad (26)$$

---

<sup>7</sup>Gibbard further refined this result by replacing *strategyproofness* with *localized* and *non-perverse*.

for all bijections  $\pi : N \rightarrow N$  and  $\varpi : M \rightarrow M$ , which implies

$$\varphi = \sum_{\pi, \varpi} \frac{1}{n!m!} \varphi^{\pi, \varpi}. \quad (27)$$

Since  $\varphi^{\pi, \varpi}$  is strategyproof, we can use Fact 1 to decompose it into  $K^{\pi, \varpi}$  strategyproof unilateral and duple mechanisms, i.e.,

$$\varphi^{\pi, \varpi} = \sum_{k=1}^{K^{\pi, \varpi}} \beta_k^{\pi, \varpi} \varphi^{k, \pi, \varpi}. \quad (28)$$

By symmetry, the decomposition can be chosen such that for any pair of renamings  $\pi, \varpi$  and  $k \in \{1, \dots, K^{\pi, \varpi}\}$  we have

- $K^{\text{uni}} + K^{\text{dup}} = K^{\pi, \varpi}$ ,
- $\beta_k = \beta_k^{\pi, \varpi}$ ,
- if  $\varphi^k = \varphi^{k, \text{id}, \text{id}}$  is a strategyproof unilateral (or duple) mechanism, then  $\varphi^{k, \pi, \varpi}$  is a strategyproof unilateral (or duple) mechanism with  $\varphi_{\varpi(j)}^k(\mathbf{P}^{\pi, \varpi}) = \varphi_j^{k, \pi, \varpi}(\mathbf{P})$ ,
- without loss of generality,  $\varphi^k$  is unilateral for  $k \in \{1, \dots, K^{\text{uni}}\}$  and duple for  $k \in \{K^{\text{uni}} + 1, \dots, K^{\text{uni}} + K^{\text{dup}}\}$ .

Averaging over all renamings  $\pi, \varpi$  we get

$$\varphi = \sum_{\pi, \varpi} \left( \frac{1}{n!m!} \right) \sum_{k=1}^{K^{\pi, \varpi}} \beta_k^{\pi, \varpi} \varphi^{k, \pi, \varpi} = \sum_{k=1}^{K^{\text{uni}} + K^{\text{dup}}} \sum_{\pi, \varpi} \left( \frac{\beta_k}{n!m!} \right) \varphi^{k, \pi, \varpi}. \quad (29)$$

If  $\varphi^{k, \text{id}, \varpi}$  is strategyproof and duple, then

$$\text{dup}^{k, \varpi} = \sum_{\pi} \left( \frac{1}{n!} \right) \varphi^{k, \pi, \varpi} \quad (30)$$

is strategyproof and duple as well. Similarly, if  $\varphi^{k, \pi, \text{id}}$  is strategyproof and unilateral, then

$$\text{uni}^{k, \pi} = \sum_{\varpi} \left( \frac{1}{m!} \right) \varphi^{k, \pi, \varpi} \quad (31)$$

is strategyproof and unilateral as well. With this we can rewrite (29) as

$$\varphi = \sum_{k=1}^{K^{\text{uni}}} \sum_{\pi} \frac{\beta_k}{n!} \text{uni}^{k,\pi} + \sum_{k=K^{\text{uni}}+1}^{K^{\text{uni}}+K^{\text{dup}}} \sum_{\varpi} \frac{\beta_k}{m!} \text{dup}^{k,\varpi}, \quad (32)$$

which is precisely the symmetric decomposition (25).  $\square$

The symmetric decomposition (25) is a consequence of Gibbard’s strong characterization and the fact that for any anonymous, neutral mechanism we have  $\varphi = \varphi^{\pi,\varpi}$ . It is *symmetric* in the sense that for any component  $\text{uni}^k$  (or  $\text{dup}^k$ ) that occur with coefficient  $\beta_k$ , the corresponding component  $\text{uni}^{k,\pi}$  (or  $\text{dup}^{k,\varpi}$ ) occur with the same coefficient. (25) decomposes  $\varphi$  into two parts: a neutral part on the left, that gets “anonymized” by randomization, and an anonymous part on the right that gets “neutralized” by randomization.

## C.5. Connection to Gibbard’s Strong Characterization of Strategyproof Random Mechanisms

In Section 9.2 we have studied the Pareto frontier for two different desiderata, namely Plurality and Veto scoring. Juxtaposing the two examples shows an interesting connection to earlier work that characterized strategyproof random mechanisms: Gibbard (1977) showed that in a domain with three or more alternatives and strict preferences, any mechanism that is strategyproof must be a probability mixture (i.e., a hybrid with an arbitrary number of components) of strategyproof *unilateral* mechanisms (i.e., whose random outcomes depend only on the preference order of a single agent) and strategyproof *duple* mechanisms (i.e., which assign positive probability to at most two alternatives).

Our examples show that the optimal strategyproof mechanisms for the particular problems are exactly the two extremes of this representation: when the desideratum is based on Plurality scoring, optimal strategyproof mechanism arise by randomizing over unilateral mechanisms only, namely dictatorships. Conversely, when the desideratum is based on Veto scoring, optimal strategyproof mechanisms arise by randomizing over duple mechanisms only, namely the majority vote between two alternatives. Thus, our two examples teach us in what sense unilateral and duple mechanisms can be understood to reside on “opposite ends” of the spectrum of strategyproof mechanisms.

---

**ALGORITHM 1: FINDBOUNDS**

---

**Input:** bound  $\underline{\varepsilon} > 0$

**Variables:** set of supporting manipulability bounds `supp.bnds`, stacks of *unverified*, *verified*, and *outer* segments `segments.u`, `segments.v`, `segments.o`

**begin**

```
    supp.bnds  $\leftarrow$  {0}
    segments.u  $\leftarrow$  {{signature( $\underline{\varepsilon}$ ), signature(1)}}
    segments.v  $\leftarrow$   $\emptyset$ 
    segments.o  $\leftarrow$  {{signature(0), signature( $\underline{\varepsilon}$ )}, {signature(1), signature(2)}}
    while segments.u  $\neq$   $\emptyset$  do
        ( $P^-, P^+$ )  $\leftarrow$  pop(segments.u)
        ( $P^{--}, P^-$ ), ( $P^+, P^{++}$ )  $\in$  segments.v  $\cup$  segments.u  $\cup$  segments.o
         $e \leftarrow$  (affine.hull({ $P^{--}, P^-$ })  $\cap$  affine.hull({ $P^+, P^{++}$ })) $_{\varepsilon}$ 
         $P \leftarrow$  signature( $e$ )
        if  $P \in$  (affine.hull({ $P^{--}, P^-$ })  $\cap$  affine.hull({ $P^+, P^{++}$ })) then
            | supp.bnds  $\leftarrow$  supp.bnds  $\cup$  { $P$ }
            | segments.v  $\leftarrow$  segments.v  $\cup$  {( $P^{--}, P^-$ ), ( $P^+, P^{++}$ ), ( $P^-, P$ ), ( $P, P^+$ )}
        end
        else if  $P \in$  affine.hull( $P^-, P^+$ ) then
            | segments.v  $\leftarrow$  segments.v  $\cup$  {( $P^-, P^+$ )}
        end
        else
            | segments.u  $\leftarrow$  segments.u  $\cup$  {( $P^-, P$ ), ( $P, P^+$ )}
        end
    end
    return supp.bnds
```

**end**

---

## D. The Algorithm FindBounds

Knowing the structure of the Pareto frontier from Theorem 3, we can take a more refined approach to the problem of identifying the Pareto frontier (than applying FINDOPT to various manipulability bounds). The algorithm FINDBOUNDS (Algorithm 1) applies FINDOPT sequentially to determine the signatures of optimal mechanisms at different manipulability bounds. It uses the information obtained in each step to interpolate the most promising candidates for supporting manipulability bounds.

Before we state our formal result about the correctness, completeness, and runtime of FINDBOUNDS, we provide an outline of how the algorithm works. Observe that by Theorem 3, the problem of computing all supporting manipulability bounds is equivalent

---

**ALGORITHM 2:** FINDLOWER

---

**Variables:** signatures  $\text{sign}^0, \text{sign}, \text{sign}^+$ , bound  $\underline{\varepsilon}$   
**begin**  
   $\underline{\varepsilon} \leftarrow 1/2$   
   $\text{sign}^0 \leftarrow \text{signature}(0), \text{sign}^+ \leftarrow \text{signature}(1), \text{sign} \leftarrow \text{signature}(\underline{\varepsilon})$   
  **while**  $\text{sign} \notin \text{affine.hull}(\text{sign}^0, \text{sign}^+)$  **do**  
     $\text{sign}^+ \leftarrow \text{sign}, \underline{\varepsilon} \leftarrow \underline{\varepsilon}/2, \text{sign} \leftarrow \text{signature}(\underline{\varepsilon})$   
  **end**  
  **return**  $\underline{\varepsilon}$   
**end**

---

to identifying the path of the monotonic, convex, piecewise linear function  $\varepsilon \mapsto \delta(\varepsilon)$ . FINDBOUNDS keeps an inventory of “known signatures” and “verified segments” on this path. It uses these to interpolate and verify new segments.

**Interpolation:** Suppose that we know four points  $s_0 = (\varepsilon_0, \delta_0), s_1 = (\varepsilon_1, \delta_1), s_2 = (\varepsilon_2, \delta_2), s_3 = (\varepsilon_3, \delta_3)$  with  $\varepsilon_0 < \varepsilon_1 < \varepsilon_2 < \varepsilon_3$  on the path. If the two segments  $\overline{s_0 s_1}$  and  $\overline{s_2 s_3}$  are linearly independent, their affine hulls have a unique point of intersection  $\varepsilon'$  over the interval  $[\varepsilon_1, \varepsilon_2]$ . This point of intersection is the candidate for a new supporting manipulability bound that FINDBOUNDS considers. The left plot in Figure 8 illustrates the geometry of this step.

**Verification:** Once we have identified the candidate  $\varepsilon'$ , we use FINDOPT to compute the deficit  $\delta' = \delta(\varepsilon')$ . The signature  $(\varepsilon', \delta')$  is either on or below the straight line connecting the signatures  $(\varepsilon_1, \delta_2)$  and  $(\varepsilon_2, \delta_2)$ . In the first case, if all three signatures lie on a single straight line segment, then we can infer that this line segment must be part of the path of  $\varepsilon \mapsto \delta(\varepsilon)$ . We can mark this segment as “verified,” because there are no supporting manipulability bounds in the open interval  $(\varepsilon_1, \varepsilon_2)$  (but possibly at its limits). In the second case, if  $(\varepsilon', \delta')$  lies strictly below the line segment, then there must exist at least one supporting manipulability bound in the open interval  $(\varepsilon_1, \varepsilon_2)$  at which the path has a real “kink.” The right plot in Figure 8 illustrates the geometry of this step. The signature  $(\varepsilon', \delta')$  must lie somewhere on the gray vertical line at  $\varepsilon'$ .

In this way FINDBOUNDS repeatedly identifies and verifies line segments on the path of  $\varepsilon \mapsto \delta(\varepsilon)$ . If no additional supporting manipulability bounds exist on a segment, then

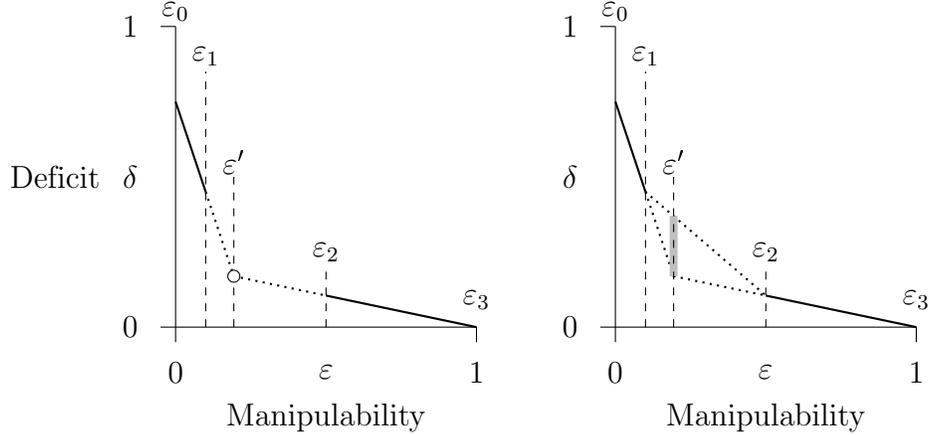


Figure 8: Illustration of interpolation (left) and verification step (right).

this segment is verified. Otherwise, one more point  $(\varepsilon', \delta')$  is added to the collection of known signatures, and the algorithm continues to interpolate.

Next, we prove Proposition 3.

*Proof of Proposition 3.* Given a problem  $(N, M, \delta)$ , the algorithms `FINDLOWER` and `FINDBOUNDS` require at most  $4K + \log_2(1/\varepsilon_1) - 1$  executions of the linear program `FINDOPT` to determine all supporting manipulability bounds of the Pareto frontier, where  $K$  is the number of supporting manipulability bounds and  $\varepsilon_1$  is the smallest non-trivial supporting manipulability bound.<sup>8</sup>

We defined  $signature(\varepsilon)$  (or  $sign(\varepsilon)$  for short) as a function that uses the linear program `FINDOPT` to determine the signature of some mechanism that is optimal at  $\varepsilon$ , i.e.,  $(\varepsilon, \delta(\varepsilon))$ . Algorithm `FINDLOWER` initially calls to the function  $sign$  3 times, for 0, 1, and  $1/2$ . Now suppose that  $\varepsilon_1 \in [\frac{1}{2^n}, \frac{1}{2^{n-1}})$  for some  $n \geq 1$ . Then  $sign$  will be called for  $\frac{1}{4}, \frac{1}{8}, \dots, \frac{1}{2^{n+1}}$  until  $\underline{\varepsilon} = \frac{1}{2^{n+1}}$  is returned. Note that while  $\underline{\varepsilon} = \frac{1}{2^n}$  would have been sufficiently small, the algorithm needs to try  $\underline{\varepsilon} = \frac{1}{2^{n+1}}$  to verify this. Thus, it takes  $n = \log_2(1/\varepsilon_1)$  calls to  $sign$ .

The remainder of the proof is concerned with the algorithm `FINDBOUNDS`. Any *segment* consists of two points  $(sign(\varepsilon), sign(\varepsilon'))$ . Initially, there are two *outer* segments  $(sign(0), sign(\underline{\varepsilon}))$  and  $(sign(1), sign(2))$ , which are needed to help the algorithm get

<sup>8</sup>For technical reasons, `FINDBOUNDS` assumes knowledge of some value  $\underline{\varepsilon} \in (0, \varepsilon_1]$ , i.e., some strictly positive lower bound for the smallest non-trivial supporting manipulability bound. This is needed to initialize the interpolation process. Algorithm `FINDLOWER` (Algorithm 2 identifies a suitable value of  $\underline{\varepsilon}$ .

started. The algorithm maintains a decomposition of the interval  $[\underline{\varepsilon}, 1]$ , which initially consists of a single *unverified* segment. In every execution of the **while**-loop, **FINDBOUNDS** selects an unverified segment  $(\text{sign}(\varepsilon^-), \text{sign}(\varepsilon^+))$ . Then it uses the segments to the left and right  $(\text{sign}(\varepsilon^{--}), \text{sign}(\varepsilon^-))$  and  $(\text{sign}(\varepsilon^+), \text{sign}(\varepsilon^{++}))$  to “guess” the position of a new supporting manipulability bound between  $\varepsilon^-$  and  $\varepsilon^+$ . This guess  $\mathbf{e}$  is the  $\varepsilon$ -value of the point of intersection of the affine hulls of the two segments, i.e.,

$$\mathbf{e} = \left( \text{affine.hull}(\{\text{sign}(\varepsilon^{--}), \text{sign}(\varepsilon^-)\}) \cap \text{affine.hull}(\{\text{sign}(\varepsilon^+), \text{sign}(\varepsilon^{++})\}) \right)_{\varepsilon}, \quad (33)$$

where *affine.hull* denotes the affine hull. This value is unique and lies inside the open interval  $(\varepsilon^-, \varepsilon^+)$  (by Claim 1). Now  $\mathbf{P} = \text{sign}(\mathbf{e})$  is computed using the linear program **FINDOPT**, and one of three cases can occur:

1.  $\mathbf{P}$  may be equal to the unique point of intersection

$$\text{affine.hull}(\{\text{sign}(\varepsilon^{--}), \text{sign}(\varepsilon^-)\}) \cap \text{affine.hull}(\{\text{sign}(\varepsilon^+), \text{sign}(\varepsilon^{++})\}). \quad (34)$$

In this case,  $\mathbf{e}$  is a supporting manipulability bound. Furthermore, the segments  $(\text{sign}(\varepsilon^{--}), \text{sign}(\varepsilon^-))$ ,  $(\text{sign}(\varepsilon^-), \mathbf{P})$ ,  $(\mathbf{P}, \text{sign}(\varepsilon^+))$ , and  $(\text{sign}(\varepsilon^+), \text{sign}(\varepsilon^{++}))$  are all part of the signature of the Pareto frontier (by Claim 3), and there are no other supporting manipulability bound in the interval  $(\varepsilon^{--}, \varepsilon^{++})$ . **FINDBOUNDS** marks the four segments as *verified* and includes  $\mathbf{e}$  in the collection of supporting manipulability bounds.

2.  $\mathbf{P}$  lies in the affine hull of the segment  $(\text{sign}(\varepsilon^-), \text{sign}(\varepsilon^+))$ . Then by Claim 2,

$$\text{Conv}(\{\text{sign}(\varepsilon^-), \mathbf{P}\}) \cup \text{Conv}(\{\mathbf{P}, \text{sign}(\varepsilon^+)\}) \quad (35)$$

is part of the signature of the Pareto frontier, and there are no supporting manipulability bounds inside the interval  $(\varepsilon^-, \varepsilon^+)$ . **FINDBOUNDS** marks the segment  $(\text{sign}(\varepsilon^-), \text{sign}(\varepsilon^+))$  as *verified*.

3. In any other case, **FINDBOUNDS** splits the segment  $(\text{sign}(\varepsilon^-), \text{sign}(\varepsilon^+))$  by creating two new unverified segments

$$(\text{sign}(\varepsilon^-), \mathbf{P}) \text{ and } (\mathbf{P}, \text{sign}(\varepsilon^+)). \quad (36)$$

We first show correctness of `FINDBOUNDS`, then completeness:

**Correctness** `FINDBOUNDS` stops running when there are no more unverified segments.

Assume towards contradiction that there exists a supporting manipulability bound  $\varepsilon \in (\underline{\varepsilon}, 1)$  that has not been identified. Then this supporting manipulability bound lies in some segment  $[\varepsilon^-, \varepsilon^+]$  that was verified.

If the verification happened in a case 1, Claim 3 ensures that there is no other supporting manipulability bound, i.e., the supporting manipulability bound would have been added to the collection during the analysis of the interval  $(\varepsilon^{--}, \varepsilon^{++})$ .

If the verification happened in a case 2, Claim 2 ensures that  $\varepsilon \notin (\varepsilon^-, \varepsilon^+)$ , so that  $\varepsilon = \varepsilon^-$  (without loss of generality we can assume  $\varepsilon = \varepsilon^-$ , the case  $\varepsilon = \varepsilon^+$  is analogous). The segment  $(\text{sign}(\varepsilon^{--}), \text{sign}(\varepsilon^-))$  was not a verified segment at this time, otherwise this would be a case 1. Thus, at some future step some segment  $(\text{sign}(\tilde{\varepsilon}), \text{sign}(\varepsilon^-))$  with a right end-point in  $\text{sign}(\varepsilon^-)$  was verified. But at this step  $\text{sign}(\varepsilon^-)$  was on the intersection of the affine hulls of the segments  $(\text{sign}(\tilde{\varepsilon}), \text{sign}(\tilde{\varepsilon}))$  and  $(\text{sign}(\varepsilon^-), \text{sign}(\varepsilon^+))$ . This creates a case 1, and thus  $\varepsilon$  was identified as a supporting manipulability bound in this step.

**Completeness** It remains to be shown that `FINDBOUNDS` stops at some point. By Claim 4, for ever two adjacent supporting manipulability bounds  $\varepsilon_k, \varepsilon_{k+1}$ , `FINDBOUNDS` computes at most three points  $\text{sign}(\varepsilon')$ ,  $\text{sign}(\varepsilon'')$ ,  $\text{sign}(\varepsilon''')$  with  $\varepsilon', \varepsilon'', \varepsilon''' \in (\varepsilon_k, \varepsilon_{k+1})$ . Since there is a finite number of supporting manipulability bounds, `FINDBOUNDS` loops at most  $3 + 1 = 4$  times per interval, which establishes completeness. The run-time bound follows by observing that there exist  $K - 2$  intervals between the smallest non-trivial bound  $\varepsilon_1$  and the largest bound  $\bar{\varepsilon}$ . However, we may need to check the interval  $(\hat{\varepsilon}, 1)$  if  $\hat{\varepsilon} < 1$ . Thus, using  $\underline{\varepsilon}$  from `FINDLOWER` we require at most  $4K - 4$  executions of `FINDOPT` to run `FINDBOUNDS`.

**Claim 1.**  $\mathbf{e} \in (\varepsilon^-, \varepsilon^+)$

*Proof.* By convexity of  $\varepsilon \mapsto \delta(\varepsilon)$ , we get that  $\mathbf{e} \in [\varepsilon^-, \varepsilon^+]$ .

Now suppose that  $\mathbf{e} = \varepsilon^-$ . Then  $\varepsilon^- \in \text{affine.hull}(\{\text{sign}(\varepsilon^+), \text{sign}(\varepsilon^{++})\})$ . Since  $\varepsilon^- < \varepsilon^+$ , the segments  $(\text{sign}(\varepsilon^-), \text{sign}(\varepsilon^+))$  and  $(\text{sign}(\varepsilon^+), \text{sign}(\varepsilon^{++}))$  would have been verified in a previous step. But this is a contradiction to the assumption that  $(\text{sign}(\varepsilon^-), \text{sign}(\varepsilon^+))$  was an unverified segment.  $\square$

**Claim 2.** For  $0 \leq \varepsilon_{-1} < \varepsilon_0 < \varepsilon_1 \leq 1$ , if

$$\text{sign}(\varepsilon_0) \in \text{Conv}(\{\text{sign}(\varepsilon_{-1}), \text{sign}(\varepsilon_1)\}), \quad (37)$$

then

$$\text{Conv}(\{\text{sign}(\varepsilon_{-1}), \text{sign}(\varepsilon_1)\}) \quad (38)$$

is part of the signature of the Pareto frontier.

*Proof.* Assume towards contradiction that  $\text{Conv}(\{\text{sign}(\varepsilon_{-1}), \text{sign}(\varepsilon_1)\})$  is not part of the signature of the Pareto frontier. Then by convexity of  $\varepsilon \mapsto \delta(\varepsilon)$  there exists  $\gamma \in (0, 1)$  with

$$\delta(\varphi_{(1-\gamma)\varepsilon_{-1} + \gamma\varepsilon_1}) < (1 - \gamma)\delta(\varphi_{\varepsilon_{-1}}) + \gamma\delta(\varphi_{\varepsilon_1}). \quad (39)$$

If  $\varepsilon' = (1 - \gamma)\varepsilon_{-1} + \gamma\varepsilon_1 > \varepsilon_0$ , then for  $\gamma' = \frac{\varepsilon_0 - \varepsilon_{-1}}{\varepsilon' - \varepsilon_{-1}}$  we get

$$(1 - \gamma')\delta(\varphi_{\varepsilon_{-1}}) + \gamma'\delta(\varphi_{\varepsilon'}) < \delta(\varphi_{\varepsilon_0}) = \delta(\varphi_{(1-\gamma')\varepsilon_{-1} + \gamma'\varepsilon'}), \quad (40)$$

a contradiction to convexity of  $\varepsilon \mapsto \delta(\varepsilon)$ . The argument for  $\varepsilon' < \varepsilon_0$  is analogous.  $\square$

**Claim 3.** For  $0 \leq \varepsilon_{-2} < \varepsilon_{-1} < \varepsilon_0 < \varepsilon_1 < \varepsilon_2 \leq 1$ , if

$$\{\text{sign}(\varepsilon)\} = \text{affine.hull}(\{\text{sign}(\varepsilon_{-2}), \text{sign}(\varepsilon_{-1})\}) \cap \text{affine.hull}(\{\text{sign}(\varepsilon_1), \text{sign}(\varepsilon_2)\}), \quad (41)$$

then

$$\text{Conv}(\{\text{sign}(\varepsilon_0), \text{sign}(\varepsilon_{\pm 2})\}) \quad (42)$$

is part of the signature of the Pareto frontier.

*Proof.* The claim follows by applying Claim 2 twice.  $\square$

**Claim 4.** For any two adjacent supporting manipulability bounds  $\varepsilon_k, \varepsilon_{k+1}$ , **FINDBOUNDS** computes at most three points  $\text{sign}(\varepsilon')$ ,  $\text{sign}(\varepsilon'')$ ,  $\text{sign}(\varepsilon''')$  with  $\varepsilon', \varepsilon'', \varepsilon''' \in (\varepsilon_k, \varepsilon_{k+1})$ .

*Proof.* Suppose that  $\text{sign}(\varepsilon')$ ,  $\text{sign}(\varepsilon'')$ , and  $\text{sign}(\varepsilon''')$  are computed in this order. If  $\varepsilon''' < \min(\varepsilon', \varepsilon'')$ , then  $\varepsilon'$  must be a supporting manipulability bound by convexity of  $\varepsilon \mapsto \delta(\varepsilon)$ , which is a contradiction. The same holds if  $\varepsilon''' > \max(\varepsilon', \varepsilon'')$ . If  $\varepsilon''' \in$

$(\min(\varepsilon', \varepsilon''), \max(\varepsilon', \varepsilon''))$ , the segment

$$(\text{sign}(\min(\varepsilon', \varepsilon'')), \text{sign}(\max(\varepsilon', \varepsilon''))) \quad (43)$$

is verified (case (2)). Another guess  $\varepsilon'''$  that lies within  $[\varepsilon_k, \varepsilon_{k+1}]$  involve the segment  $(\text{sign}(\varepsilon'), \text{sign}(\varepsilon''))$ . Thus, by convexity of  $\varepsilon \mapsto \delta(\varepsilon)$ ,  $\varepsilon'''$  is a supporting manipulability bound.  $\square$

This concludes the proof of Proposition 3  $\square$

The interpolation process in `FINDBOUNDS` is set up in such a way that at most 4 executions of `FINDOPT` are required for finding each supporting manipulability bound. Thus, assuming that representatives of  $\text{OPT}(\varepsilon_k)$  must be determined by executing `FINDOPT` at each  $\varepsilon_k$ , `FINDBOUNDS` is at most a constant factor 4 slower than an algorithm that correctly “guesses” all supporting manipulability bounds and uses `FINDOPT` for these only.

## E. Alternative Domains and Desiderata

With the exception of the examples in Sections 9.2, we have formulated our results for the general domain with weak preferences. In the following Section E.1, we explain how our findings continue to hold in many other interesting domains. Subsequently, in Section E.2, we discuss the expressiveness of desideratum functions for the purpose of encoding desiderata.

### E.1. Domain Restrictions

We have proven our results (Theorems 1 through 3 and Proposition 2) for the unrestricted preference domain with indifferences. However, all these results continue to hold when the domain is restricted by only admitting a subset of the possible preference profiles. Formally, let  $\mathcal{P} \subseteq \mathcal{P}^N$  be *any* subset of the set of preference profiles; then all our results still hold, even if we admit only preference profiles from  $\mathcal{P}$ . This extension is possible because none of our proofs make use of the richness of the full domain.

Important domains that can be expressed as domain restrictions  $\mathcal{P} \subset \mathcal{P}^N$  are:

- The *full domain of strict ordinal preferences*, which is the basis for the classical impossibility results pertaining to strategyproofness of deterministic mechanisms (Gibbard, 1973; Satterthwaite, 1975) and random mechanisms (Gibbard, 1977).
- The *assignment domain*, where each agent must be assigned to a single, indivisible object (e.g., a seat at a public school) (Abdulkadiroğlu and Sönmez, 2003).
- The *two-sided matching domain*, where agents of two different types must be matched to each other to form pairs (e.g., men and women, doctors and hospitals, mentors and mentees), and variations of this domain, such as *matching with couples* (Roth, 1984).
- The *combinatorial assignment domain*, where agents receive bundles of indivisible objects (e.g., schedules with multiple courses) (Budish, 2011)

Of course, the fact that our results continue to hold in these domains does not mean that the actual Pareto frontiers will be equal across the different domains. On the contrary, one would expect optimal mechanisms to be highly adapted to the particular domain in which they are designed to operate.

*Remark 7* (Full Utility Assumption). One essential ingredient to our results is the fact that  $\varepsilon$ -approximate strategyproofness can be equivalently expressed via a finite set of linear constraints. Theorem 1, which yields this equivalence, relies on the assumption of *full utility*: given a preference order  $P_i$ , the agent  $i$  can have vNM utility function  $u_i \in U_{P_i}$  that represents this preference order (and is bounded between 0 and 1). A domain restriction that excludes *subsets of utility functions* would violate this condition. For example, suppose that we imposed the additional restriction that the agent’s utility are *normalized* (i.e.,  $\|u_i\|_2 = \sqrt{\sum_{j \in M} u_i(j)^2} = 1$ ). This restriction limits the gain that an agent can possibly obtain from misreporting in a non-linear way. The linear constraints from Theorem 1 would be sufficient for  $\varepsilon$ -approximate strategyproofness. However, they would no longer be necessary, and we would lose *equivalent*.

## E.2. Encoding Desiderata via Desideratum Functions

desideratum functions are a very versatile means to express different design goals that the mechanism designer may have. We have already seen how they can be used to reflect desiderata based on Plurality scoring (Example 1), Condorcet consistency (Examples

2), or Veto scoring (Section 9.2). We now give additional examples that illustrate how desideratum functions can be used to encode common desiderata.

**Binary Desiderata** Many desirable properties of mechanisms are simply properties of the alternatives they select at any given preference profile. For instance, *Pareto optimality* requires the mechanism to choose alternatives that are not Pareto dominated by any other alternative at the respective preference profile. In general, let  $\Theta$  be a desirable property that an alternative can have at some preference profile (e.g., Pareto optimality). Desideratum functions can capture  $\Theta$  by setting

$$d(j, \mathbf{P}) = \begin{cases} 1, & \text{if alternative } j \text{ has property } \Theta \text{ at } \mathbf{P}, \\ 0, & \text{else.} \end{cases} \quad (44)$$

The  $d$ -deficit of a mechanism at a given preference profile,

$$\delta_d(\varphi(\mathbf{P}), \mathbf{P}) = \max_{j \in M} d(j, \mathbf{P}) - \sum_{j \in M} \varphi_j(\mathbf{P}) \cdot d(j, \mathbf{P}), \quad (45)$$

has a straightforward interpretation: it is simply the “probability that an alternative selected by  $\varphi$  at  $\mathbf{P}$  violates property  $\Theta$  at  $\mathbf{P}$ .” Consequently, if  $\delta$  is the worst-case deficit, then the probability that  $\varphi$  selects an alternative that violates  $\Theta$  is at most  $\delta(\varphi)$  across all preference profiles. Alternatively, if  $\delta^{\mathbb{P}}$  is the ex-ante deficit, then  $\delta^{\mathbb{P}}(\varphi)$  is the ex-ante probability (i.e., before the preferences have been reported) that  $\varphi$  selects an alternative that violates  $\Theta$ , given the prior distribution  $\mathbb{P}$  over agents’ preference profiles.

Common binary desiderata include

- *unanimity*: if all agents agree on a first choice, this alternative should be selected,
- *Condorcet consistency*: if some alternatives weakly dominate any other alternatives in a pairwise majority comparison, one of these alternatives should be selected,
- *Pareto optimality*: only undominated alternatives should be selected,
- *egalitarian fairness*: let  $R(j) = \max_{i \in N} \text{rank}_{P_i}(j)$  be the rank of alternative  $j$  in the preference order of the agent who likes it least, then only alternatives with minimal  $R(j)$  should be selected.

While not all of these desiderata are in conflict with strategyproofness individually,

their combinations may be. It is possible to encode combinations of binary properties: given two desideratum functions  $d_1$  and  $d_2$ , the minimum  $d = \min(d_1, d_2)$  is again a desideratum function, and it expresses the desideratum that both properties (encoded by  $d_1$  and  $d_2$ ) should hold simultaneously; and the maximum  $d = \max(d_1, d_2)$  expresses that at least one of the properties should hold.

Even in the absence of a clear-cut objective, a desideratum may be *implicitly* specified via a target mechanism. A *target mechanism*  $\phi$  is a mechanism that one would like to use if there were no concerns about manipulability (Birrell and Pass, 2011). The desideratum function induced by  $\phi$  takes value 1 at the alternatives that are selected by  $\phi$  at the respective preference profiles, or equivalently,  $d(j, \mathbf{P}) = \phi_j(\mathbf{P})$ . An analogous construction is possible for a *target correspondence*  $\Phi$ , which is a mechanism that selects a *set of alternatives*. To reflect a target correspondence, the desideratum function can be chosen as  $d(j, \mathbf{P}) = 1$  if  $j \in \Phi(\mathbf{P})$  and  $d(j, \mathbf{P}) = 0$  otherwise, where we denote by  $\Phi(\mathbf{P})$  the set of alternatives selected by the correspondence  $\Phi$  at  $\mathbf{P}$ .

In the assignment domain, our notion of Pareto optimality turns into *ex-post efficiency*, which is an important baseline requirement for many assignment mechanisms. Featherstone (2011) introduced *v-rank efficiency*, which is a refinement of ex-post efficiency (and implies ordinal efficiency). He showed that for any *rank valuation*  $v$ , the set of  $v$ -rank efficient assignments is equal to the convex hull of the set of deterministic  $v$ -rank efficient assignments. Thus,  $v$ -rank efficiency is representable by a desideratum function by setting  $d(a, \mathbf{P}) = 1$  for all alternatives  $a$  that correspond to a  $v$ -rank efficient deterministic assignment. In two-sided matching, it is often important to select *stable matchings* to prevent blocking pairs from matching outside the mechanism and thereby causing unraveling of the market. Since stability is a property of deterministic matchings, we can encode this desideratum by a desideratum function.

With any such desideratum function in hand, we can then define a notion of deficit and use the results in this paper to understand the respective Pareto frontier.

*Remark 8* (Linearity Assumption). One example of a desideratum that cannot be represented by a desideratum function is the intermediate efficiency requirement of *ordinal efficiency* for assignment mechanisms. To see this, note that ordinal efficiency coincides with ex-post efficiency on deterministic assignments. However, the convex combination of two ordinally efficient assignments is not necessarily again ordinally efficient. In general, to be able to encode a desideratum via a desideratum function, the

societal value of any random outcome  $x$  must depend in a linear fashion on the societal value of the individual alternatives (see Definition 5). Ordinal efficiency violates this requirement.

**Quantifiable Desiderata** In some situations, there exists a quantified notion of the desirability of different alternatives at different preference profiles. Straightforward examples are desiderata based on positional scoring. Formally, a *scoring function*  $v : \{1, \dots, m\} \rightarrow \mathbb{R}$  is a mapping that associates a *score*  $v(r)$  with selecting some agent's  $r$ th choice alternative. Typical examples include *Borda* (i.e.,  $v(r) = m - r$  for all  $r \in \{1, \dots, m\}$ ), *Plurality* (i.e.,  $v(1) = 1$  and  $v(r) = 0$  for  $r \neq 1$ ), and *Veto* (i.e.,  $v(r) = 0$  for all  $r \neq m$  and  $v(m) = -1$ ). The score of an alternative at a preference profile is then determined by summing the scores of this alternative under the different preference orders of the agents. Formally,

$$\text{sc}_v(j, \mathbf{P}) = \sum_{i \in N} v(\text{rank}_{P_i}(j)), \quad (46)$$

where  $\text{rank}_{P_i}(j)$  is the rank of alternative  $j$  in preference order  $P_i$ . For any scoring-based desideratum, the respective desideratum function arises by scaling  $\text{sc}_v$  as follows:

$$d(j, \mathbf{P}) = \begin{cases} \frac{\text{sc}_v(j, \mathbf{P}) - \min_{j' \in M} \text{sc}_v(j', \mathbf{P})}{\max_{j' \in M} \text{sc}_v(j', \mathbf{P}) - \min_{j' \in M} \text{sc}_v(j', \mathbf{P})}, & \text{if } \max_{j' \in M} \text{sc}_v(j', \mathbf{P}) > \min_{j' \in M} \text{sc}_v(j', \mathbf{P}) \\ 0, & \text{else.} \end{cases}$$

In the assignment domain, the rank valuation  $v$  is used to determine the  $v$ -rank value of any assignment. This induces a scoring function on the corresponding alternatives. Consequently, instead of viewing  $v$ -rank efficiency as a binary desideratum, the quantified desideratum could be to maximize the  $v$ -rank value. This natural measure of societal value was put forward by Featherstone (2011).

In quasi-linear domains, a common objective is to maximize the sum of the agents' cardinal utilities. For situations where this aggregate value is interesting, Procaccia and Rosenschein (2006) proposed to maximize a *conservative estimate* of this value instead. This (slightly unorthodox) desideratum is also representable by a desideratum function. Even though the vNM utilities in our model do not have to be comparable across agents, one may choose to maximize the sum of these utilities. However, since agents only report

ordinal preferences, we cannot maximize their utilities directly. Instead, we can maximize a *conservative estimate* for this sum. Formally, the respective desideratum would be to maximize

$$d(j, \mathbf{P}) = \inf_{u_1 \in U_{P_1}, \dots, u_n \in U_{P_n}} \sum_{i \in N} u_i(j). \quad (47)$$

Recall that we bounded utilities between 0 and 1, and let us assume that  $\max_{j \in M} u_i(j) = 1$ . Then the aggregate utility from any outcome  $x \in \Delta(M)$  is lowest if agents have utility (close to) 0 for all except their first choices. Thus, maximizing the conservative lower bound for aggregate utility in our model corresponds to selecting an alternative that is the first choice of as many agents as possible. This is equivalent to the quantitative measure of societal value induced by Plurality scoring.

## F. Omitted Proofs

### F.1. Proof of Theorem 1

*Proof of Theorem 1.* Given a setting  $(N, M)$ , a bound  $\varepsilon \in [0, 1]$ , and a mechanism  $\varphi$ , the following are equivalent:

1.  $\varphi$  is  $\varepsilon$ -approximately strategyproof in  $(N, M)$ .
2. For any agent  $i \in N$ , any preference profile  $(P_i, P_{-i}) \in \mathcal{P}^N$ , any misreport  $P'_i \in \mathcal{P}$ , and any rank  $r \in \{1, \dots, m\}$ , we have

$$\sum_{j \in M: \text{rank}_{P'_i}(j) \leq r} \varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i}) \leq \varepsilon. \quad (48)$$

Fix an agent  $i$ , a preference profile  $(P_i, P_{-i})$ , and a misreport  $P'_i$ . The admissible set of utility functions for agent  $i$  is  $U_{P_i}$ , i.e., all the utilities  $u_i : M \rightarrow [0, 1]$  for which  $u_i(j) \geq u_i(j')$  whenever  $P_i : j \geq j'$ . Let  $B^{\{0,1\}}(P_i)$  denote the set of *binary utilities associated with  $P_i$* , i.e.,

$$B^{\{0,1\}}(P_i) = \{u : M \rightarrow \{0, 1\} \mid u(j) \geq u(j') \text{ whenever } P_i : j \geq j'\}. \quad (49)$$

We first show the direction “ $\Leftarrow$ ,” i.e., the condition 2 implies  $\varepsilon$ -approximate strate-

gyproofness (1). Let  $u \in B^{\{0,1\}}(P_i)$ , then the incentive constraint (3) for this particular utility function has the form

$$\varepsilon(u, (P_i, P_{-i}), P'_i, \varphi) = \sum_{j \in M} u(j) \cdot (\varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i})) \quad (50)$$

$$= \sum_{j \in M: u(j)=1} \varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i}) \quad (51)$$

$$= \sum_{j \in M: \text{rank}_{P_i}(j) \leq r} \varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i}) \quad (52)$$

for some  $k \in \{1, \dots, m\}$ . By (48) from the condition 2, this term is always upper bounded by  $\varepsilon$ . By Lemma 1,  $U_{P_i} \subseteq \text{Conv}(B^{\{0,1\}}(P_i))$ , which means that any  $u_i \in U_{P_i}$  that represents the preference order  $P_i$  can be written as a convex combination of utility functions in  $B^{\{0,1\}}(P_i)$ , i.e.,

$$u_i = \sum_{l=1}^L \alpha_l u^l \quad (53)$$

for  $u^l \in B^{\{0,1\}}(P_i)$  and  $\alpha_l \geq 0$  for all  $l \in \{1, \dots, L\}$  and  $\sum_{l=1}^L \alpha_l = 1$ . By linearity of the incentive constraint (3) we get that

$$\varepsilon(u_i, (P_i, P_{-i}), P'_i, \varphi) = \sum_{l=1}^L \alpha_l \varepsilon(u^l, (P_i, P_{-i}), P'_i, \varphi) \quad (54)$$

$$\leq \sum_{l=1}^L \alpha_l \varepsilon = \varepsilon. \quad (55)$$

This proves the direction “ $\Leftarrow$ .”

Next, we prove the direction “ $\Rightarrow$ ”. Towards contradiction, assume that the constraint (48) is violated for some  $k \in \{1, \dots, m\}$ , i.e.,

$$\sum_{j \in M: \text{rank}_{P_i}(j) \leq k} \varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i}) = \varepsilon + \delta \quad (56)$$

with  $\delta > 0$ . Let  $u \in B^{\{0,1\}}(P_i)$  be the binary utility function with

$$u(j) = \begin{cases} 1, & \text{if } \text{rank}_{P_i}(j) \leq k, \\ 0, & \text{else.} \end{cases} \quad (57)$$

Then

$$\varepsilon(u, (P_i, P_{-i}), P'_i, \varphi) = \varepsilon + \delta. \quad (58)$$

Choose any utility function  $u' \in U_{P_i}$  and let  $\beta = \frac{\delta/2}{\varepsilon + \delta + 1}$ . The utility function constructed by  $\tilde{u} = (1 - \beta)u + \beta u'$  represents  $P_i$  and we have

$$\varepsilon(\tilde{u}, (P_i, P_{-i}), P'_i, \varphi) = (1 - \beta)\varepsilon(u, (P_i, P_{-i}), P'_i, \varphi) + \beta\varepsilon(u', (P_i, P_{-i}), P'_i, \varphi) \quad (59)$$

$$\geq (1 - \beta)(\varepsilon + \delta) - \beta \quad (60)$$

$$= -\beta(\varepsilon + \delta + 1) + (\varepsilon + \delta) = \varepsilon + \delta/2, \quad (61)$$

since the change in utility from manipulation is lower bounded by  $-1$ . Thus, the  $\varepsilon$ -approximate strategyproofness constraint is violated (for the utility function  $\tilde{u}$ , a contradiction).

**Lemma 1.** *For any preference order  $P \in \mathcal{P}$  define the set of binary utilities associated with  $P$  by*

$$B^{\{0,1\}}(P) = \{u : M \rightarrow \{0, 1\} \mid u(j) \geq u(j') \text{ whenever } P : j \geq j'\}. \quad (62)$$

Then  $U_P \subseteq \text{Conv}(B^{\{0,1\}}(P))$ , where  $\text{Conv}(\cdot)$  denotes the convex hull of a set.

*Proof of Lemma 1.* First, suppose that the preference order  $P$  is strict, i.e.,  $P : j > j'$  or  $P : j' > j$  for all  $j \neq j'$ , and without loss of generality,

$$P : j_1 > j_2 > \dots > j_m. \quad (63)$$

In this case,  $B^{\{0,1\}}(P)$  consists of all the functions  $u^r : M \rightarrow \{0, 1\}$  with

$$u^k(j_r) = \begin{cases} 1, & \text{if } r \leq k, \\ 0, & \text{else.} \end{cases} \quad (64)$$

With  $k \in 0, \dots, m$ . Let  $\Delta u(0) = 1 - u(j_1)$ ,  $\Delta u(k) = u(j_k) - u(j_{k+1})$  for all  $k \in \{1, \dots, m-1\}$ , and  $\Delta u(m) = u(j_m)$ . Then represent  $u$  by

$$u(j_r) = \sum_{k=r}^m \Delta u(k). \quad (65)$$

Now we construct the utility function

$$\tilde{u} = \sum_{k=0}^m \Delta u(k) \cdot u^k. \quad (66)$$

First note that  $\sum_{k=0}^m \Delta u(k) = 1$  and  $\Delta u(k) \geq 0$  for all  $k$ , so that  $\tilde{u}$  is a convex combination of elements of  $B^{\{0,1\}}(P)$ . Furthermore, for any  $j_r \in M$ , we have that

$$\tilde{u}(j_r) = \sum_{k=0}^m \Delta u(k) \cdot u^k(j_r) \quad (67)$$

$$= \sum_{k=r}^m \Delta u(k) = u(j_r), \quad (68)$$

$$(69)$$

i.e.,  $u = \tilde{u}$ . This establishes the Lemma for strict preference orders.

For arbitrary preference orders (i.e., with indifferences) the proof can be easily extended by combining all alternatives about which an agent with preference order  $P$  is indifferent into a single virtual alternative. Then we apply the proof for the strict case. The utility functions in  $B^{\{0,1\}}(P)$  will be exactly those that put equal value on the groups of alternatives between which an agent with preference order  $P$  is indifferent. This concludes the proof of the Lemma.  $\square$

This concludes the proof of Theorem 1.  $\square$

## F.2. Proof of Proposition 1

*Proof of Proposition 1.* Given a problem  $(N, M, \delta)$  and a manipulability bound  $\varepsilon \in [0, 1]$ , there exists at least one mechanism that is optimal at  $\varepsilon$ .

A strategyproof mechanism always exists (e.g., the constant mechanism), and any strategyproof mechanism is also  $\varepsilon$ -approximately strategyproof. Thus, the set of candidates for  $\text{OPT}(\varepsilon)$  is never empty. Since the deficit of any mechanism is upper bounded by 1, we get  $\delta(\varepsilon) \leq 1$  for all  $\varepsilon \in [0, 1]$ . Next, for some  $\varepsilon \in [0, 1]$  let

$$\delta(\varepsilon) = \inf \{ \delta(\varphi) \mid \varphi \text{ } \varepsilon\text{-approximately strategyproof} \}. \quad (70)$$

By definition (of the infimum), there exists a sequence of mechanisms  $(\varphi_k)_{k \geq 1}$  such that  $\delta(\varphi^k) \rightarrow \delta(\varepsilon)$  as  $k \rightarrow \infty$ . Since all  $\varphi_k$  are mechanisms, they are uniformly bounded functions from the finite set  $\mathbb{P}^N$  of preference profiles to the compact set  $\Delta(M)$ . Therefore, some sub-sequence of  $(\varphi_k)_{k \geq 1}$  converges uniformly to a limit  $\tilde{\varphi}$ , which must itself be a mechanism.

By Theorem 1,  $\varepsilon$ -approximate strategyproofness is equivalent to a finite set of weak, linear inequalities. Since these are all satisfied by the elements of the sequence  $(\varphi_k)_{k \geq 1}$ , they are also satisfied by  $\tilde{\varphi}$ , i.e.,  $\tilde{\varphi}$  is an  $\varepsilon$ -approximately strategyproof mechanism. The deficit of a mechanism  $\varphi$  is either a maximum (worst-case deficit) or a weighted average (ex-ante deficit) of a finite set of values, but in any case, it is a uniformly continuous projection of  $\varphi$  onto the compact interval  $[0, 1]$ . Thus, by uniform convergence, the deficit of  $\delta(\tilde{\varphi})$  must be the limit of the sequence of deficits  $(\delta(\varphi_k))_{k \geq 1}$ , i.e.,  $\delta(\varepsilon)$ , which yields  $\tilde{\varphi} \in \text{OPT}(\varepsilon)$ .  $\square$

### F.3. Proof of Theorem 2

*Proof of Theorem 2.* Given a problem  $(N, M, \delta)$ , mechanisms  $\varphi, \psi$ , and  $\beta \in [0, 1]$ , we have

$$\varepsilon(h_\beta) \leq (1 - \beta)\varepsilon(\varphi) + \beta\varepsilon(\psi), \quad (71)$$

$$\delta(h_\beta) \leq (1 - \beta)\delta(\varphi) + \beta\delta(\psi). \quad (72)$$

The manipulability of a mechanism is determined by the maximum manipulability across all agents  $i \in N$ , preference profiles  $(P_i, P_{-i}) \in \mathcal{P}^N$ , and misreports  $P'_i \in \mathcal{P}$ . Using

the triangle inequality for this max-operator, we get that

$$\varepsilon(h_\beta) = \max \varepsilon(u, (P_i, P_{-i}), P'_i, h_\beta) \quad (73)$$

$$= \max \sum_{j \in M} u_i(j) \cdot ((h_\beta)_j(P'_i, P_{-i}) - (h_\beta)_j(P_i, P_{-i})) \quad (74)$$

$$\leq (1 - \beta) \max \sum_{j \in M} u_i(j) \cdot (\varphi_j(P'_i, P_{-i}) - \varphi_j(P_i, P_{-i})) \quad (75)$$

$$+ \beta \max \sum_{j \in M} u_i(j) \cdot (\psi_j(P'_i, P_{-i}) - \psi_j(P_i, P_{-i})) \quad (76)$$

$$= (1 - \beta)\varepsilon(\varphi) + \beta\varepsilon(\psi). \quad (77)$$

For worst-case deficit, observe that the extended desideratum function for random outcomes is linear in the first argument, i.e.,  $d((1 - \beta)x + \beta y, \mathbf{P}) = (1 - \beta)d(x, \mathbf{P}) + \beta d(y, \mathbf{P})$  for any outcomes  $x, y$  preference profile  $\mathbf{P}$ . Thus,

$$\delta(h_\beta) = \max \delta(h_\beta(\mathbf{P}), \mathbf{P}) \leq (1 - \beta)\varepsilon(\varphi) + \beta\varepsilon(\psi) \quad (78)$$

follows analogously. Since the ex-ante deficit is simply a single linear function of the mechanism, we get exact equality, i.e.,  $\delta(h_\beta) = (1 - \beta)\varepsilon(\varphi) + \beta\varepsilon(\psi)$ .  $\square$

#### F.4. Proof of Corollary 1

*Proof of Corollary 1.* Given a problem  $(N, M, \delta)$ , the mapping  $\varepsilon \mapsto \delta(\varepsilon)$  is monotonically decreasing and convex.

Monotonicity follows (almost) by definition: observe a mechanism  $\varphi \in \text{OPT}(\varepsilon)$  is also a candidate for  $\text{OPT}(\varepsilon')$  for any larger manipulability bound  $\varepsilon' \geq \varepsilon$ . The only reason for  $\varphi$  to not be optimal at  $\varepsilon'$  is that some other  $\varepsilon'$ -approximately strategyproof mechanism has strictly lower deficit. Thus, the mapping  $\varepsilon \mapsto \delta(\varepsilon)$  is weakly monotonic and decreasing.

To see convexity, assume towards contradiction that the mapping is not convex. Then there must exist bounds  $\varepsilon, \varepsilon' \in [0, 1]$  and  $\beta \in [0, 1]$ , such that for  $\varepsilon_\beta = (1 - \beta)\varepsilon + \beta\varepsilon'$ , we have

$$\delta(\varepsilon_\beta) > (1 - \beta)\delta(\varepsilon) + \beta\delta(\varepsilon'). \quad (79)$$

Let  $\varphi \in \text{OPT}(\varepsilon)$  and  $\varphi' \in \text{OPT}(\varepsilon')$  and consider the hybrid  $h_\beta = (1 - \beta)\varphi + \beta\varphi'$ . By Theorem 2, this hybrid has a manipulability of at most  $\varepsilon_\beta$  and a deficit of at most

$(1 - \beta)\delta(\varepsilon) + \beta\delta(\varepsilon')$ . Thus,

$$\delta(\varepsilon_\beta) \leq \delta(h_\beta) \leq (1 - \beta)\delta(\varepsilon) + \beta\delta(\varepsilon') < \delta(\varepsilon_\beta), \quad (80)$$

a contradiction.  $\square$

## F.5. Proof of Theorem 3

*Proof of Theorem 3.* Given a problem  $(N, M, \delta)$ , there exists a finite set of supporting manipulability bounds

$$\varepsilon_0 = 0 < \varepsilon_1 < \dots < \varepsilon_{K-1} < \varepsilon_K = \bar{\varepsilon}, \quad (81)$$

such that for any  $k \in \{1, \dots, K\}$  and any  $\varepsilon \in [\varepsilon_{k-1}, \varepsilon_k]$  with  $\beta = \frac{\varepsilon - \varepsilon_{k-1}}{\varepsilon_k - \varepsilon_{k-1}}$  we have that

$$\text{OPT}(\varepsilon) = (1 - \beta)\text{OPT}(\varepsilon_{k-1}) + \beta\text{OPT}(\varepsilon_k), \quad (82)$$

$$\delta(\varepsilon) = (1 - \beta)\delta(\varepsilon_{k-1}) + \beta\delta(\varepsilon_k). \quad (83)$$

From Section 6 we know that for each  $\varepsilon \in [0, \bar{\varepsilon}]$  we can write  $\text{OPT}(\varepsilon)$  as the set of solutions to a linear program, i.e.,

$$\text{OPT}(\varepsilon) = \underset{x}{\text{argmin}} \langle v, x \rangle \quad (84)$$

$$\text{s.t.} \quad Dx \leq d, \quad (85)$$

$$Ax \leq \varepsilon, \quad (86)$$

where  $D$  and  $A$  are matrices,  $v$  and  $d$  are vectors,  $\varepsilon$  is a vector with all entries equal to  $\varepsilon$ , and  $x$  is a vector of variables of dimension  $L$ . Observe that  $\varepsilon$  enters the constraints only as the upper bound in a number of linear inequalities. The proof utilizes this characterization of  $\text{OPT}(\varepsilon)$ .

Before we proceed with the proof of Theorem 3, we require a number of definitions. We denote by  $F_\varepsilon$  the set of *feasible points at  $\varepsilon$* , i.e.,

$$F_\varepsilon = \{x \mid Dx \leq d, Ax \leq \varepsilon\}, \quad (87)$$

and we denote by  $S_\varepsilon$  the set of *solutions at  $\varepsilon$* , i.e.,

$$S_\varepsilon = \operatorname{argmin}_{x \in F_\varepsilon} \langle v, x \rangle. \quad (88)$$

A constraint is a row  $C_l$  of either the matrix  $A$  or the matrix  $D$  with the corresponding bound  $c_l$  equal to respective entry of  $d$  or  $\varepsilon$ . A feasible point  $x \in F_\varepsilon$  is an *extreme point of  $F_\varepsilon$*  if there exist  $L$  independent constraints  $C_1, \dots, C_L$  such that

$$C_l x = c_l \quad (89)$$

for all  $l \in \{1, \dots, L\}$ , i.e., the constraints are satisfied with equality at  $x$ .  $x$  is then said to be an extreme point of  $F_\varepsilon$  *with respect to  $(C_1, \dots, C_L)$* . We say that the set of constraints  $(C_1, \dots, C_L)$  is *restrictive at  $\varepsilon$*  if they are independent and there exists an extreme point in  $F_\varepsilon$  with respect to these constraints. Let

$$\mathcal{R}(\varepsilon) = \{(C_1, \dots, C_L) \mid (C_1, \dots, C_L) \text{ is restrictive at } \varepsilon\} \quad (90)$$

be the set of all sets of constraints that are restrictive at  $\varepsilon$ . A set of restrictive constraints  $\mathbf{C} \in \mathcal{R}(\varepsilon)$  is *binding at  $\varepsilon$*  if the extreme point  $x \in F_\varepsilon$  where the constraints of  $\mathbf{C}$  are satisfied with equality is a solution, i.e.,  $x \in S_\varepsilon$ . Let

$$\mathcal{B}(\varepsilon) = \{\mathbf{C} \in \mathcal{R}(\varepsilon) \mid \mathbf{C} \text{ is binding at } \varepsilon\} \quad (91)$$

be the set of all sets of constraints that are binding at  $\varepsilon$ . We denote by  $E(F_\varepsilon)$  and  $E(S_\varepsilon)$  the extreme points of  $F_\varepsilon$  and  $S_\varepsilon$ , respectively.

Observe that since  $F_\varepsilon$  and  $S_\varepsilon$  are polytopes and bounded by finitely many hyperplanes, the extreme points  $E(F_\varepsilon)$  and  $E(S_\varepsilon)$  form minimal  $\mathcal{V}$ -representations of  $F_\varepsilon$  and  $S_\varepsilon$  (see, e.g., p.51ff in (Grünbaum, 2003)). Thus,  $\operatorname{Conv}(E(F_\varepsilon)) = F_\varepsilon$  and  $\operatorname{Conv}(E(S_\varepsilon)) = S_\varepsilon$ . Furthermore, since  $S_\varepsilon \subseteq F_\varepsilon$ , any extreme point of  $S_\varepsilon$  is also an extreme point of  $F_\varepsilon$ , i.e.,  $E(S_\varepsilon) \subseteq E(F_\varepsilon)$ . Finally, each extreme point is uniquely determined by the set of constraints with respect to which it is extreme, i.e., if there exists an extreme point with respect to a set of constraints  $\mathbf{C} \in \mathcal{R}(\varepsilon)$ , then this point is unique.

**Claim 5.** For  $\varepsilon_0, \varepsilon_1 \in [0, \bar{\varepsilon}]$  with  $\varepsilon_0 < \varepsilon_1$ , if  $x_0 \in F_{\varepsilon_0}$  and  $x_1 \in F_{\varepsilon_1}$ , then for any  $\gamma \in [0, 1]$

and  $\varepsilon = (1 - \gamma)\varepsilon_0 + \gamma\varepsilon_1$  we have that

$$x = (1 - \gamma)x_0 + \gamma x_1 \in F_\varepsilon. \quad (92)$$

*Proof.* By assumption,  $D_k x_0 \leq d_k$  and  $D_k x_1 \leq d_k$  for all  $k$ . Thus,

$$D_k((1 - \gamma)x_0 + \gamma x_1) = (1 - \gamma)D_k x_0 + \gamma D_k x_1 \leq d_k. \quad (93)$$

Furthermore,  $A_k x_0 \leq \varepsilon_0$  and  $A_k x_1 \leq \varepsilon_1$  for all  $k$ , which implies

$$A_k((1 - \gamma)x_0 + \gamma x_1) = (1 - \gamma)A_k x_0 + \gamma A_k x_1 \leq (1 - \gamma)\varepsilon_0 + \gamma\varepsilon_1 = \varepsilon. \quad (94)$$

□

**Claim 6.** For  $\varepsilon_0, \varepsilon_1 \in [0, \bar{\varepsilon}]$  with  $\varepsilon_0 < \varepsilon_1$ , if  $\mathbf{C} \in \mathcal{R}(\varepsilon_0)$  and  $\mathbf{C} \in \mathcal{R}(\varepsilon_1)$ , then for any  $\gamma \in [0, 1]$  and  $\varepsilon = (1 - \gamma)\varepsilon_0 + \gamma\varepsilon_1$  we have that  $\mathbf{C} \in \mathcal{R}(\varepsilon)$ , and the  $\gamma$ -convex combination of the extreme points at  $\varepsilon_0$  and  $\varepsilon_1$  with respect to  $\mathbf{C}$  is the unique extreme point at  $\varepsilon$  with respect to  $\mathbf{C}$ .

*Proof.* Since  $\mathbf{C} \in \mathcal{R}(\varepsilon_0)$  and  $\mathbf{C} \in \mathcal{R}(\varepsilon_1)$ , there exist unique extreme points  $x_0 \in E(F_{\varepsilon_0})$  and  $x_1 \in E(F_{\varepsilon_1})$  with respect to  $\mathbf{C}$ . By Claim 5, the point  $x = (1 - \gamma)x_0 + \gamma x_1$  is feasible at  $\varepsilon$ . For any  $l \in \{1, \dots, L\}$  if  $C_l = D_k$  for some  $k$  we have that

$$C_l x = (1 - \gamma)D_k x_0 + \gamma D_k x_1 = (1 - \gamma)d_k + \gamma d_k = d_k = c_l, \quad (95)$$

i.e.,  $x$  satisfies the constraint  $C_l$  with equality. If  $C_l = A_k$ , then the constraint is also satisfied with equality, since

$$C_l x = (1 - \gamma)A_k x_0 + \gamma A_k x_1 = (1 - \gamma)\varepsilon_0 + \gamma\varepsilon_1 = \varepsilon. \quad (96)$$

Consequently,  $x$  is the unique extreme point at  $\varepsilon$  with respect to  $\mathbf{C}$ . This in turn implies that  $\mathbf{C}$  is restrictive at  $\varepsilon$ . □

**Claim 7.** There exists a finite decomposition

$$0 = \varepsilon_0 < \varepsilon_1 < \dots < \varepsilon_K = \bar{\varepsilon} \quad (97)$$

of the interval  $[0, \bar{\varepsilon}]$ , such that on each interval  $[\varepsilon_{k-1}, \varepsilon_k]$  we have that  $\mathcal{R}(\varepsilon) = \mathcal{R}(\varepsilon')$  for all  $\varepsilon, \varepsilon' \in [\varepsilon_{k-1}, \varepsilon_k]$ .

*Proof.* By Claim 6, if some set of  $L$  independent constraints  $\mathbf{C}$  is restrictive at some  $\varepsilon \in [0, \bar{\varepsilon}]$ , then the set of  $\varepsilon' \in [0, \bar{\varepsilon}]$  where  $\mathbf{C}$  is also restrictive must be compact interval  $[\varepsilon_{\mathbf{C}}^-, \varepsilon_{\mathbf{C}}^+] \subseteq [0, \bar{\varepsilon}]$  with  $\varepsilon \in [\varepsilon_{\mathbf{C}}^-, \varepsilon_{\mathbf{C}}^+]$ . Since there is a finite number of constraints, there is also a finite number of constraint sets  $\mathbf{C}$ . Consider the set

$$\{\varepsilon_0, \dots, \varepsilon_K\} = \bigcup_{\mathbf{C} \text{ set of } L \text{ indep. constraints}} \{\varepsilon_{\mathbf{C}}^-, \varepsilon_{\mathbf{C}}^+\}. \quad (98)$$

Observe that by construction, a set of  $L$  independent constraints  $\mathbf{C}$  becomes restrictive or un-restrictive only at one of the finitely many  $\varepsilon_{k-1}$ . This proves the claim.  $\square$

**Claim 8.** *On each interval  $[\varepsilon_{k-1}, \varepsilon_k]$  from Claim 7 and for any  $\varepsilon \in [\varepsilon_{k-1}, \varepsilon_k]$ , if  $\mathbf{C} \in \mathcal{R}(\varepsilon)$ , then  $\mathbf{C} \in \mathcal{R}(\varepsilon_{k-1}) \cap \mathcal{R}(\varepsilon_k)$ .*

*Proof.* Assume towards contradiction that  $\mathbf{C} \in \mathcal{R}(\varepsilon)$ , but  $\mathbf{C} \notin \mathcal{R}(\varepsilon_{k-1})$ . Then there exists an  $\varepsilon' \in (\varepsilon_{k-1}, \varepsilon) \subseteq (\varepsilon_{k-1}, \varepsilon_k)$ , where  $\mathbf{C}$  become restrictive for the first time. Then  $\varepsilon' \in \{\varepsilon_0, \varepsilon_{k-1}\}$ , and therefore  $[\varepsilon_{k-1}, \varepsilon_k]$  cannot be one of the intervals in the decomposition. Instead, it would be split by  $\varepsilon'$ , a contradiction.  $\square$

**Claim 9.** *On each interval  $[\varepsilon_{k-1}, \varepsilon_k]$  from Claim 7 and for any  $\gamma \in [0, 1]$  with  $\varepsilon = (1 - \gamma)\varepsilon_{k-1} + \gamma\varepsilon_k$  we have that*

$$F_\varepsilon = (1 - \gamma)F_{\varepsilon_{k-1}} + \gamma F_{\varepsilon_k}, \quad (99)$$

*i.e., the set of feasible points at  $\varepsilon$  is equal to the  $\gamma$ -convex combination of feasible points at  $\varepsilon_{k-1}$  and  $\varepsilon_k$ .*

*Proof.* By Claim 5 we have

$$F_\varepsilon \supseteq (1 - \gamma)F_{\varepsilon_{k-1}} + \gamma F_{\varepsilon_k}. \quad (100)$$

By Claim 8 the extreme points of  $F_\varepsilon$  are the  $\gamma$ -convex combinations of extreme points of  $F_{\varepsilon_{k-1}}$  and  $F_{\varepsilon_k}$ . Since  $F_\varepsilon = \text{Conv}(E(F_\varepsilon))$ , this implies

$$F_\varepsilon \subseteq (1 - \gamma)F_{\varepsilon_{k-1}} + \gamma F_{\varepsilon_k}, \quad (101)$$

which concludes the proof of the claim.  $\square$

**Claim 10.** *On each interval  $[\varepsilon_{k-1}, \varepsilon_k]$  from Claim 7, if  $\mathbf{C} \in \mathcal{B}(\varepsilon)$  for some  $\varepsilon \in (\varepsilon_{k-1}, \varepsilon_k)$ , then  $\mathbf{C} \in \mathcal{B}(\varepsilon_{k-1}) \cap \mathcal{B}(\varepsilon_k)$ . Furthermore, the extreme point of  $S_\varepsilon$  with respect to  $\mathbf{C}$  is the  $\gamma$ -convex combination of the extreme points of  $S_{\varepsilon_{k-1}}$  and  $S_{\varepsilon_k}$  with respect to  $\mathbf{C}$  with  $\gamma = \frac{\varepsilon - \varepsilon_{k-1}}{\varepsilon_k - \varepsilon_{k-1}}$ .*

*Proof.* Since  $\mathbf{C} \in \mathcal{B}(\varepsilon)$ , there exists a unique extreme point  $x \in E(S_\varepsilon)$  such that  $x$  is extreme at  $\varepsilon$  with respect to  $\mathbf{C}$ . By Claim 9, we can represent  $x = (1 - \gamma)x_0 + \gamma x_1$  with  $x_0 \in F_{\varepsilon_{k-1}}$ ,  $x_1 \in F_{\varepsilon_k}$ . By Claim 6,  $x_0$  and  $x_1$  are extreme points of  $F_{\varepsilon_{k-1}}$  and  $F_{\varepsilon_k}$ , respectively.

Suppose towards contradiction that  $x_0 \notin S_{\varepsilon_{k-1}}$ . Then there exists  $x'_0 \in S_{\varepsilon_{k-1}}$  such that

$$\langle v, x'_0 \rangle < \langle v, x_0 \rangle. \quad (102)$$

By Claim 5,  $x' = (1 - \gamma)x'_0 + \gamma x_1$  is in  $F_\varepsilon$  and

$$\langle v, x' \rangle = (1 - \gamma)\langle v, x'_0 \rangle + \gamma\langle v, x_1 \rangle \quad (103)$$

$$< (1 - \gamma)\langle v, x_0 \rangle + \gamma\langle v, x_1 \rangle \quad (104)$$

$$= \langle v, x \rangle, \quad (105)$$

i.e.,  $x'$  is feasible at  $\varepsilon$  and has lower objective than  $x$ . This contradicts the assumption that  $x \in S_\varepsilon$ . A similar argument yields  $x_1 \in S_{\varepsilon_k}$ , which concludes the proof of the claim.  $\square$

**Claim 11.** *On each interval  $[\varepsilon_{k-1}, \varepsilon_k]$  from Claim 7 and any  $\gamma \in [0, 1]$  with  $\varepsilon = (1 - \gamma)\varepsilon_{k-1} + \gamma\varepsilon_k$  we have that*

$$S_\varepsilon = (1 - \gamma)S_{\varepsilon_{k-1}} + \gamma S_{\varepsilon_k}, \quad (106)$$

*i.e., the set of solutions at  $\varepsilon$  is equal to the  $\gamma$ -convex combination of solutions at  $\varepsilon_{k-1}$  and  $\varepsilon_k$ .*

*Proof.* By Claim 5 we have

$$S_\varepsilon \supseteq (1 - \gamma)S_{\varepsilon_{k-1}} + \gamma S_{\varepsilon_k}. \quad (107)$$

By Claim 10 the extreme points of  $S_\varepsilon$  are the  $\gamma$ -convex combinations of extreme points of  $S_{\varepsilon_{k-1}}$  and  $S_{\varepsilon_k}$ . Since  $S_\varepsilon = \text{Conv}(E(S_\varepsilon))$ , this implies

$$S_\varepsilon \subseteq (1 - \gamma)S_{\varepsilon_{k-1}} + \gamma S_{\varepsilon_k}, \quad (108)$$

which concludes the proof of the claim.  $\square$

Claim 11 is the main step in the proof of Theorem 3: every solution  $x \in S_\varepsilon$  corresponds to some optimal mechanism  $\varphi \in \text{OPT}(\varepsilon)$ . Furthermore, by the nature of the representation of mechanisms in the linear program, the convex combination of two solutions corresponds to the hybrid mechanism of the two mechanisms. Thus,

$$S_\varepsilon = (1 - \gamma)S_{\varepsilon_{k-1}} + \gamma S_{\varepsilon_k}, \quad (109)$$

implies

$$\text{OPT}(\varepsilon) = (1 - \gamma)\text{OPT}(\varepsilon_{k-1}) + \gamma\text{OPT}(\varepsilon_k). \quad (110)$$

Since the objective value  $\mathbf{d}$  is a variable in the solution, we get

$$\delta(\varepsilon) = (1 - \gamma)\delta(\varepsilon_{k-1}) + \gamma\delta(\varepsilon_k). \quad (111)$$

$\square$

## F.6. Proof of Proposition 4

*Proof of Proposition 4.* In a problem  $(N, M, \delta^{\text{PLU}})$  with  $n = 3$  agents,  $m = 3$  alternatives, strict preferences, and where  $\delta^{\text{PLU}}$  is the worst-case  $d^{\text{PLU}}$ -deficit, the following hold:

1. The Pareto frontier has two supporting manipulability bounds 0 and  $1/3$ .
2. Random Dictatorship is a representative of  $\text{OPT}(0)$ .
3. Uniform Plurality is a representative of  $\text{OPT}(1/3)$ .

We first prove that Random Dictatorship is optimal at  $\varepsilon_0 = 0$ . Since Random Dictatorship is by construction a lottery of unilateral, strategyproof mechanisms, it is obviously strategyproof. At any preference profile where all agents have the same first choice, Random Dictatorship will select this alternative and achieve zero deficit. At any

preference profile where all agents have different first choices, all alternatives have the same  $d^{\text{PLU}}$ -value and therefore, any outcome has zero deficit. Finally, consider the case where two agents agree on a first choice,  $a$  say, but the third agent has a different first choice,  $b$  say. In this case, selecting  $a$  would yield a maximal  $d^{\text{PLU}}$ -value of  $\frac{2}{3}$ . However, Random Dictatorship will only select alternative  $a$  with probability  $\frac{2}{3}$  and  $b$  otherwise. This leads to an outcome with  $d^{\text{PLU}}$ -value of  $\frac{2}{3} \cdot \frac{2}{3} + \frac{1}{3} \cdot \frac{1}{3} = \frac{5}{9}$ , and therefore, the deficit of Random Dictatorship is  $\frac{1}{9}$ .

It remains to be shown that any strategyproof mechanism will have a deficit of at least  $\frac{1}{9}$ . Following the discussion of the second takeaway from Theorem 2 in Section 7, we can restrict our attention to mechanisms  $\varphi$  that are strategyproof, anonymous, and neutral. By Theorem 4,  $\varphi$  has a symmetric decomposition, i.e., it can be represented as a lottery over neutral, strategyproof unilaterals and anonymous, strategyproof duples. Suppose, it contains a some anonymous, strategyproof duple  $\text{dup}_{a,b}$ . By the characterization of strategyproofness via swap monotonicity, upper invariance, and lower invariance (Theorem 1 in (Mennle and Seuken, 2017b)), it follows that the outcome of  $\text{dup}_{a,b}$  can only depend on the relative rankings of  $a$  and  $b$ , so that  $\text{dup}_{a,b}$  has the form

$$\text{dup}_{a,b}(\mathbf{P}) = \begin{cases} (p_3, 1 - p_3, 0), & \text{if } P_i : a > b \text{ for all agents } i, \\ (p_2, 1 - p_2, 0), & \text{if } P_i : a > b \text{ for two agents } i, \\ (p_1, 1 - p_1, 0), & \text{if } P_i : a > b \text{ for one agents } i, \\ (p_0, 1 - p_0, 0), & \text{if } P_i : a > b \text{ for one agents } i, \end{cases} \quad (112)$$

where the vector  $(p, 1 - p, 0) = (\varphi_a(\mathbf{P}), \varphi_b(\mathbf{P}), \varphi_c(\mathbf{P}))$  denotes the outcome and  $p_3 \geq p_2 \geq \frac{1}{2}$  and  $p_0 \leq p_1 \leq \frac{1}{2}$ . Again by symmetry of the symmetric decomposition, it must also contain the anonymous duple  $\text{dup}_{a,b}^{\varpi}$  for any permutation of the alternatives  $\varpi : M \rightarrow M$ . Consider the preference profile

$$P_1 : \quad a > b > c, \quad (113)$$

$$P_2 : \quad a > c > b, \quad (114)$$

$$P_3 : \quad b > c > a. \quad (115)$$

The following table shows what outcomes of the different duples  $\{\text{dup}_{a,b}^{\varpi} \mid \varpi : M \rightarrow M \text{ permutation}\}$  at this preference profile.

Duple	$a$	$b$	$c$
dup $_{a,b}$	$p_2$	$1 - p_2$	0
dup $_{a,c}$	$p_2$	0	$1 - p_2$
dup $_{b,c}$	0	$p_2$	$1 - p_2$
dup $_{b,a}$	$1 - p_1$	$p_1$	0
dup $_{c,a}$	$1 - p_1$	0	$p_1$
dup $_{c,b}$	0	$1 - p_1$	$p_1$

A uniform lottery over these duples assigns probability  $2 \cdot \frac{1}{6} = \frac{1}{3}$  to alternative  $b$ . Since  $p_1 \geq p_2$ , this mechanism selects an outcome with  $d^{\text{PLU}}$ -value at most  $\frac{2}{3} \cdot \frac{2}{3} + \frac{1}{3} \cdot \frac{1}{3} = \frac{5}{9}$ . Since the best alternative  $a$  has  $d^{\text{PLU}}$ -value  $\frac{2}{3}$ , the mechanism must have a deficit of at least  $\frac{1}{9}$  at this particular preference profile. This is the same deficit that Random Dictatorship has at this profile, which means that including any strategyproof duples in the symmetric decomposition will not improve the deficit of  $\varphi(\mathbf{P})$ .

Suppose now that the symmetric decomposition of  $\varphi$  contains a neutral, strategyproof unilateral  $\text{uni}_i$ . As before, it follows that  $u_i$  must pick an outcome  $(p_1, p_2, 1 - p_1 - p_2)$  where  $p_1 \geq p_2 \geq 1 - p_1 - p_2$ . Again by symmetry of the symmetric decomposition, it must also contain the neutral unilateral  $\text{uni}_i^\pi$  for any permutation of the agents  $\pi : N \rightarrow N$ . Consider the same preference profile as before, with

$$P_1 : \quad a > b > c, \quad (116)$$

$$P_2 : \quad a > c > b, \quad (117)$$

$$P_3 : \quad b > c > a. \quad (118)$$

The following table shows what outcomes the different unilaterals  $\{\text{uni}_i^\pi \mid \pi : N \rightarrow N \text{ permutation}\}$  will select.

Unilateral	$a$	$b$	$c$
uni $_1$	$p_1$	$p_2$	$1 - p_1 - p_2$
uni $_2$	$p_1$	$1 - p_1 - p_2$	$p_2$
uni $_3$	$1 - p_1 - p_2$	$p_1$	$p_2$

A uniform lottery over these unilaterals assigns probability  $\frac{1}{3}$  to alternative  $b$ , and consequently, this mechanism has a deficit of at least  $\frac{1}{9}$  at this profile. This is the same

deficit that Random Dictatorship has at this profile, which means that including any strategyproof unilaterals in the symmetric decomposition will not improve the deficit of  $\varphi(\mathbf{P})$ .

Since for all mechanisms the worst-case deficit was attained at the same preference profile, there exists no strategyproof, anonymous, neutral mechanism that has a lower deficit at  $\mathbf{P}$  than  $\frac{1}{9}$ , and therefore, Random Dictatorship has minimal deficit among all strategyproof mechanisms.

Next, we show that the manipulability of Uniform Plurality is minimal among all  $d^{\text{PLU}}$ -maximizing mechanisms. Without loss of generality, we restrict attention to anonymous, neutral, and  $d^{\text{PLU}}$ -maximizing mechanisms. At any preference profile where an alternative is ranked first by two agents or more, this alternative is implemented with certainty. Thus, if two agents have the same first choice, the third agent has no opportunity to manipulate, because it cannot change the outcome. The two agents with the same first choice are already receiving their favorite outcome, which makes manipulation useless for them as well.

Thus, any manipulability will arise at a preference profile where all agents have different first choices. Consider the preference profile  $\mathbf{P}$  with

$$P_1 : \quad a > b > c, \tag{119}$$

$$P_2 : \quad b > c > a, \tag{120}$$

$$P_3 : \quad c > a > b. \tag{121}$$

Renaming the alternatives is equivalent to renaming the agents. Thus, an anonymous and neutral mechanisms has to treat all alternatives equally and must therefore select each alternative with probability  $\frac{1}{3}$ . Now suppose that agent 1 is almost indifferent between  $a$  and  $b$ , but strongly dislikes  $c$ , i.e., its utility function is “close to” the binary utility  $u(a) = u(b) = 1, u(c) = 0$ . Then, by swapping  $a$  and  $b$ , agent 1 can enforce the implementation  $b$  with certainty. Its gain from this manipulation is close to

$$1 \cdot u(b) - \frac{1}{3} (u(a) + u(b) + u(c)) = \frac{1}{3}. \tag{122}$$

Thus, any  $d^{\text{PLU}}$ -maximizing mechanism will have manipulability  $\varepsilon(\varphi) \geq \frac{1}{3}$ .

Now consider the Uniform Plurality mechanism. At the above preference profile agents

cannot change the outcome, unless they change their first choice. By anonymity and neutrality, it suffices to show that agent 1 cannot do any better than  $\frac{1}{3}$  by manipulating. However, the only other possible misreport that has any effect on the outcome is to bring  $c$  forward and enforce it as the outcome, which would yield no benefit for agent 1. Consequently, Uniform Plurality has minimal manipulability of  $\frac{1}{3}$ .

Finally, we show that there are no additional supporting manipulability bounds besides 0 and  $1/3$ . So far we have that Random Dictatorship is in  $\text{OPT}(0)$  with a deficit of  $\delta(0) = \frac{1}{9}$ . Furthermore, Uniform Plurality is in  $\text{OPT}(1/3)$  and no  $d^{\text{PLU}}$ -maximizing mechanism has strictly lower manipulability. To complete the proof, we will show that for  $\varepsilon = \frac{1}{6}$ , all optimal mechanisms have deficit  $\frac{1}{18}$ . By convexity of the mapping  $\varepsilon \mapsto \delta(\varepsilon)$ , the signature of the Pareto frontier must therefore be a straight line between  $(0, 1/9)$  and  $(1/3, 0)$ . Considering the performance guarantees for hybrid mechanisms (Theorem 2), this implies optimality of the hybrids of Random Dictatorship and Uniform Plurality.

Suppose that  $\varphi$  is  $\frac{1}{6}$ -approximately strategyproof, anonymous, and neutral. It follows from anonymity and neutrality that at the preference profile

$$P_1 : \quad a > b > c, \quad (123)$$

$$P_2 : \quad b > c > a, \quad (124)$$

$$P_3 : \quad c > a > b, \quad (125)$$

the outcome must be  $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$  for  $a, b, c$ , respectively. If agent 1 changes its report to

$$P'_1 : \quad b > a > c, \quad (126)$$

the outcome changes to  $(\frac{1}{3} - \underline{\varepsilon}_a, \frac{1}{3} + \underline{\varepsilon}_a + \underline{\varepsilon}_c, \frac{1}{3} - \underline{\varepsilon}_c)$  for some values  $\underline{\varepsilon}_a \leq \frac{1}{3}, \underline{\varepsilon}_c \leq \frac{1}{3}$ . Suppose now that agent 1 has a utility close to indifference between  $a$  and  $b$ , i.e., close to  $u(a) = u(b) = 1, u(c) = 0$ . Then the utility gain for agent 1 from misreporting  $P'_1$  is (arbitrarily close to)

$$1 \cdot (-\underline{\varepsilon}_a) + 1 \cdot (\underline{\varepsilon}_a + \underline{\varepsilon}_c) + 0 \cdot (-\underline{\varepsilon}_c) \approx \underline{\varepsilon}_c. \quad (127)$$

If the mechanism is  $\frac{1}{6}$ -approximately strategyproof, then  $\underline{\varepsilon}_c \leq \frac{1}{6}$ . The  $d^{\text{PLU}}$ -value of  $\varphi(\mathbf{P}')$

at  $\mathbf{P}' = (P'_1, P_2, P_3)$  is

$$\langle \varphi(\mathbf{P}'), d^{\text{PLU}}(\cdot, \mathbf{P}') \rangle = \varphi_a(\mathbf{P}') \cdot 0 + \varphi_b(\mathbf{P}') \cdot \frac{2}{3} + \varphi_c(\mathbf{P}') \cdot \frac{1}{3} \quad (128)$$

$$= \frac{2}{3} \left( \frac{1}{3} + \varepsilon_a + \varepsilon_c \right) + \frac{1}{3} \left( \frac{1}{3} - \varepsilon_c \right) \quad (129)$$

$$= \frac{6}{18} + \frac{4}{6}\varepsilon_a + \frac{1}{3}\varepsilon_c \quad (130)$$

$$\leq \frac{6}{18} + \frac{4}{18} + \frac{1}{18} = \frac{11}{18}. \quad (131)$$

However, the only  $d^{\text{PLU}}$ -maximizing alternative at  $\mathbf{P}'$  is  $b$  with  $d^{\text{PLU}}(b, \mathbf{P}') = \frac{2}{3} = \frac{12}{18}$ . Thus, any  $\frac{1}{6}$ -approximately strategyproof mechanism must incur a deficit of at least  $\frac{1}{18}$ .  $\square$

## F.7. Proof of Proposition 5

*Proof of Proposition 5.* In a problem  $(N, M, \delta^{\text{VETO}})$  with  $n = 3$  agents,  $m = 3$  alternatives, strict preferences, and where  $\delta^{\text{VETO}}$  is the worst-case  $d^{\text{Veto}}$ -deficit, the following hold:

1. The Pareto frontier has four supporting manipulability bounds  $0, 1/21, 1/12, 1/2$ .
2. Random Duple is a representative of  $\text{OPT}(0)$ .
3. Uniform Veto is a representative of  $\text{OPT}(1/2)$ .
4. Hybrids of Random Duple and Uniform Veto not optimal for  $\beta \in (0, 1/2)$ .

First, we prove that Random Duple is optimal at  $\varepsilon_0 = 0$ . Since it is a lottery over strategyproof duple mechanisms, it is obviously strategyproof.

At any preference profile where all agents agree on the last choice, Random Duple selects one of the other alternatives, each of which gives maximal  $d^{\text{VETO}}$ -value. At any preference profile where all agents have different last choices, any outcome has zero deficit. Finally, consider a preference profile with

$$P_1, P_2 : \quad \dots > c, \quad (132)$$

$$P_3 : \quad \dots > b. \quad (133)$$

The  $d^{\text{VETO}}$ -value of  $a$  is 1 and the  $d^{\text{VETO}}$ -value of  $b$  is  $\frac{2}{3}$ , so that the maximum  $d^{\text{VETO}}$ -value is 1. The worst case for Random Duple is that agents 1 and 2 rank  $b$  first, in which case  $a$  will be selected with probability  $\frac{1}{3}$  and  $b$  with probability  $\frac{2}{3}$ . Thus, the deficit of Random Duple at  $\mathbf{P}$  is

$$1 - \frac{1}{3} \cdot 1 - \frac{2}{3} \cdot \frac{2}{3} = \frac{2}{9}. \quad (134)$$

In particular, this deficit is attained by Random Duple at the preference profile  $\mathbf{P}$  with

$$P_1, P_2 : \quad b > a > c, \quad (135)$$

$$P_3 : \quad c > a > b. \quad (136)$$

It remains to be proven whether any strategyproof, anonymous, neutral mechanism  $\varphi$  can achieve a lower deficit (where anonymity and neutrality are without loss of generality, similar to the proof of Proposition 4). Consider the preference profile  $\mathbf{P}$  and let  $\text{uni}_i$  be a strategyproof and neutral unilateral component in the symmetric decomposition of  $\varphi$ .  $\text{uni}_i$  must pick an outcome  $(p_1, p_2, 1 - p_1 - p_2)$  with  $p_1 \geq p_2 \geq 1 - p_1 - p_2$ , where  $p_k$  denotes the probability of agent  $i$ 's  $k$ th choice. The symmetric decomposition implies that  $\text{uni}_1$ ,  $\text{uni}_2$ , and  $\text{uni}_3$  are equally likely to be chosen. Analogous to the proof of Proposition 4, we get

Unilateral	$a$	$b$	$c$
$\text{uni}_1$ or $\text{uni}_2$	$p_1$	$p_2$	$1 - p_1 - p_2$
$\text{uni}_3$	$p_2$	$1 - p_1 - p_2$	$p_1$

Thus, alternative  $b$  is selected with probably at least  $\frac{1}{3}$ , which means that the deficit of  $\varphi$  at the preference profile  $\mathbf{P}$  is not reduced by including any unilaterals in the symmetric decomposition.

Similarly, if  $\text{dup}_{a,b}$  is a strategyproof, anonymous duple in the symmetric decomposition of  $\varphi$ , it has the form

$$\text{dup}_{a,b}(\mathbf{P}) = \begin{cases} (p_3, 1 - p_3, 0), & \text{if } P_i : a > b \text{ for all agents } i, \\ (p_2, 1 - p_2, 0), & \text{if } P_i : a > b \text{ for two agents } i, \\ (p_1, 1 - p_1, 0), & \text{if } P_i : a > b \text{ for one agents } i, \\ (p_0, 1 - p_0, 0), & \text{if } P_i : a > b \text{ for one agents } i, \end{cases} \quad (137)$$

where  $p_3 \geq p_2 \geq \frac{1}{2}$  and  $p_0 \leq p_1 \leq \frac{1}{2}$ . Again by symmetry of the symmetric decomposition, it must also contain the anonymous duple  $\text{dup}_{a,b}^{\varpi}$  for any permutation of the alternatives  $\varpi : M \rightarrow M$ . The following table shows what outcomes the different duples  $\{\text{dup}_{a,b}^{\varpi} \mid \varpi : M \rightarrow M \text{ permutation}\}$  will select at  $\mathbf{P}$ .

Duple	$a$	$b$	$c$
$\text{dup}_{a,b}$	$p_1$	$1 - p_1$	0
$\text{dup}_{a,c}$	$p_2$	0	$1 - p_2$
$\text{dup}_{b,c}$	0	$p_2$	$1 - p_2$
$\text{dup}_{b,a}$	$1 - p_2$	$p_2$	0
$\text{dup}_{c,a}$	$1 - p_1$	0	$p_1$
$\text{dup}_{c,b}$	0	$1 - p_1$	$p_1$

Thus, since  $p_2 \geq p_1$ ,  $b$  is selected with probability of at least  $\frac{1}{3}$ , and therefore, including any other duples in the symmetric decomposition of  $\varphi$  will not improve the deficit at  $\mathbf{P}$ . Consequently, there exists no strategyproof mechanism with a lower deficit than Random Duple.

Next, we show that Uniform Veto has minimal manipulability of among all  $d^{\text{VETO}}$ -maximizing mechanisms. First, observe that Uniform Veto is  $\frac{1}{2}$ -approximately strategyproof. To see this, consider the preference profile  $\mathbf{P}$

$$P_1 : \quad a > b > c, \quad (138)$$

$$P_2, P_3 : \quad \dots > c. \quad (139)$$

Uniform Veto selects  $a$  and  $b$  with probability  $\frac{1}{2}$  each. To manipulate, agent 1 can rank  $b$  last and obtain  $a$  with certainty. Its gain from this manipulation would be

$$u_1(a) - \frac{1}{2}(u_1(a) + u_1(b)) = \frac{1}{2}(u_1(a) - u_1(b)), \quad (140)$$

which is at most  $\frac{1}{2}$  for  $u_1(a) = 1$  and  $u_1(b)$  close to 0. Thus, the manipulability of Uniform Veto is at least  $\frac{1}{2}$ .

Suppose now that all agents have different last choices. In this case, Uniform Veto selects any of the alternatives with probability  $\frac{1}{3}$ . By ranking another alternative last, an agent could only ensure the implementation of its third choice with certainty, which

is not a beneficial manipulation.

Finally, suppose that two agents have the same last choice, while a third agent has a different last choice. We have the following cases from the perspective of agent 1.

- Case I:

$$P_1 : \quad a > b > c, \quad (141)$$

$$P_2, P_3 : \quad \dots > a. \quad (142)$$

In this case, Uniform Veto implements  $b$  with certainty. Agent 1 can only enforce  $c$  by ranking  $b$  last, or rank  $a$  last and obtain  $b$  and  $c$  with probabilities  $\frac{1}{2}$  each. Neither of these moves will make agent 1 better off.

- Case II:

$$P_1 : \quad a > b > c, \quad (143)$$

$$P_2, P_3 : \quad \dots > b. \quad (144)$$

In this case, Uniform Veto implements  $a$  with certainty, which is already agent 1's first choice.

- Case III:

$$P_1 : \quad a > b > c, \quad (145)$$

$$P_2 : \quad \dots > b, \quad (146)$$

$$P_3 : \quad \dots > c. \quad (147)$$

In this case, Uniform Veto implements  $a$  with certainty, which is already agent 1's first choice.

- Case IV:

$$P_1 : \quad a > b > c, \quad (148)$$

$$P_2 : \quad \dots > a, \quad (149)$$

$$P_3 : \quad \dots > c. \quad (150)$$

In this case, Uniform Veto implements  $b$  with certainty. By ranking  $b$  last, agent 1 could obtain a probability of  $\frac{1}{3}$  for each alternative instead. Its gain from this manipulation is

$$\frac{1}{3}(u_1(a) + u_1(b) + u_1(c)) - u_1(b) \leq \frac{1}{3} - \frac{2}{3}u_1(b) \leq \frac{1}{3}. \quad (151)$$

Further observe that by renaming agents and alternatives, the above cases I through IV cover all possible constellations with 2 different alternatives ranked as last choices from the perspective of any agent. Thus, Uniform Veto is  $\frac{1}{2}$ -approximately strategyproof.

Having shown that the manipulability of Uniform Veto is exactly  $\frac{1}{2}$ , we now must show that any other  $d^{\text{VETO}}$ -maximizing mechanism  $\varphi$  has manipulability  $\varepsilon(\varphi) \geq \frac{1}{2}$ . We assume without loss of generality that  $\varphi$  is also anonymous and neutral, and we consider the preference profile

$$P_1, P_2 : \quad a \succ b \succ c, \quad (152)$$

$$P_3 : \quad b \succ a \succ c. \quad (153)$$

At this profile,  $\varphi$  has to select  $a$  with some probability  $p_a$  and  $b$  with probability  $1 - p_a$ . If  $p_a \geq \frac{1}{2}$ , agent 3 can rank  $a$  last and enforce selection of  $b$ , the only remaining alternative with full  $d^{\text{VETO}}$ -value. If  $u_3(b) = 1$ ,  $u_3(c) = 0$ , and  $u_3(a)$  is close to 0, its gain will be

$$1 - (1 - p_a) = p_a \geq \frac{1}{2}. \quad (154)$$

If  $p_a < \frac{1}{2}$ , agent 1 can enforce  $a$  by ranking  $b$  last and obtain a gain of

$$1 - p_a > \frac{1}{2} \quad (155)$$

with a similar utility function. Thus, any  $d^{\text{VETO}}$ -maximizing mechanism has manipulability of at least  $\frac{1}{2}$ .

Last, we show that the hybrids of Random Duples and Uniform Veto do not lie on the Pareto frontier, except for the extreme cases. In fact, the signatures of their hybrids form a straight line: observe that the deficit of Random Duple is attained at the preference

Algorithm	$\varepsilon$	$\delta(\varepsilon)$	Comment
FINDLOWER	0	2/9	found s.m.b. at $\varepsilon = 0$
FINDLOWER	1/2	0	
FINDLOWER	1/4	1/12	
FINDLOWER	1/8	1/8	
FINDLOWER	1/16	65/432	
FINDLOWER	1/32	13/72	
FINDLOWER	1/64	29/144	found $\underline{\varepsilon} = 1/64 < \varepsilon_1$
FINDBOUNDS	1/6	1/9	
FINDBOUNDS	1/18	25/162	
FINDBOUNDS	3/47	28/187	
FINDBOUNDS	1/12	5/36	found s.m.b. at $\varepsilon = 1/12$ and $\varepsilon = 1/2$
FINDBOUNDS	1/21	10/63	found s.m.b. at $\varepsilon = 1/21$

Table 1: Executions of the linear program FINDOPT when using FINDLOWER and FINDBOUNDS to determine the Pareto frontier.

profile  $\mathbf{P}$  with

$$P_1, P_2, P_3 : \quad a > b > c. \quad (156)$$

Since Uniform Veto is  $d^{\text{VETO}}$ -maximizing, its deficit is zero at all preference profiles. Consequently, by linearity of the  $d^{\text{VETO}}$ -value  $d(x, \mathbf{P})$  in the outcome  $x$ , the deficit of any hybrid  $h_\beta$  of Random Duple and Uniform Veto is determined by the deficit at  $\mathbf{P}$ . Furthermore, the manipulability of Uniform Veto is highest at the same preference profile if agent 1 swaps  $b$  and  $c$  to enforce  $a$  and has a utility close to  $u_1(a) = 1$ ,  $u_1(c) = 0$ , and  $u_1(b)$  close to 0. This misreport leaves the outcome of Random Duple unchanged. Therefore, the manipulability of any hybrid will also be determined by this preference profile and this potential misreport. By linearity of the incentive constraints from Theorem 1 it is evident that the signatures of the hybrids of Random Duple and Uniform Veto for a straight line between the signatures  $(0, 2/9)$  and  $(1/2, 0)$  of the respective component mechanism. Consequently, if the Pareto frontier is not linear, then these hybrids will not be on the Pareto frontier for any  $\beta \neq 0, 1$ .

To find the supporting manipulability bounds of the Pareto frontier, we used the algorithm FINDLOWER to determine a lower bound for the smallest non-zero supporting manipulability bound and then applied FINDBOUNDS with this value of  $\underline{\varepsilon}$ . Table 1 gives

Preference Profile			$\varepsilon_1 = 1/21$			$\varepsilon_2 = 1/12$		
$P_1$	$P_2$	$P_3$	$a$	$b$	$c$	$a$	$b$	$c$
$a > b > c$	$a > b > c$	$a > b > c$	11/21	3/7	1/21	7/12	1/3	1/12
$a > b > c$	$a > b > c$	$a > c > b$	4/7	3/7	0	2/3	1/4	1/12
$a > b > c$	$a > b > c$	$b > a > c$	11/21	10/21	0	1/2	1/2	0
$a > b > c$	$a > b > c$	$b > c > a$	10/21	11/21	0	5/12	7/12	0
$a > b > c$	$a > b > c$	$c > a > b$	4/7	8/21	1/21	2/3	1/3	0
$a > b > c$	$a > b > c$	$c > b > a$	10/21	11/21	0	5/12	7/12	0
$a > b > c$	$a > c > b$	$a > c > b$	4/7	0	3/7	2/3	1/12	1/4
$a > b > c$	$a > c > b$	$b > a > c$	4/7	3/7	0	7/12	5/12	0
$a > b > c$	$a > c > b$	$b > c > a$	3/7	8/21	4/21	1/3	5/12	1/4
$a > b > c$	$a > c > b$	$c > a > b$	4/7	0	3/7	7/12	0	5/12
$a > b > c$	$a > c > b$	$c > b > a$	3/7	4/21	8/21	1/3	1/4	5/12
$a > b > c$	$b > a > c$	$b > a > c$	10/21	11/21	0	1/2	1/2	0
$a > b > c$	$b > a > c$	$b > c > a$	3/7	4/7	0	5/12	7/12	0
$a > b > c$	$b > a > c$	$c > a > b$	4/7	8/21	1/21	7/12	5/12	0
$a > b > c$	$b > a > c$	$c > b > a$	8/21	4/7	1/21	5/12	7/12	0
$a > b > c$	$b > c > a$	$b > c > a$	1/21	4/7	8/21	0	2/3	1/3
$a > b > c$	$b > c > a$	$c > a > b$	1/3	1/3	1/3	1/3	1/3	1/3
$a > b > c$	$b > c > a$	$c > b > a$	1/21	4/7	8/21	0	7/12	5/12
$a > b > c$	$c > a > b$	$c > a > b$	11/21	0	10/21	7/12	0	5/12
$a > b > c$	$c > a > b$	$c > b > a$	8/21	4/21	3/7	5/12	1/4	1/3
$a > b > c$	$c > b > a$	$c > b > a$	0	11/21	10/21	0	7/12	5/12

Table 2: Optimal mechanisms at  $\varepsilon_1 = 1/21$  and  $\varepsilon_2 = 1/12$  (extended to other preference profiles via the anonymity and neutrality extension).

the signatures on the Pareto frontier that were determined using the *signature*-function in the order in which they were computed.

Table 2 gives two mechanisms that are optimal at  $\varepsilon_1 = 1/21$  and  $\varepsilon_2 = 1/12$ , respectively. For preference profiles that are not listed, rename the agents and alternatives to obtain one of the listed preference profiles, and select the respective outcome (renaming the alternatives again).  $\square$