JSPS Grant-in-Aid for Scientific Research (S)

Central Bank Communication Design

working paper series

No.38 (January 2022)

Robust Voting under Uncertainty

Satoshi Nakada
Shmuel Nitzan
Takashi Ui

Research Project on Central Bank Communication
702 Faculty of Economics, The University of Tokyo,
7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
Tel: +81-3-5841-5595 E-mail: watlab@e.u-tokyo.ac.jp
http://www.centralbank.e.u-tokyo.ac.jp/en/

Robust Voting under Uncertainty*

Satoshi Nakada† Shmuel Nitzan‡ Takashi Ui§

December 2021

Abstract

This paper proposes normative criteria for voting rules under uncertainty about individual preferences to characterize a weighted majority rule (WMR). The criteria stress the significance of responsiveness, i.e., the probability that the social outcome coincides with the realized individual preferences. A voting rule is said to be robust if, for any probability distribution of preferences, the responsiveness of at least one individual is greater than one-half. This condition is equivalent to the seemingly stronger condition requiring that, for any probability distribution of preferences and any deterministic voting rule, the responsiveness of at least one individual is greater than that under the deterministic voting rule. Our main result establishes that a voting rule is robust if and only if it is a WMR without ties. This characterization of a WMR avoiding the worst possible outcomes provides a new complement to the well-known characterization of a WMR achieving the optimal outcomes, i.e., efficiency in the set of all random voting rules.

JEL classification numbers: D71, D81.

Keywords: majority rule, weighted majority rule, responsiveness, belief-free criterion.

---

*We thank the seminar participants at the Workshop on Decision Making and Contest Theory (Ein Gedi, January 2017), Game Theory Workshop, Mathematical Economics at Keio, RUD, JEA Spring Meeting, SING, APET, Tokyo Metropolitan University, Hitotsubashi University, Tokyo University of Science, and Kwansei Gakuin University for their comments or for their suggestions. Ui acknowledges financial support by MEXT, Grant-in-Aid for Scientific Research (grant numbers 15K03348, 26245024, 18H05217). Nakada acknowledges financial support by MEXT, Grant-in-Aid for Scientific Research (grant numbers 16J03250, 19K13651). Nitzan is grateful for Hitotsubashi Institute for Advanced Study, Hitotsubashi University, for enabling the completion of this study. Any update will be available at http://www1.econ.hit-u.ac.jp/oui/rvu.pdf.

†School of Management, Department of Business Economics, Tokyo University of Science; s.nakada0911@gmail.com.

‡Department of Economics, Bar-Ilan University; nitzans@biu.ac.il.

§Department of Economics, Hitotsubashi University; oui@econ.hit-u.ac.jp.

1 Introduction

Consider the choice of a voting rule on a succession of two alternatives (such as "yes" or "no") by a group of individuals. When a voting rule is chosen, the alternatives to come in the future are unknown, and the individuals are uncertain about their future preferences. An individual votes sincerely being concerned with the probability that the outcome agrees with his or her preference, which is referred to as responsiveness (Rae, 1969). More specifically, an individual prefers a voting rule with higher responsiveness because he or she can expect that a favorable alternative is more likely to be chosen. For example, if an individual has a von Neumann-Morgenstern (VNM) utility function such that the utility from the passage of a favorable issue is one and that of an unfavorable issue is zero, then the expected utility equals the responsiveness.

Imagine that a social planner proposes a voting rule under which the responsiveness of every individual is less than one-half. Then, the individuals may unanimously reject this proposal not only because the collective decision reflects minority preferences on average but also because there exists a voting rule that Pareto dominates this rule in terms of responsiveness. To demonstrate it, consider a voting rule whose collective decision always disagrees with that of the original rule, which is referred to as the inverse rule. Then, the responsiveness of every individual under the inverse rule is greater than one-half, i.e., greater than that under the original rule, because the sum of the former responsiveness and the latter responsiveness equals one. Therefore, the social planner regards it as a minimum requirement for a voting rule that the responsiveness of at least one individual should be greater than one-half, or equivalently, greater than the responsiveness under the inverse rule.

To find out whether the minimum requirement is satisfied, the social planner must know the true probability distribution over preferences. In some circumstances, however, not only the true distribution is unknown but also individuals' subjective beliefs are unobservable. Then, the social planner asks the following question: what voting rule satisfies this minimum requirement under any probability distribution of preferences? We call such a voting rule robustly undominated by the inverse rule, or robust for short,[1] and propose it as a normative criterion for voting rules. From the social planner's viewpoint, robustness is understood as a conservative criterion to avoid the worst-case scenario of violating the minimum requirement. In one sense, robustness looks very weak because it concerns the minimum requirement; in

[1]We borrow the term "robustness" from robust statistics, statistics with good performance for data drawn from a wide range of probability distributions (Huber, 1981).

another sense, it looks very strong because the minimum requirement must always be satisfied under any probability distribution of preferences.

This paper introduces two types of robustness criteria and characterizes the classes of voting rules satisfying them. The criteria require that a voting rule should avoid the following worst-case scenarios. The first worst-case scenario is that the true responsiveness of every individual is less than or equal to one-half, or equivalently, the inverse rule weakly Pareto dominates this rule. A voting rule is said to be robust if, for any probability distribution of preferences, it avoids this scenario. It is easily shown that, for any probability distribution of preferences, not only the inverse rule but also any other deterministic voting rule never weakly Pareto dominates a robust rule. Thus, robustness requires a voting rule to be robustly maximal with respect to the Pareto order over the set of deterministic voting rules. By replacing "less than or equal to" with "strictly less than" in the first worst-case scenario, we obtain a slightly more severe scenario. The second worst-case scenario is that the true responsiveness of every individual is strictly less than one-half, or equivalently, the inverse rule strictly Pareto dominates this rule. A voting rule is said to be weakly robust if, for any probability distribution of preferences, it avoids this scenario. Under weakly robust rules, the responsiveness of every individual can be less than or equal to one-half, as long as the responsiveness of at least one individual is equal to one-half. In this case, a collective decision is at best neutral to each individual's choice on average, which is never the case under robust rules.

Our criteria are based upon the notion of Pareto dominance when individuals are assumed to have common beliefs rather than heterogeneous beliefs. This is because Pareto dominance under heterogeneous beliefs is not as normatively compelling as Pareto dominance under common beliefs (Gilboa et al., 2004; Mongin, 2016).[2] To illustrate it, imagine that individuals vote on some amendment and there are two types of individuals. An individual of type 1 prefers the approval and believes that more than two-thirds of the individuals have the same preferences and vote yes. An individual of type 2 prefers the disapproval and believes that more than one-third of the individuals have the same preferences and vote no. When the two-thirds rule is proposed, where the amendment is approved if and only if more than two-thirds of the individuals vote yes, both types of individuals unanimously agree to adopt it because they believe the preferred outcome to be chosen on the basis of incompatible beliefs. As argued by Mongin (2016), such

[2]Blume et al. (2018) compare complete markets and incomplete markets when individuals have heterogeneous beliefs and demonstrate potential social benefits from restrictions on trade that make markets incomplete.

spurious unanimity is normatively inappropriate.

The main result establishes that a voting rule is robust if and only if it is a weighted majority rule (WMR) without ties; we also show that a voting rule is weakly robust if and only if it is a WMR allowing ties with an arbitrary tie-breaking rule. By this result together with the property of WMRs, we find that, for any probability distribution of preferences, there exists no random voting rule that weakly (strictly) Pareto dominates a (weakly) robust rule. The proof is based upon the theorem of alternatives due to von Neumann and Morgenstern (1944), which is also known as a corollary of Farkas' lemma.

We apply the above result to anonymous rules, which are considered to be fair, and obtain the following characterization of robust anonymous rules. A simple majority rule (SMR) is a unique robust anonymous rule when the number of individuals is odd, whereas no anonymous rule is robust when the number of individuals is even. In the latter case, however, a SMR with an anonymous tie-breaking rule is a weakly robust anonymous rule. This implies that we face a trade-off between robustness and anonymity when the number of individuals is even: we must be content with a nonanonymous rule if we require robustness and we must be content with a weakly robust rule if we require anonymity.

To illustrate the difference between robustness and weak robustness as well as the trade-off between robustness and anonymity, assume that the number of individuals is even and consider SMRs. A SMR with some tie-breaking rule is robust if and only if it is represented as a WMR allowing no ties. For example, a SMR with a casting (tie-breaking) vote is a robust nonanonymous rule. On the other hand, a SMR with any tie-breaking rule is weakly robust. In particular, a SMR with the status quo tie-breaking rule (i.e., the status quo is followed whenever there is a tie) is a weakly robust anonymous rule. Because no anonymous rule is robust when the number of individuals is even, the choice between the two SMRs depends upon which criterion to prioritize, and notice that, in the real world, both rules are widely used. A SMR with a casting vote is used by legislatures such as the United States Senate, the Australian House of Representatives, and the National Diet of Japan. A SMR with the status quo tie-breaking rule is used by legislatures such as the New Zealand House of Representatives, the British House of Commons, and the Australian Senate.

Our characterization of a WMR complements the well-known characterization of a WMR due to Rae (1969), Taylor (1969), and Fleurbaey (2008). Their result, which we call the Rae-Taylor-Fleurbaey (RTF) theorem, states that a voting rule is a WMR if and only if it maximizes the corresponding weighted sum of responsiveness over all individuals, where a probability distribution of preferences is assumed to be known. The normative implication of the RTF theorem is efficiency of WMRs under a fixed probability distribution of preferences;[3] that is, a voting rule is efficient in the set of random voting rules in terms of responsiveness if and only if it is a WMR.

We emphasize that efficiency and robustness are distinct sufficient conditions for WMRs. A voting rule is efficient if, under a fixed probability distribution of preferences, any random voting rule never Pareto dominates this rule. In contrast, a voting rule is robust if, under any probability distribution of preferences, the inverse rule never Pareto dominates this rule. Thus, to check efficiency, we need to consider all random voting rules, i.e., all probability distributions of voting rules, ensuring that the optimal outcome is achieved. On the other hand, to check robustness, we need to consider all probability distributions of preferences, ensuring that the worst outcome is avoided. Because robustness and efficiency are different requirements, the sufficiency of robustness for WMRs is not implied by the RTF theorem.

This paper not only contributes to the literature on the axiomatic foundations of a SMR or a WMR (May, 1952; Fishburn, 1973), but also joins a recently growing literature on the worst-case approach to economic design. Most studies in the latter literature, however, have focused on design of a mechanism maximizing the worst-case value of an objective function, where details of the environment are not fully known to the designer. For example, Chung and Ely (2007) consider a revenue maximization problem in a private value auction where the auctioneer does not know the agents' exact belief structures[4] and show that the optimal auction rule is a dominant-strategy mechanism when the auctioneer evaluates rules by their worst-case performance. On the other hand, Carroll (2015) considers a moral hazard problem where the principal does not know the agent's set of possible actions exactly and shows that the optimal contract is linear when the principal evaluates contracts by their worst-case performance.[5]

At the same time, our criterion requiring "Pareto-ranking robustness" is different from the maximin criterion requiring "performance robustness" adopted by the above papers. Mathe-

[3]See also Schmitz and Tröger (2012) and Azrieli and Kim (2014).

[4]This is a standard assumption in robust mechanism design (Bergemann and Morris, 2005).

[5]Other recent examples of economic design with worst-case objectives include Bergemann and Schlag (2011), Yamashita (2015), Bergemann et al. (2016), Carroll (2017), Chen and Li (2018), Auster (2018), Carrasco et al. (2018), Du (2018), and Yamashita and Zhu (2021) among others. See also a recent survey by Carroll (2019). In decision theory, Gilboa and Schmeidler (1989) is a seminal paper.

matically, our criterion is rather close to a Pareto criterion avoiding the problem of spurious unanimity, in particular, the belief-neutral Pareto criterion introduced by Brunnermeier et al. (2014) in their study of financial markets.[6] Let us translate it into our model. Assume that individuals' subjective beliefs are observable to a social planner. Under this assumption, a voting rule is said to be belief-neutral efficient if it is efficient under any belief in the convex hull of the individuals' beliefs. Thus, any convex combination of the individuals' beliefs is considered to be possible in the belief-neutral Pareto criterion, whereas all beliefs are considered to be possible in our criterion, although we focus on the comparison of a deterministic voting rule and its inverse rule. In other words, our criterion is understood as a comprehensively belief-free criterion.[7]

The rest of this paper is organized as follows. Section 2 summarizes properties of WMRs which will be used in the subsequent sections. Section 3 introduces the concepts of robustness. Section 4 characterizes robust rules. Section 5 compares our result and the RTF theorem. We discuss several extensions in Section 6 and conclude the paper in Section 7.

2 Weighted majority rules

Consider a group of individuals $N = \{1, \dots, n\}$ that faces a choice between two alternatives -1 and 1. The choice of individual $i \in N$ is represented by a decision variable $x_i \in \{-1, 1\}$. The choices of the group members are summarized by a decision profile $x = (x_i)_{i \in N}$. Let $\mathcal{X} = \{-1, 1\}^n$ denote the set of all possible decision profiles.

A deterministic voting rule, or a voting rule for short, is a mapping $\phi : \mathcal{X} \to \{-1, 1\}$, which assigns a collective decision $\phi(x) \in \{-1, 1\}$ to each $x \in \mathcal{X}$. The set of all voting rules is denoted by Φ. For $\phi \in \Phi$, let $-\phi \in \Phi$ be a voting rule whose collective decision always disagrees with that of ϕ, i.e., $(-\phi)(x) = -\phi(x)$ for all $x \in \mathcal{X}$. This voting rule is referred to as the inverse rule of ϕ. Although we mainly study a deterministic voting rule, we occasionally consider a random voting rule as well. A random voting rule is a mapping $\bar{\phi} : \mathcal{X} \to [-1, 1]$, which is interpreted as follows: $\bar{\phi}$ assigns a collective decision $+1$ with probability $(1 + \bar{\phi}(x))/2$ and -1 with probability $(1 - \bar{\phi}(x))/2$ to each $x \in \mathcal{X}$. The set of all random voting rules is denoted by $\Delta(\Phi)$.

[6]For other criteria, see Gilboa et al. (2004), Gilboa et al. (2014), and Gayer et al. (2014).

[7]Admittedly, our criterion can be too extreme in some situations. See the discussion in Section 6.

A voting rule $\phi \in \Phi$ is a weighted majority rule (WMR)[8] if there exists a nonzero weight vector[9] $w = (w_i)_{i \in N} \in \mathbb{R}^n$ satisfying

$$\phi(x) = \begin{cases} 1 & \text{if } \sum_{i \in N} w_i x_i > 0, \\ -1 & \text{if } \sum_{i \in N} w_i x_i < 0. \end{cases} \tag{1}$$

A simple majority rule (SMR) is a special case with positive equal weights, i.e., $w_i = w_j > 0$ for all $i, j \in N$. When there is a tie, i.e. $\sum_{i \in N} w_i x_i = 0$, a tie-breaking rule is used to determine a WMR. For example, a SMR requires a tie-breaking rule if n is even. A tie-breaking rule can be either deterministic or random. A WMR with a random tie-breaking rule is a random voting rule $\bar{\phi} \in \Delta(\Phi)$ satisfying (1) (with ϕ replaced by $\bar{\phi}$).

A voting rule $\phi \in \Phi$ is anonymous if it is symmetric in its n variables; that is, $\phi(x) = \phi(x^\pi)$ for each $x \in \mathcal{X}$ and each permutation $\pi : N \to N$, where $x^\pi = (x_{\pi(i)})_{i \in N}$. A SMR is anonymous if n is odd or if n is even and its tie-breaking rule is anonymous, i.e., symmetric in its n variables. A WMR with nonnegative weights is anonymous if and only if it is an anonymous SMR.

The following characterization of WMRs, which is immediate from the definition, plays an important role in the subsequent analysis.

Lemma 1. *A voting rule $\phi \in \Phi$ is a WMR with a weight vector $w \in \mathbb{R}^n$ if and only if*

$$\phi(x) \sum_{i \in N} w_i x_i \geq 0 \textit{ for all } x \in \mathcal{X}.$$

A voting rule $\phi \in \Phi$ is a WMR with a weight vector $w \in \mathbb{R}^n$ allowing no ties if and only if

$$\phi(x) \sum_{i \in N} w_i x_i > 0 \textit{ for all } x \in \mathcal{X}. \tag{2}$$

This lemma states that ϕ is a WMR allowing no ties (allowing ties) if and only if the corresponding weighted sum of $\phi(x)x_i$ over $i \in N$ is positive (nonnegative) for all $x \in \mathcal{X}$ (because the left-hand side of the inequality is $\sum_{i \in N} w_i(\phi(x)x_i)$).

Note that a weight vector w representing a WMR is a solution to the system of linear inequalities in Lemma 1. This observation leads us to the next lemma, which shows that the set of WMRs with nonnegative weights (i.e., $w_i \geq 0$ for all $i \in N$) coincides with that with positive weights (i.e., $w_i > 0$ for all $i \in N$) if there are no ties.

[8]For an axiomatic foundation of weighted majority rules, see Fishburn (1973), Fishburn and Gehrlein (1977), Nitzan and Paroush (1981), Einy and Lehrer (1989), and Taylor and Zwicker (1992) among others.

[9]We allow negative weights, which appear in Proposition 5.

Lemma 2. *A WMR with nonnegative weights allowing no ties is represented as a WMR with positive weights allowing no ties.*

Proof. Let ϕ be a WMR allowing no ties. By Lemma 1, the set of all weight vectors representing ϕ is $W \equiv \{w \in \mathbb{R}^n : \sum_{i\in N} w_i(\phi(x)x_i) > 0 \text{ for all } x \in \mathcal{X}\}$, which is an open convex polyhedron. If there exists a nonnegative weight vector representing ϕ, then W contains a nonnegative vector, so W also contains a positive vector. □

3 Voting under uncertainty

Assume that $x \in \mathcal{X}$ is randomly drawn according to a probability distribution $p \in \Delta(\mathcal{X}) \equiv \{p' \in \mathbb{R}_+^{\mathcal{X}} : \sum_{x\in\mathcal{X}} p'(x) = 1\}$. For a deterministic voting rule $\phi \in \Phi$, let

$$p(\phi(x) = x_i) \equiv p(\{x \in \mathcal{X} : \phi(x) = x_i\}) = \sum_{x\in\mathcal{X}:\phi(x)=x_i} p(x)$$

be the probability that i's choice agrees with the collective decision, which is referred to as responsiveness or the Rae index (Rae, 1969). It is calculated as

$$p(\phi(x) = x_i) = (E_p[\phi(x)x_i] + 1)/2 \tag{3}$$

because

$$E_p[\phi(x)x_i] = \sum_{x:\phi(x)=x_i} p(x) - \sum_{x:\phi(x)\neq x_i} p(x) = 2p(\phi(x) = x_i) - 1. \tag{4}$$

An individual is assumed to vote sincerely and prefer a voting rule with higher responsiveness because he or she expects that a favorable alternative is more likely to be chosen. A discussion and justification of these assumptions appear in Online Appendix A.[10]

We consider the following Pareto order over Φ. We say that a voting rule $\phi \in \Phi$ is weakly Pareto-preferred to $\phi' \in \Phi$ under $p \in \Delta(\mathcal{X})$ if $p(\phi(x) = x_i) \geq p(\phi'(x) = x_i)$ for all $i \in N$. Similarly, we say that ϕ is Pareto-preferred to ϕ' under $p \in \Delta(\mathcal{X})$ if $p(\phi(x) = x_i) \geq p(\phi'(x) = x_i)$ for all $i \in N$ with strict inequality holding for at least one individual, and that ϕ is strictly Pareto-preferred to ϕ' under $p \in \Delta(\mathcal{X})$ if $p(\phi(x) = x_i) > p(\phi'(x) = x_i)$ for all $i \in N$.

[10]We give two alternative definitions of robustness. Redefining robustness as a requirement for a voting mechanism, where individuals can vote insincerely, we show that sincere voting is weakly dominant for every individual if the mechanism satisfies robustness. Redefining robustness using the expected utility of individuals, we provide justification for the original definition which is based upon responsiveness.

Imagine that a social planner proposes a voting rule $\phi \in \Phi$ such that the responsiveness of every individual is less than or equal to one-half:

$$p(\phi(x) = x_i) \leq 1/2 \text{ for all } i \in N. \tag{5}$$

The collective decision under ϕ reflects minority preferences on average. Moreover, there exists a voting rule that is weakly Pareto-preferred to ϕ, which is shown to be the inverse rule $-\phi$ in the next lemma.

Lemma 3. *A deterministic voting rule $\phi \in \Phi$ satisfies* (5) *if and only if the inverse rule $-\phi$ is weakly Pareto-preferred to ϕ under p; that is,*

$$p((-\phi)(x) = x_i) \geq p(\phi(x) = x_i) \textit{ for all } i \in N. \tag{6}$$

Proof. This lemma is an immediate consequence of $p(\phi(x) = x_i) + p((-\phi)(x) = x_i) = p(\phi(x) = x_i) + p(\phi(x) = -x_i) = 1$. □

The above discussion implies that the social planner regards it as a minimum requirement for a voting rule that the responsiveness of at least one individual should be greater than one-half. In fact, this requirement is considered as minimum because, even if ϕ satisfies it, there may exist a deterministic voting rule other than the inverse rule that is weakly Pareto-preferred to ϕ.

If the probability distribution of preferences is unknown, evaluating the responsiveness is impossible. Then, the collective worst-case scenario for adopting ϕ is that the true responsiveness of every individual is less than or equal to one-half, or equivalently, the inverse rule of ϕ is weakly Pareto-preferred to ϕ under the true probability distribution. We say that ϕ is robustly undominated by the inverse rule, or robust for short, if the individuals can avoid this worst-case scenario whatever the true probability distribution is.[11]

Definition 1. A deterministic voting rule $\phi \in \Phi$ is *robust* if the inverse rule of ϕ is not weakly Pareto-preferred to ϕ under any $p \in \Delta(\mathcal{X})$, or equivalently, for each $p \in \Delta(\mathcal{X})$, the responsiveness of at least one individual is strictly greater than one-half.

For example, a WMR with nonnegative weights is robust if there are no ties. In fact, by Lemma 1, $\sum_{i \in N} w_i E_p[\phi(x)x_i] > 0$ for all $p \in \Delta(\mathcal{X})$, so there exists $i \in N$ such that $w_i E_p[\phi(x)x_i] > 0$, i.e., $p(\phi(x) = x_i) > 1/2$.

[11] In Section 6, we discuss several extensions of Definition 1.

As shown in the next lemma, a robust rule is robustly undominated not only by the inverse rule but also by any other deterministic voting rule. In Section 5, we show that a robust rule is robustly undominated by any other random voting rule as well using the main result reported in Section 4.

Lemma 4. *A deterministic voting rule $\phi \in \Phi$ is robust if and only if, for any $p \in \Delta(\mathcal{X})$ and $\phi' \in \Phi$ such that ϕ' is weakly Pareto-preferred to ϕ under p, we must have $\phi'(x) = \phi(x)$ for all $x \in \mathcal{X}$ with $p(x) > 0$.*

Proof. It is enough to show the "only if" part. Seeking a contradiction, suppose that ϕ is robust and there exist p and ϕ' such that ϕ' is weakly Pareto-preferred to ϕ under p and $p(\{x \in \mathcal{X} : \phi'(x) \neq \phi(x)\}) > 0$. Define $p' \in \Delta(\mathcal{X})$ such that

$$p'(x) = \begin{cases} p(x)/\sum_{x':\phi'(x')\neq\phi(x')} p(x') & \text{if } \phi(x) \neq \phi'(x), \\ 0 & \text{otherwise.} \end{cases}$$

Then,

$$E_{p'}[\phi(x)x_i] - E_{p'}[(-\phi)(x)x_i] = (E_p[\phi(x)x_i] - E_p[(\phi'(x))x_i]) / \sum_{x':\phi(x')\neq\phi'(x')} p(x') \leq 0$$

for all $i \in N$. Thus, by (3), the inverse rule $-\phi$ is weakly Pareto-preferred to ϕ under p', which is a contradiction to the robustness of ϕ. □

By Lemma 4, a voting rule ϕ is not robust if there exists $p \in \Delta(\mathcal{X})$ such that ϕ is not maximal with respect to the Pareto order over Φ under p, which is the worst-case Pareto-ranking of ϕ. Thus, in our criterion of robustness, each voting rule is evaluated based upon the worst-case Pareto-ranking.

We also consider a weaker version of robustness by considering the following more severe collective worst-case scenario for adopting ϕ: the true responsiveness of every individual is *strictly* less than one-half, or equivalently, the inverse rule of ϕ is *strictly* Pareto-preferred to ϕ under the true probability distribution. We say that ϕ is weakly robust if the individuals can avoid this worst-case scenario whatever the true distribution is.

Definition 2. A deterministic voting rule $\phi \in \Phi$ is *weakly robust* if the inverse rule of ϕ is not strictly Pareto-preferred to ϕ under any $p \in \Delta(\mathcal{X})$, or equivalently, for each $p \in \Delta(\mathcal{X})$, the responsiveness of at least one individual is greater than or equal to one-half.

If $\phi \in \Phi$ is weakly robust but not robust, there exists $p \in \Delta(\mathcal{X})$ such that $\max_{i \in N} p(\phi(x) = x_i) = 1/2$. If such p is the true probability distribution, a collective decision is at best neutral to each individual's choice on average, which is never the case with robust rules.

For example, a WMR with nonnegative weights is weakly robust even if there are ties. In fact, by Lemma 1, $\sum_{i \in N} w_i E_p[\phi(x)x_i] \geq 0$ for all $p \in \Delta(\mathcal{X})$, so there exists $i \in N$ such that $w_i > 0$ and $E_p[\phi(x)x_i] \geq 0$, i.e., $p(\phi(x) = x_i) \geq 1/2$.

A weakly robust rule also has the property analogous to that of a robust rule in Lemma 4.

Lemma 5. *A deterministic voting rule ϕ is weakly robust if and only if, for any $p \in \Delta(\mathcal{X})$, no deterministic voting rule is strictly Pareto-preferred to ϕ under p.*

Proof. The proof is essentially the same as that of Lemma 4, so it is omitted. □

The following numerical example illustrates the requirement of robustness.

Example 1. Assume that $n = 5$ and let ϕ be a SMR with veto power of individual 1:

$$\phi(x) = \begin{cases} -x_1 & \text{if } x \in \mathcal{A}, \\ x_1 & \text{otherwise,} \end{cases} \tag{7}$$

where $\mathcal{A} \equiv \{x \in \mathcal{X} : x_1 = 1, \ \#\{i : x_i = 1\} \leq 2\}$. Then, ϕ is not weakly robust because the inverse rule $-\phi$ is strictly Pareto-preferred to ϕ under $p \in \Delta(\mathcal{X})$ given by

$$p(x) = \begin{cases} 1/16 & \text{if } x = (1, -1, -1, -1, -1), \\ 2/16 & \text{if } x \in \{(1, 1, -1, -1, -1), (1, -1, 1, -1, -1), (1, -1, -1, 1, -1), (1, -1, -1, -1, 1)\}, \\ 7/16 & \text{if } x = (-1, 1, 1, 1, 1), \\ 0 & \text{otherwise.} \end{cases}$$

In fact, $E_p[\phi(x)x_i] = 7/16 - 9/16 = -1/8 < 0$ for all i, and thus $E_p[-\phi(x)x_i] = -E_p[\phi(x)x_i] > E_p[\phi(x)x_i]$ for all i. This implies that ϕ does not satisfy the requirement of weak robustness as well as that of robustness by (3).

4 Main results

In this section, we present our main result characterizing robust and weakly robust rules. The results are stated and discussed in Section 4.1. The proofs are given in Section 4.2.

4.1 Characterizations

First, we characterize robust rules. A WMR with nonnegative weights allowing no ties is robust as discussed in Section 3. Our first main result establishes that every robust rule must be such a WMR, which is robustly undominated not only by any other deterministic voting rule but also by any other random voting rule as will be discussed in the next section.

Proposition 1. *A deterministic voting rule is robust if and only if it is a WMR with nonnegative weights such that there are no ties.*

Next, we characterize weakly robust rules. A WMR with nonnegative weights is weakly robust even if there are ties as discussed in Section 3. Our second main result establishes that every weakly robust rule must be such a WMR.

Proposition 2. *A deterministic voting rule is weakly robust if and only if it is a WMR with nonnegative weights.*

As a corollary of Proposition 1, we characterize robust anonymous rules. If n is odd, a SMR is the unique rule that is both robust and anonymous. In Appendix B.1, we give another characterization of a SMR with odd n using a stronger version of robustness. However, if n is even, no anonymous rule is robust. That is, there is a trade-off between robustness and anonymity.

Corollary 3. *Suppose that n is odd. Then, a voting rule is robust and anonymous if and only if it is a SMR. Suppose that n is even. Then, no voting rule is both robust and anonymous.*

Proof. A voting rule is robust and anonymous if and only if it is an anonymous WMR with nonnegative weights allowing no ties, which is a SMR with odd n. □

Corollary 3 implies that an anonymous rule is not robust if n is even or if it is not a SMR. In Appendix A, we introduce robustness of random voting rules and demonstrate that not only anonymous deterministic rules but also anonymous random rules cannot be robust when n is even.

For example, a supermajority rule is not robust because it is anonymous. To illustrate it by a numerical example, consider a two-thirds rule with a very large number of individuals. Suppose that $x_i = 1$ with probability $p \in (1/2, 2/3)$ independently and identically for each $i \in N$. By the law of large numbers, the group decision is -1 with probability close to one, so responsiveness of each individual is close to $1 - p < 1/2$, which implies that this rule is not robust.[12]

As a corollary of Proposition 2, we characterize weakly robust anonymous rules. Although no anonymous rule is robust when n is even, there exists a weakly robust anonymous rule regardless of n, which is an anonymous SMR (a SMR with an anonymous tie-breaking rule).

Corollary 4. *A voting rule is weakly robust and anonymous if and only if it is an anonymous SMR.*

Proof. A voting rule is weakly robust and anonymous if and only if it is an anonymous WMR with nonnegative weights, which is an anonymous SMR. □

To illustrate the difference between robustness and weak robustness as well as the trade-off between robustness and anonymity, suppose that n is even and consider SMRs. By Proposition 1, a SMR with some tie-breaking rule is robust if and only if it is represented as a WMR allowing no ties and, by Corollary 3, such a SMR is nonanonymous. For example, a SMR with a casting (tie-breaking) vote is a robust nonanonymous rule. To see this, we consider two cases. First, assume that the presiding officer with a casting vote is a member of a group of n individuals. This rule is equivalent to a WMR such that the presiding officer's weight is slightly greater than the others' weights. Next, assume that the presiding officer is not a member of a group of n individuals and that he or she votes only when there is a tie. This rule is equivalent to a WMR with $n + 1$ individuals including the presiding officer such that the presiding officer's weight is very close to zero. Each of these WMRs does not have ties and is robust

On the other hand, a SMR with any tie-breaking rule is weakly robust by Proposition 2. In particular, a SMR with the status quo tie-breaking rule (i.e., the status quo is followed whenever there is a tie) is a weakly robust anonymous rule, but it is not robust by Corollary 3.

4.2 Proofs

This section provides the proofs of Propositions 1 and 2. In the proofs, we use the following inequality symbols. For vectors ξ and η, we write $\xi \geq \eta$ if $\xi_i \geq \eta_i$ for each i, $\xi > \eta$ if $\xi_i \geq \eta_i$ for each i and $\xi \neq \eta$, and $\xi \gg \eta$ if $\xi_i > \eta_i$ for each i.

[12]Even if n is not so large, we can find $p \in \Delta(\mathcal{X})$ such that the responsiveness of every individual is less than or equal to one-half.

We enumerate elements in $\mathcal{X}$ as $\{x^j\}_{j \in M}$, where $M \equiv \{1, \ldots, m\}$ is an index set with $m \equiv 2^n$. Consider an $n \times m$ matrix

$$L = [l_{ij}]_{n \times m} = \left[\phi(x^j) x_i^j\right]_{n \times m}.$$

Note that l_{ij} equals $+1$ if i's choice agrees with the collective decision and -1 otherwise. Using this matrix, we can restate the conditions in Proposition 1 as follows.

(a) By Lemma 1, a voting rule ϕ is a WMR with nonnegative weights allowing no ties if and only if there exists $w = (w_i)_{i \in N} \geq 0$ such that

$$\sum_{i \in N} w_i l_{ij} = \sum_{i \in N} w_i \big(\phi(x^j) x_i^j\big) > 0$$

for each $j \in M$, or equivalently, $w^\top L \gg 0$.

(b) By definition and (4), a voting rule is *not* robust if and only if there exists $p = (p_j)_{j \in M} > 0$ such that

$$\sum_{j \in M} l_{ij} p_j = \sum_{j : \phi(x^j) = x_i^j} p_j - \sum_{j : \phi(x^j) \neq x_i^j} p_j \leq 0$$

for each $i \in N$, or equivalently, $Lp \leq 0$.

Proposition 1 states that exactly one of (a) and (b) holds. The following theorem of alternatives due to von Neumann and Morgenstern (1944)[13] guarantees that this is true. The same result also appears in the work of Gale (1960, Theorem 2.10) as a corollary of Farkas' lemma.

Lemma 6. *Let A be an $n \times m$ matrix. Exactly one of the following alternatives holds.*

- *There exists $\xi \in \mathbb{R}^n$ satisfying*

$$\xi^\top A \gg 0, \ \xi \geq 0.$$

- *There exists $\eta \in \mathbb{R}^m$ satisfying*

$$A\eta \leq 0, \ \eta > 0.$$

Proof of Proposition 1. Plug L, w, and p into A, ξ, and η in Lemma 6, respectively. Then, Lemma 6 implies that exactly one of (a) and (b) holds. □

[13] von Neumann and Morgenstern (1944) use this result to prove the minimax theorem.

We can interpret Lemma 6 as a corollary of the fundamental theorem of asset pricing,[14] which is equivalent to Farkas' lemma. Thus, we can explain why Proposition 1 is true in terms of arbitrage-free pricing in an imaginary asset market, which is discussed in Appendix C.[15]

We can prove Proposition 2 similarly by restating the conditions in the proposition as follows.

(a′) By Lemma 1, a voting rule ϕ is a WMR with nonnegative weights if and only if there exists $w = (w_i)_{i \in N} > 0$ such that

$$\sum_{i \in N} w_i l_{ij} = \sum_{i \in N} w_i \big(\phi(x^j) x_i^j\big) \geq 0$$

for each $j \in M$, or equivalently, $w^\top L \geq 0$.

(b′) By definition, a voting rule is *not* weakly robust if and only if there exists $p = (p_j)_{j \in M} > 0$ such that

$$\sum_{j \in N} l_{ij} p_j = \sum_{j : \phi(x^j) = x_i^j} p_j - \sum_{j : \phi(x^j) \neq x_i^j} p_j < 0$$

for each $i \in N$, or equivalently, $Lp \ll 0$.

Proposition 2 states that exactly one of (a′) and (b′) holds. To prove it, we use Lemma 6 again, but in another way.

Proof of Proposition 2. Plug $-L^\top$, p, and w into A, ξ, and η in Lemma 6, respectively, where we replace (n, m) with (m, n). Then, Lemma 6 implies that exactly one of (a′) and (b′) holds. □

5 Robustness vs. efficiency

Rae (1969) and Taylor (1969) were the first to use responsiveness to characterize voting rules, followed by Straffin (1977) and Fleurbaey (2008). Using their characterization, we show that a robust rule is robustly undominated not only by any other deterministic voting rule but also by any other random voting rule. We also clarify the connection between the requirement of robustness and that of efficiency.

[14] For details on the fundamental theorem of asset pricing, see Dybvig and Ross (2003, 2008) and references therein.

[15] Gilboa and Samuelson (2021) ask which collections of uncertain-act evaluations can be simultaneously justified under the maxmin expected utility criterion and draw connections to the fundamental theorem of asset pricing.

Note that, by Lemma 1, $\phi \in \Phi$ is a WMR with a weight vector $w \in \mathbb{R}^n$ if and only if $\phi(x)\sum_{i\in N} w_i x_i = |\phi(x)\sum_{i\in N} w_i x_i|$ for all $x \in \mathcal{X}$, which is equivalent to the following inequality: for all $\phi' \in \Phi$ and $x \in \mathcal{X}$,

$$\phi(x)\sum_{i\in N} w_i x_i = \left|\phi(x)\sum_{i\in N} w_i x_i\right| = \left|\phi'(x)\sum_{i\in N} w_i x_i\right| \geq \phi'(x)\sum_{i\in N} w_i x_i. \tag{8}$$

This is true if and only if, for all $p \in \Delta(\mathcal{X})$,

$$\sum_{i\in N} w_i E_p[\phi(x)x_i] = \max_{\phi'\in\Phi}\sum_{i\in N} w_i E_p[\phi'(x)x_i], \tag{9}$$

or equivalently,

$$\sum_{i\in N} w_i p(\phi(x) = x_i) = \max_{\phi'\in\Phi}\sum_{i\in N} w_i p(\phi'(x) = x_i). \tag{10}$$

That is, a necessary and sufficient condition for a voting rule to be a WMR is that it maximizes the corresponding weighted sum of responsiveness over all voting rules for each $p \in \Delta(\mathcal{X})$. This result is summarized in the following proposition due to Fleurbaey (2008),[16] where the sufficient condition is weaker.[17]

Proposition 5. *If ϕ is a WMR with a weight vector w, then* (10) *holds for each $p \in \Delta(\mathcal{X})$. For fixed $p \in \Delta(\mathcal{X})^\circ \equiv \{p \in \Delta(\mathcal{X}) : p(x) > 0 \text{ for each } x \in \mathcal{X}\}$, where every x is possible, if* (10) *holds, then ϕ is a WMR with a weight vector w.*

We call the above result the Rae-Taylor-Fleurbaey (RTF) theorem because it generalizes the Rae-Taylor theorem[18] which focuses on a SMR. Note that a WMR in the RTF theorem can have negative weights. In Appendix B.2, we characterize a WMR with possibly negative weights by introducing a further weaker version of robustness.

The normative implication of the RTF theorem is efficiency of WMRs in the set of random voting rules under a fixed probability distribution.[19]

Definition 3. Fix $p \in \Delta(\mathcal{X})^\circ$. A deterministic voting rule $\phi \in \Phi$ is strictly efficient if any other random voting rule is not weakly Pareto-preferred to ϕ in terms of responsiveness under p. A

[16]See also Brighouse and Fleurbaey (2010), who discuss the implication of this result for democracy.

[17]To see why a weaker condition suffices, suppose that ϕ is not a WMR. Then, (8) does not hold for some $\phi' \in \Phi$ and $x \in \mathcal{X}$, which contradicts (9) and (10) for each $p \in \Delta(\mathcal{X})^\circ$.

[18]See Rae (1969), Taylor (1969), Straffin (1977), and references in Fleurbaey (2008).

[19]This issue is not formally discussed in Fleurbaey (2008). Instead, Fleurbaey (2008) considers the optimality of a WMR by assuming that w_i is proportional to i's utility, where the weighted sum of responsiveness is the total sum of expected utilities.

deterministic voting rule $\phi \in \Phi$ is efficient if any random voting rule is not Pareto-preferred to ϕ in terms of responsiveness under p. A deterministic voting rule $\phi \in \Phi$ is weakly efficient if any random voting rule is not strictly Pareto-preferred to ϕ in terms of responsiveness under p.

It should be noted that, even if no deterministic voting rule is Pareto-preferred to ϕ, there may exist a random voting rule that is Pareto-preferred to ϕ, which will be demonstrated later.

Using the RTF theorem, we can obtain the following normative characterizations of WMRs.

Proposition 6. *Fix $p \in \Delta(\mathcal{X})^\circ$. A deterministic voting rule is strictly efficient if and only if it is a WMR with positive weights such that there are no ties. A deterministic voting rule is efficient if and only if it is a WMR with positive weights. A deterministic voting rule is weakly efficient if and only if it is a WMR with nonnegative weights.*

Proof. To the best of the authors' knowledge, the proofs as well as the above formal statements have never appeared in the literature. We give the proofs in Appendix E. □

Propositions 1 and 6 characterize a WMR using responsiveness. In particular, the class of robust rules coincides with that of strictly efficient rules.[20] This implies that a robust rule is robustly undominated by any other random voting rule.

Let us elaborate on the difference between efficiency and robustness. Note that if a voting rule ϕ is a WMR with nonnegative weights, then, under every $p \in \Delta(\mathcal{X})$, no random voting rule is strictly Pareto-preferred to ϕ, as implied by (9). The converse is also true, but a weaker condition suffices, which is each of weak efficiency and weak robustness. Weak efficiency requires that, under a fixed $p \in \Delta(\mathcal{X})^\circ$, any random voting rule is not strictly Pareto-preferred to ϕ. Thus, we need to consider all random voting rules (i.e., all probability distributions of voting rules) to check weak efficiency. On the other hand, weak robustness requires that, under every $p \in \Delta(\mathcal{X})$, the inverse rule of ϕ is not strictly Pareto-preferred to ϕ. Thus, we need to consider all probability distributions of decision profiles to check weak robustness. The following table summarizes the difference between the two sufficient conditions by listing possible combinations of X, Y, and Z in the sentence: "a voting rule ϕ is X if Y is not strictly

[20]A robust rule is a WMR with nonnegative weights allowing no ties, which can be represented as a WMR with positive weights allowing no ties by Lemma 2.

Pareto-preferred to ϕ under Z."

X	Y	Z
weakly efficient	any $\bar{\phi} \in \Delta(\Phi)$	fixed $p \in \Delta(\mathcal{X})^\circ$
weakly robust	the inverse rule $-\phi \in \Phi$	any $p \in \Delta(\mathcal{X})$

We demonstrate the above difference using a numerical example.

Example 2. Let ϕ be the voting rule discussed in Example 1, where we demonstrate that ϕ is not weakly robust because the inverse rule is strictly Pareto-preferred to ϕ under some p. Here, we demonstrate that ϕ is not weakly efficient under some p because there exists a random voting rule that is strictly Pareto-preferred to ϕ. Fix $p \in \Delta(\mathcal{X})$ with

$$p(x) = \begin{cases} \alpha & \text{if } x \in \mathcal{A}, \\ \beta & \text{otherwise,} \end{cases} \tag{11}$$

where $0 < \alpha < 1/140$ and $\beta = (1 - 5\alpha)/27$. Then, no deterministic voting rule is Pareto-preferred to ϕ under p, as shown in Appendix D. However, ϕ is not weakly efficient because the following random voting rule $\bar{\phi}$ is strictly Pareto-preferred to ϕ under p:

$$\bar{\phi}(x) = \begin{cases} -1 + \delta & \text{if } x_1 = -1 \text{ and } x_i = 1 \text{ for all } i \neq 1, \\ -1 + \gamma & \text{if } x_1 = 1 \text{ and } \#\{i : x_i = 1\} = 2, \\ \phi(x) & \text{otherwise,} \end{cases}$$

where $\delta, \gamma \in (0, 1)$ and $\beta\delta/(4\alpha) < \gamma < \beta\delta/(2\alpha)$. In fact,

$$E_p[\phi'(x)x_1] - E_p[\phi(x)x_1] = -\beta\delta + 4\alpha\gamma > 0,$$
$$E_p[\phi'(x)x_i] - E_p[\phi(x)x_i] = \beta\delta - 2\alpha\gamma > 0 \text{ for each } i \neq 1.$$

Thus, ϕ does not satisfy the requirement of weak efficiency by (3).

Before closing this section, we emphasize that our characterization of WMRs is not a restatement of the RTF theorem. As discussed in Section 3, it is obvious that a WMR is robust. Thus, Proposition 6 implies that a strictly efficient rule is robust. However, neither Proposition 6 nor the RTF theorem says anything about whether a robust rule is strictly efficient. Our contribution is to identify the set of all robust rules, which is not implied by the RTF theorem.

6 Discussion

In this section, we discuss several extensions of Definition 1 and also compare our criterion with other criteria.

Random voting rules

We focus on deterministic voting rules in Definition 1, but we can also apply the same requirement to random voting rules. This requirement is very weak for random voting rules because robust random rules do not have the property in Lemma 4; that is, some deterministic voting rule can be Pareto-preferred to a robust random rule under some probability distribution of preferences. In Appendix A, we demonstrate it and characterize the class of robust random rules.

Heterogeneous priors

We adopt the notion of Pareto dominance under common beliefs in Definition 1. This is because Pareto dominance under heterogeneous beliefs is not so normatively compelling as discussed in the introduction. Moreover, when we adopt Pareto dominance under heterogeneous beliefs, the stronger version of robustness implies dictatorship. Let us say that a voting rule $\phi \in \Phi$ is robust under heterogeneous priors if, for any $p^1, \dots, p^n \in \Delta(\mathcal{X})$, it holds that $p^i(\phi(x) = x_i) > 1/2$ for at least one $i \in N$. We show that there exists a dictator, i.e., individual $i \in N$ with $\phi(x) = x_i$ for all $x \in \mathcal{X}$. Seeking a contradiction, suppose that no individual is a dictator. Then, for each $i \in N$, there exists $x^i \in \mathcal{X}$ and $p^i \in \Delta(\mathcal{X})$ such that $\phi(x^i) \neq x^i_i$ and $p^i(x^i) = 1$; that is, $p^i(\phi(x^i) = x^i_i) = 0 \leq 1/2$. Therefore, ϕ is not robust under heterogeneous priors. On the other hand, if individual i is a dictator, then $p(\phi(x) = x_i) = 1 > 1/2$ for all $p \in \Delta(\mathcal{X})$, so ϕ is robust under heterogeneous priors.

A proper subset of $\Delta(\mathcal{X})$

We take account of all beliefs in $\Delta(\mathcal{X})$ in Definition 1 because a social planner is assumed to be very conservative having no information about the true distribution. That is, our criterion is a comprehensively belief-free criterion. However, a proper subset $P \subsetneq \Delta(\mathcal{X})$ may be considered to be reasonable in some cases. Then, we can consider the following weaker requirement: we say that $\phi \in \Phi$ is P-robust if the inverse rule $-\phi$ is not weakly Pareto-preferred to ϕ under any $p \in P$. This criterion is similar to, but different from, the belief-neutral Pareto criterion introduced by Brunnermeier et al. (2014) in their study of financial markets. To translate it into our model, assume that individuals' subjective beliefs $p_1, \ldots, p_n \in \Delta(\mathcal{X})$ are observable to a social planner and the set of reasonable beliefs is the convex hull of them, i.e., $\mathbf{co}\{p_1, \ldots, p_n\}$. Under this assumption, ϕ is said to be belief-neutral efficient if it is efficient under any $p \in \mathbf{co}\{p_1, \ldots, p_n\}$; that is, any random voting rule is not Pareto-preferred to ϕ under any $p \in \mathbf{co}\{p_1, \ldots, p_n\}$. Note that, by the RTF theorem, ϕ is belief-neutral efficient if and only if it is a WMR with positive weights. Contrastingly, in P-robustness with $P = \mathbf{co}\{p_1, \ldots, p_n\}$, the inverse rule $-\phi$ is not weakly Pareto-preferred to ϕ under any $p \in \mathbf{co}\{p_1, \ldots, p_n\}$. In Online Appendix B, we demonstrate that a P-robust rule is not necessarily a WMR, and we also characterize the class of P-robust rules.

The maximin criterion

If a social planner has a cardinal utility function over the set of voting rules rather than the Pareto order, the planner can apply the maximin criterion to obtain the optimal rule. However, there may exist many sensible utility functions, and the optimal rule depends upon the choice of them; that is, it may be difficult to characterize the class of *all* WMRs using a single utility function. For example, let the weighted sum of each individual's responsiveness, $W(\phi, p) \equiv \sum_{i \in N} w_i p(\phi(x) = x_i)$, be the planner's utility function. We say that a voting rule $\phi \in \Phi$ satisfies the maximin criterion if $\phi \in \arg\max_{\phi' \in \Phi} \min_{p \in \Delta(\mathcal{X})} W(\phi', p)$. Then, a voting rule is a WMR with the same weights $w \equiv (w_i)_{i \in N}$, which is denoted by ϕ_w, if and only if it satisfies the maximin criterion. Let us prove it here. By Lemma 1 and (3), ϕ_w achieves the following maximum of $W(\phi, p)$ for each $p \in \Delta(\mathcal{X})$: $\max_{\phi \in \Phi} W(\phi, p) = \max_{\phi \in \Phi}(E_p[\phi(x) \sum_{i \in N} w_i x_i] + \sum_{i \in N} w_i)/2 = (E_p[|\sum_{i \in N} w_i x_i|] + \sum_{i \in N} w_i)/2$, which implies that ϕ_w satisfies the maximin criterion. Conversely, if ϕ satisfies the maximin criterion, then ϕ must be a WMR with the same weights w because $\min_{p \in \Delta(\mathcal{X})} W(\phi_w, p) = \min_{p \in \Delta(\mathcal{X})} W(\phi, p)$, and thus $\min_{p \in \Delta(\mathcal{X})} E_p[\phi_w(x) \sum_{i \in N} w_i x_i] = \min_{p \in \Delta(\mathcal{X})} E_p[\phi(x) \sum_{i \in N} w_i x_i] \geq 0$ by Lemma 1, which implies that ϕ is a WMR with the same weight w by Lemma 1 again.

7 Conclusion

The justification of WMRs and, in particular, SMRs based on efficiency arguments or axiomatic characterizations has yielded some of the celebrated contributions to the social choice and voting literature. The two paramount examples rationalizing a SMR within a dichotomous setting are Condorcet's jury theorem and May's theorem,[21] where the rationalization of a voting rule is based on asymptotic (i.e., infinite-individual) probabilistic criteria or deterministic criteria. An alternative approach based on non-asymptotic (i.e., finite-individual) probabilistic criteria was pioneered by Rae (1969), who suggested aggregate responsiveness as a meaningful criterion for evaluating the performance of a voting rule in the constitutional stage, namely, where the veil of ignorance prevails.

This paper contributes to the latter literature by introducing normative criteria for voting rules, robustness and weak robustness. Robustness requires that a voting rule should avoid the worst-case scenario in which the true responsiveness of every individual is less than or equal to one-half, and weak robustness requires that a voting rule should avoid the worst-case scenario in which the true responsiveness of every individual is less than one-half. The justification of the robustness property is based on two arguments: (i) it does not allow the collective decision to reflect just the minority preferences on average, and (ii) it does not permit Pareto inferiority in the set of all deterministic voting rules. Because a robust rule is never Pareto inferior under every common belief, it not only bypasses the problem associated with "spurious unanimity" under heterogeneous beliefs but also achieves "Pareto-ranking robustness" under common beliefs.

We establish that a voting rule is robust if and only if it is a WMR without ties and that a voting rule is weakly robust if and only if it is a WMR allowing ties with any tie-breaking rule. We also find that we face a trade-off between robustness and anonymity when the number of individuals is even: we must be content with a nonanonymous rule if we require robustness and we must be content with a weakly robust rule if we require anonymity. Our result and the RTF theorem (Rae, 1969; Taylor, 1969; Fleurbaey, 2008) have in common that both examine WMRs using responsiveness. However, the RTF theorem characterizes WMRs as efficient rules achieving the optimal outcomes, whereas our result characterizes WMRs as robust rules avoiding the worst outcomes. Hence, our result complements the renowned RTF theorem by providing a dual characterization of WMRs and, in particular, of SMRs in terms of responsiveness.

[21]See May (1952), Fishburn (1973), and Dasgupta and Maskin (2008).

A Robust random voting

In this appendix, we define and characterize a robust random voting rule. Robustness in this case is a very weak requirement. However, no anonymous random voting rule is robust when the number of individuals is even.

Recall that a random voting rule is a mapping $\bar{\phi} : \mathcal{X} \to [-1, 1]$ assigning a collective decision 1 with probability $(1 + \bar{\phi}(x))/2$ and -1 with probability $(1 - \bar{\phi}(x))/2$ to each $x \in \mathcal{X}$. Thus, the conditional responsiveness of individual $i \in N$ given $x \in \mathcal{X}$ is $(1 + \bar{\phi}(x))/2$ if $x_i = 1$ and $(1 - \bar{\phi}(x))/2$ if $x_i = -1$, which is rewritten as $(\bar{\phi}(x)x_i + 1)/2$. Therefore, when $x \in \mathcal{X}$ is drawn according to $p \in \Delta(\mathcal{X})$, the responsiveness of individual $i \in N$ is calculated as $(E_p[\bar{\phi}(x)x_i] + 1)/2$, which is analogous to (3).

We define a robust random rule by directly applying Definition 1 and provide a characterization, which is an extension of Proposition 1 to a random voting rule.

Definition A. A random voting rule $\bar{\phi} \in \Delta(\Phi)$ is *robust* if, for each $p \in \Delta(\mathcal{X})$, the responsiveness of at least one individual is strictly greater than one-half, or equivalently, $E_p[\bar{\phi}(x)x_i] > 0$ for at least one $i \in N$.

Proposition A. *A random voting rule $\bar{\phi} \in \Delta(\Phi)$ is robust if and only if there exists $w \in \mathbb{R}^n_+$ satisfying*

$$\bar{\phi}(x) \gtreqless 0 \iff \sum_{i \in N} w_i x_i \gtreqless 0. \tag{12}$$

Proof. The proof is essentially the same as that of Proposition 1. Consider an $n \times m$ matrix $\bar{L} = [\bar{l}_{ij}]_{n\times m} = \left[\bar{\phi}(x^j)x_i^j\right]_{n\times m}$. Then, (12) is rewritten as $\sum_{i\in N} w_i \bar{l}_{ij} = \sum_{i \in N} w_i\big(\bar{\phi}(x^j)x_i^j\big) > 0$ for each $j \in M$, or equivalently, $w^\top \bar{L} \gg 0$. By Lemma 6, this is true if and only if there does not exist $p = (p_j)_{j\in M} > 0$ such that $\bar{L}p \leq 0$, or equivalently, $\sum_{j\in M} \bar{l}_{ij} p_j = \sum_{j \in M} \bar{\phi}(x^j)x_i^j p_j \leq 0$ for each $i \in N$; that is, $\bar{\phi}$ is robust. □

Not only a WMR but also many other random voting rules satisfy (12). For example, assume that n is odd and let $\bar{\phi} \in \Delta(\Phi)$ be a random voting rule such that the collective decision is the majority's vote with probability $0.5 + \varepsilon$, where $0 < \varepsilon < 1/2$; that is,

$$\bar{\phi}(x) = \begin{cases} +2\varepsilon & \text{if } \sum_{i\in N} x_i > 0, \\ -2\varepsilon & \text{if } \sum_{i \in N} x_i < 0. \end{cases}$$

Note that $\bar{\phi}$ satisfies (12) with $w_i = 1$ for all $i \in N$, and thus it is robust. However, $\bar{\phi}$ can be dominated by a deterministic voting rule. To see this, let $\phi \in \Phi$ be a SMR, and let $p \in \Delta(\mathcal{X})$

be such that $p(x) = 1$ if $x_i = 1$ for all $i \in N$ and $p(x) = 0$ otherwise. Then, for each $i \in N$, $E_p[\bar{\phi}(x)x_i] - E_p[\phi(x)x_i] = 2\varepsilon - 1 < 0$, which implies that ϕ is strictly Pareto-preferred to $\bar{\phi}$ under p in terms of responsiveness.

Although robustness is a weak requirement for random voting rules, there exists no anonymous robust random rule when n is even. Let $p \in \Delta(\mathcal{X})$ be such that

$$p(x) = \begin{cases} 1/\binom{n}{n/2} & \text{if } \#\{i : x_i = 1\} = n/2, \\ 0 & \text{otherwise.} \end{cases}$$

Then, for any anonymous random voting rule $\bar{\phi} \in \Delta(\Phi)$, it holds that $E_p[\bar{\phi}(x)x_i] = E_p[\bar{\phi}(x)x_j]$ for all $i, j \in N$ and

$$\sum_{i \in N} E_p[\bar{\phi}(x)x_i] = E_p\left[\bar{\phi}(x)\sum_{i \in N} x_i\right] = 0.$$

This implies that $(E_p[\bar{\phi}(x)x_i] + 1)/2 = 1/2$ for all $i \in N$; that is, the responsiveness of every individual equals one-half. Therefore, $\bar{\phi}$ is not robust.

B Other robustness concepts

In this appendix, we consider two variants of robustness concepts and characterize them.

B.1 Strong robustness and a SMR

Even if a voting rule is robust and the responsiveness of at least one individual is strictly greater than one-half, the arithmetic mean of responsiveness of all individuals can be less than one-half, which implies that the collective decision reflects minority preferences on average. To avoid this scenario, a voting rule must satisfy the following stronger requirement.

Definition B. A voting rule $\phi \in \Phi$ is *strongly robust* if, for each $p \in \Delta(\mathcal{X})$, the arithmetic mean of responsiveness is strictly greater than one-half:

$$\sum_{i \in N} p(\phi(x) = x_i)/n > 1/2 \text{ for all } p \in \Delta(\mathcal{X}). \tag{13}$$

Clearly, a strongly robust rule is robust. In the next proposition, we show that a voting rule is strongly robust if and only if it is robust and anonymous; that is, it is a SMR with odd n.

Proposition B. *Suppose that n is odd. Then, a voting rule is strongly robust if and only if it is a SMR. Suppose that n is even. Then, no voting rule is strongly robust.*

Proof. Note that, by (3), (13) is equivalent to

$$\sum_{i\in N} E_p[\phi(x)x_i] > 0 \text{ for all } p \in \Delta(\mathcal{X}). \tag{14}$$

Suppose that n is odd and that ϕ is a SMR. Then, ϕ satisfies (14) by Lemma 1, so a SMR is strongly robust.

Suppose that ϕ is not a SMR. Then, there exist $y \in \mathcal{X}$ and $S \subsetneq N$ such that $|S| < n/2$ and $\phi(y) = y_i$ if and only if $i \in S$. Let $p \in \Delta(\mathcal{X})$ be such that $p(y) = 1$. Then, $\sum_{i\in N} E_p[\phi(x)x_i] = |S| - |N \setminus S| < 0$, violating (14). Thus, a voting rule is not strongly robust unless it is a SMR.

Suppose that n is even. Let $y \in \mathcal{X}$ be such that $y_i = 1$ for $i \leq n/2$ and $y_i = -1$ for $i \geq n/2+1$. For $p \in \Delta(\mathcal{X})$ with $p(y) = 1$ and any $\phi \in \Phi$, it holds that $\sum_{i\in N} E_p[\phi(x)x_i] = 0$, violating (14). Thus, no voting rule is strongly robust when n is even. □

B.2 Semi-robustness and a WMR

We consider the following weaker version of robustness to characterize a WMR with possibly negative weights.

Definition C. A voting rule $\phi \in \Phi$ is *semi-robust* if, for each $p \in \Delta(\mathcal{X})$, the responsiveness of at least one individual is not equal to one-half.

To understand the implication of semi-robustness, imagine that some individuals are more likely to have correct choices and other individuals are more likely to have wrong choices. However, if the responsiveness of every individual is equal to one-half, then it is difficult to extract information from individuals in order to arrive at a correct group decision. A semi-robust rule does not face this problem for any probability distribution.

The following proposition establishes the equivalence of a semi-robust rule and a WMR with possibly negative weights allowing no ties.

Proposition C. *A voting rule is semi-robust if and only if it is a WMR such that there are no ties.*

To prove Proposition C, we use the following theorem of alternatives, which is referred to as Gordan's theorem. This result also appears in the work of Gale (1960, Theorems 2.9) as a corollary of Farkas' lemma.

Lemma A. *Let A be an $n \times m$ matrix. Exactly one of the following alternatives holds.*

- *There exists $\xi \in \mathbb{R}^n$ satisfying*

$$\xi^\top A \gg 0.$$

- *There exists $\eta \in \mathbb{R}^m$ satisfying*

$$A\eta = 0, \ \eta > 0.$$

Proof of Proposition C. We can restate the conditions in Proposition C as follows.

(a″) By Lemma 1, a voting rule ϕ is a WMR allowing no ties if and only if there exists $w = (w_i)_{i \in N} \neq 0$ such that

$$\sum_{i \in N} w_i l_{ij} = \sum_{i \in N} w_i \big(\phi(x^j) x_i^j\big) > 0$$

for each $j \in M$, or equivalently, $w^\top L \gg 0$.

(b″) By definition, a voting rule is *not* semi-robust if and only if there exists $p = (p_j)_{j \in M} > 0$ such that

$$\sum_{j \in N} l_{ij} p_j = \sum_{j : \phi(x^j) = x_i^j} p_j - \sum_{j : \phi(x^j) \neq x_i^j} p_j = 0$$

for each $i \in N$, or equivalently, $Lp = 0$.

Proposition C states that exactly one of (a″) and (b″) holds, which is true by Lemma A. In fact, by plugging L, w, and p into A, ξ, and η in Lemma A, respectively, we can conclude that exactly one of (a″) and (b″) holds. □

C An imaginary asset market

We can explain why Proposition 1 is true in terms of arbitrage-free pricing in an imaginary asset market because we can interpret Lemma 6 as a corollary of the fundamental theorem of asset pricing.

Let M and N be the set of states and the set of assets, respectively. One unit of asset $i \in N$ yields a payoff l_{ij} when state $j \in M$ is realized. Recall that l_{ij} equals +1 if i's choice agrees with the collective decision and −1 otherwise. The matrix L is referred to as the payoff matrix. We denote by $q = (q_i)_{i \in N}$ the vector of prices of the n assets.

A portfolio defined by a vector $w = (w_i)_{i \in N}$ consists of w_i units of asset i for each $i \in N$. It yields a payoff $\sum_{i \in N} w_i l_{ij}$ when state $j \in M$ is realized, which is summarized in $w^\top L = (\sum_{i \in N} w_i l_{ij})_{j \in M}$. The price of the portfolio is $w^\top q = \sum_{i \in N} q_i w_i$.

We say that a price vector q is arbitrage-free if $w^\top L \geq 0$ implies $w^\top q \geq 0$; that is, the price of any portfolio yielding a nonnegative payoff in each state is nonnegative. We say that a price vector q is determined by a nonnegative linear pricing rule if there exists a nonnegative vector $p = (p_j)_{j \in M} > 0$, which is referred to as a state price, such that $q = Lp$. The fundamental theorem of asset pricing establishes the equivalence of an arbitrage-free price and the existence of a nonnegative linear pricing rule, which is immediate from Farkas' lemma.

Claim A. *A price vector q is arbitrage-free if and only if it is determined by a nonnegative linear pricing rule. That is, the set of all arbitrage-free price vectors is $\{q : q = Lp,\ p > 0\}$.*

The fundamental theorem of asset pricing has the following corollary, which is immediate from Lemma 6 (a corollary of Farkas' lemma).

Claim B. *There exists a portfolio with nonnegative weights in all assets (i.e. no short selling) yielding a strictly positive payoff in each state if and only if, for any arbitrage-free price vector, the price of at least one asset is strictly positive.*

The former condition is restated as $w^\top L \gg 0$ for some $w \geq 0$ and the latter condition is restated as $Lp \nleq 0$ for all $p > 0$. Therefore, Claim B implies the equivalence of a robust rule and a WMR with nonnegative weights allowing no ties.

This paper does not discuss how to find optimal weights of WMRs, whereas it is a central topic in modern portfolio theory to determine optimal weights. Thus, modern portfolio theory could be useful to find optimal weights of WMRs.

D On Example 2

In this appendix, we consider $\phi \in \Phi$ and $p \in \Delta(\mathcal{X})$ given by (7) and (11), respectively, and show that no other deterministic voting rule is Pareto-preferred to ϕ under p.

Let $\phi' \in \Phi$ be weakly Pareto-preferred to ϕ under p. It is enough to show that $\phi' = \phi$. We first prove that $\phi'(x) = \phi(x)$ for all $x \notin \mathcal{A}$. This is because otherwise there exists $x \notin \mathcal{A}$ such that $\phi'(x)x_1 = -1$ and thus the expected payoff of individual 1 under ϕ' is strictly less that that under ϕ:

$$\begin{aligned} E_p[\phi(x)x_1] - E_p[\phi'(x)x_1] &= \sum_{x\notin\mathcal{A}} p(x)(\phi(x)x_1 - \phi'(x)x_1) + \sum_{x\in\mathcal{A}} p(x)(\phi(x)x_1 - \phi'(x)x_1) \\ &= \sum_{x\notin\mathcal{A}} \beta(1-\phi'(x)x_1) + \sum_{x\in\mathcal{A}} \alpha(-1-\phi'(x)x_1) \\ &\geq 2\beta - 10\alpha = 2(1-140\alpha)/27 > 0. \end{aligned}$$

We next prove that $\phi'(x) = \phi(x)$ for all $x \in \mathcal{A}$. By the above property of ϕ', we obtain

$$\sum_{i\in N} E_p[\phi(x)x_i] - \sum_{i\in N} E_p[\phi'(x)x_i] = \alpha \sum_{x\in\mathcal{A}} \left(\sum_{i\in N} \phi(x)x_i - \sum_{i\in N} \phi'(x)x_i \right) \geq 0,$$

where the last inequality follows from $\sum_{i\in N}\phi(x)x_i = |\sum_{i\in N} x_i| = |\sum_{i\in N}\phi'(x)x_i| \geq \sum_{i\in N}\phi'(x)x_i$ for all $x \in \mathcal{A}$. We also have $\sum_{i\in N} E_p[\phi(x)x_i] \leq \sum_{i\in N} E_p[\phi'(x)x_i]$ because ϕ' is weakly Pareto-preferred to ϕ under p. Thus, $\sum_{i\in N} E_p[\phi(x)x_i] = \sum_{i\in N} E_p[\phi'(x)x_i]$, which implies that $\sum_{i\in N}\phi(x)x_i = \sum_{i\in N}\phi'(x)x_i$ for all $x \in \mathcal{A}$, i.e., $\phi'(x) = \phi(x)$ for all $x \in \mathcal{A}$.

E Proof of Proposition 6

Proof. We first prove the first statement, which asserts that exactly one of the following holds.

(a‴) Note that, by Lemma 1, a voting rule ϕ is a WMR with nonnegative weights allowing no ties if and only if there exists $w = (w_i)_{i\in N} \geq 0$ such that, for all $\phi' \in \Phi \setminus \{\phi\}$ and $x \in \mathcal{X}$, (8) is true with strict inequality holding for at least one decision profile x. This is true if and only if there exists $w = (w_i)_{i\in N} \geq 0$ such that

$$\sum_{i\in N} w_i E_p[\phi(x)x_i] > \sum_{i\in N} w_i E_p[\phi'(x)x_i] \text{ for all } \phi' \in \Phi \setminus \{\phi\},$$

or equivalently,

$$\sum_{i\in N} w_i(E_p[\phi(x)x_i] - E_p[\phi'(x)x_i]) > 0 \text{ for all } \phi' \in \Phi \setminus \{\phi\}.$$

(b‴) There exists a random voting rule that is weakly Pareto-preferred to ϕ if and only if there exists $\rho \in \mathbb{R}_+^{\Phi\setminus\{\phi\}}$ such that

$$E_p[\phi(x)x_i] \leq \sum_{\phi'\in\Phi\setminus\{\phi\}} E_p[\phi'(x)x_i]\rho(\phi') / \sum_{\phi'\in\Phi\setminus\{\phi\}} \rho(\phi') \text{ for all } i \in N,$$

or equivalently,

$$\sum_{\phi'\in\Phi\setminus\{\phi\}} (E_p[\phi(x)x_i] - E_p[\phi'(x)x_i])\rho(\phi') \leq 0 \text{ for all } i \in N.$$

By Lemma 6, exactly one of (a‴) and (b‴) holds. Although the same theorem of alternatives characterizes both robust rules and strictly efficient rules, its use is different from each other. In fact, as discussed in Section 5, robustness and strict efficiency are different conditions.

The second and third statements are implied by the well-known theorem of Wald (1950) on admissible decision functions or that of Pearce (1984) on undominated strategies. Proposition 5 states that a voting rule is a WMR with a weight vector w if and only if (9) holds. Mathematically, (9) is equivalent to expected utility maximization, where Φ is the set of actions, N is the set of states, and $w_i/\sum_j w_j$ is a probability of state $i \in N$. Therefore, we can apply the theorem of Wald (1950) on admissible decision functions or that of Pearce (1984) on undominated strategies. In particular, Theorems 5.2.1 and 5.2.5 in Blackwell and Girshick (1954) are useful. Theorem 5.2.1 implies that a voting rule is weakly efficient if and only if there exists a weight vector $w > 0$ such that (9) holds. Theorem 5.2.5 implies that a voting rule is efficient if and only if there exists a weight vector $w \gg 0$ such that (9) holds. Therefore, this corollary holds by Proposition 5. □

References

Auster, S., 2018. Robust contracting under common value uncertainty. Theor. Econ. 13, 175–204.

Azrieli, Y., Kim, S., 2014. Pareto efficiency and weighted majority rules. Int. Econ. Rev. 55, 1067–1088.

Bergemann, D., Brooks, B., Morris, S., 2016. Informationally robust optimal auction design. Working paper.

Bergemann, D., Morris, S., 2005. Robust mechanism design. Econometrica 73, 1771–1813.

Bergemann, D., Schlag, K., 2011. Robust monopoly pricing. J. Econ. Theory 146, 2527–2543.

Blackwell, D., Girshick, M. A., 1954. Theory of Games and Statistical Decisions. John Wiley & Sons.

Blume, L. E., Cogley, T., Easley, D. A., Sargent, T. J., Tsyrennikov, V., 2018. A case for incomplete markets. J. Econ. Theory 178, 191–221.

Brighouse, H., Fleurbaey, M., 2010. Democracy and proportionality. J. Polit. Philos. 18,137–155.

Brunnermeier, M., Simsek, A., Xiong, W., 2014. A welfare criterion for models with distorted beliefs. Q. J. Econ. 129, 1753–1797.

Carrasco, V., Luz, V. F., Kos, N., Messner, M., Monteiro, P., Moreira, H., 2018. Optimal selling mechanisms under moment conditions. J. Econ. Theory 177, 245–279.

Carroll, G., 2015. Robustness and linear contracts. Am. Econ. Rev. 105, 536–563.

Carroll, G., 2017. Robustness and separation in multidimensional screening. Econometrica 85, 453–488.

Carroll, G., 2019. Robustness in Mechanism Design and Contracting. Annu. Rev. Econ. 11, 139–166.

Chen, Y-C, Li, J., 2018. Revisiting the foundations of dominant-strategy mechanisms. J. Econ. Theory 178, 294–317.

Chung, K. S., Ely, J. C., 2007. Foundations of dominant-strategy mechanisms. Rev. Econ. Stud. 74, 447–476.

Dasgupta, P., Maskin, E., 2008. On the robustness of majority rule. J. Eur. Econ. Assoc. 6, 949–973.

Du, S., 2018. Robust mechanisms under common valuation. Econometrica 86, 1569–1588.

Dybvig, P., Ross, S., 2003. Arbitrage, state prices and portfolio theory. In: Constantinides, G., Harris, M., Stulz, R. M. (Eds.), Handbook of the Economics of Finance 1. Elsevier, pp. 605–637.

Dybvig, P., Ross, S., 2008. Arbitrage. In: Durlauf, S., Blume, L. (Eds.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, pp. 188–197.

Einy, E., Lehrer, E., 1989. Regular simple games. Int. J. Game Theory 18, 195–207.

Fishburn, P. C., 1973. The Theory of Social Choice. Princeton University Press.

Fishburn, P. C., Gehrlein, W. V., 1977. Collective rationality versus distribution of power for binary social choice functions. J. Econ. Theory 15, 72–91.

Fleurbaey, M., 2008. One stake one vote. Working paper.

Gale, D., 1960. The Theory of Linear Economic Models. University of Chicago Press.

Gayer, G., Gilboa, I., Samuelson, L., Schmeidler, D., 2014. Pareto efficiency with different beliefs. J. Legal Stud. 43, 151–171.

Gilboa, I., Samet, D., Schmeidler, D., 2004. Utilitarian aggregation of beliefs and tastes. J. Polit. Econ. 112, 932–938.

Gilboa, I., Samuelson, L., 2021. What were you thinking? Decision theory as coherence test. Forthcoming in Theor. Econ.

Gilboa, I., Samuelson, L. Schmeidler, D., 2014. No-betting-Pareto dominance. Econometrica 82, 1405–1442.

Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with a non-unique prior. J. Math. Econ. 18, 141–153.

Huber, P. J., 1981. Robust Statistics. John Wiley & Sons.

May, K., 1952. A set of independent necessary and sufficient conditions for simple majority decisions. Econometrica 20, 680–684.

Mongin, P., 2016. Spurious unanimity and the Pareto principle. Econ.Philos. 32, .511–532.

Neumann, J. V., Morgenstern, O., 1944. Theory of Games and Economic Behavior. Princeton University Press.

Nitzan, S., Paroush, J., 1981. The characterization of decisive weighted majority rules. Econ. Letters 7. 119–124.

Pearce, D. G., 1984. Rationalizable strategic behavior and the problem of perfection. Econometrica 52, 1029–1050.

Rae, D. W., 1969. Decision-rules and individual values in constitutional choice. Am. Pol. Sci. Rev. 63, 40–56.

Schmitz, P. W., Tröger, T., 2012. The (sub-)optimality of the majority rule. Games Econ. Behav. 74, 651–665.

Straffin, P. D., 1977. Majority rule and general decision rules. Theory Dec. 8, 351–360.

Taylor, A., Zwicker, W., 1992. A characterization of weighted voting. Proc. Am. Math. Soc. 115, 1089–1094.

Taylor, M., 1969. Proof of a theorem on majority rule. Behav. Sci. 14, 228–231.

Yamashita, T., 2015. Implementation in weakly undominated strategies: optimality of second-price auction and posted-price mechanism. Rev. Econ. Stud. 82, 1223–1246.

Yamashita, T., Zhu, S., 2021. On the foundations of ex post incentive compatible mechanisms. Forthcoming in Am. Econ. J. Microecon.

Wald, A., 1950. Statistical Decision Functions. Chapman & Hall.

Online Appendix to
"Robust Voting under Uncertainty"

Satoshi Nakada* Shmuel Nitzan† Takashi Ui‡

December 2021

In the main text, we introduce the concept of robustness by assuming that an individual votes sincerely and prefers a voting rule with higher responsiveness. We provide justification for these assumptions by discussing two alternative definitions of robustness in Online Appendix A. First, we redefine robustness as a requirement for a voting mechanism, where individuals can vote insincerely, and show that sincere voting is weakly dominant for every individual if the mechanism satisfies robustness. Next, we redefine robustness using the expected utility of individuals (rather than responsiveness) and provide justification for the original definition, where we consider a set of probability distributions over individuals' utility functions. We elaborate on the robustness in terms of the expected utility in Online Appendix B.

A Alternative definitions of robustness

A.1 Robustness as a requirement for a mechanism

We regard $\phi \in \Phi$ as a mechanism with a message space $\mathcal{X}$. Individual $i \in N$ has a VNM utility function $u_i : \{-1, 1\} \to \mathbb{R}$ over the set of alternatives $\{-1, 1\}$ representing a strict preference relation, i.e., $u_i(1) \neq u_i(-1)$. The preferred alternative of individual i is denoted by $\delta u_i \in \{-1, 1\}$, i.e.,

$$\delta u_i \equiv \begin{cases} 1 & \text{if } u_i(1) > u_i(-1), \\ -1 & \text{if } u_i(-1) > u_i(1). \end{cases}$$

*School of Management, Department of Business Economics, Tokyo University of Science.
†Department of Economics, Bar-Ilan University.
‡Department of Economics, Hitotsubashi University.

Let $\mathcal{U}_i \equiv \{u_i : u_i(1) \neq u_i(-1)\}$ be the set of all utility functions of individual i representing strict preferences. We write $\mathcal{U} \equiv \prod_{i \in N} \mathcal{U}_i$.

Individual $i \in N$ with a voting strategy $\sigma_i : \mathcal{U}_i \to \{-1, 1\}$ casts a vote for $\sigma_i(u_i) \in \{-1, 1\}$. A sincere voting strategy is $\delta_i : \mathcal{U}_i \to \{-1, 1\}$ satisfying $\delta_i(u_i) = \delta u_i$ for all $u_i \in \mathcal{U}_i$. Let Σ_i be the set of all strategies of individual i. We write $\Sigma = \prod_{i \in N} \Sigma_i$.

Assume that $u \equiv (u_i)_{i \in N} \in \mathcal{U}$ is randomly drawn according to a probability distribution $\lambda \in \Delta(\mathcal{U})$, where $\Delta(\mathcal{U})$ is the set of all probability distributions over $\mathcal{U}$. When the individuals follow a strategy profile $\sigma \equiv (\sigma_i)_{i \in N} \in \Sigma$ under a mechanism ϕ, the probability distribution of a decision profile $\sigma(u) \equiv (\sigma_i(u_i))_{i \in N}$ is $\lambda \circ \sigma^{-1} \in \Delta(\mathcal{X})$ with

$$\lambda \circ \sigma^{-1}(x) \equiv \lambda(\{u \in \mathcal{U} : \sigma(u) = x\}),$$

and the responsiveness of individual $i \in N$ is

$$\lambda \circ \sigma^{-1}(\phi(x) = x_i) \equiv \lambda \circ \sigma^{-1}(\{x \in \mathcal{X} : \phi(x) = x_i\}) = \lambda(\{u \in \mathcal{U} : \phi(\sigma(u)) = \sigma_i(u_i)\}). \quad \text{(A.1)}$$

We define robustness of a voting mechanism using the responsiveness (A.1) generated by the reported messages $\sigma(u)$, where the true probability distribution $\lambda \in \Delta(\mathcal{U})$ and the strategy profile $\sigma \in \Sigma$ followed by the individuals are assumed to be unknown.

Definition A. A voting rule $\phi \in \Phi$ is *robust as a mechanism* if, for each $\lambda \in \Delta(\mathcal{U})$ and each strategy profile $\sigma \in \Sigma$, $\lambda \circ \sigma^{-1}(\phi(x) = x_i) > 1/2$ for at least one $i \in N$.

If $\phi \in \Phi$ is robust as a mechanism, then it is robust in the original sense. This is because ϕ is robust if and only if, for each $\lambda \in \Delta(\mathcal{U})$, the responsiveness (A.1) is strictly greater than one-half for at least one individual when every player adopts a sincere voting strategy. However, it is straightforward to see that both concepts are equivalent.

Lemma A. *A voting rule $\phi \in \Phi$ is robust as a mechanism if and only if it is robust.*

Proof. This lemma follows from $\Delta(\mathcal{X}) = \bigcup_{\sigma \in \Sigma} \{\lambda \circ \sigma^{-1} \in \Delta(\mathcal{X}) : \lambda \in \Delta(\mathcal{U})\}$. □

Moreover, if ϕ is robust as a mechanism, then a sincere voting strategy is weakly dominant under ϕ; that is, ϕ satisfies strategyproofness.

Lemma B. *If $\phi \in \Phi$ is robust as a mechanism, then, for each individual $i \in N$, the sincere voting strategy δ_i is a best response to any strategy profile of the other individuals.*

Proof. It is enough to show that, for any decision profile of the other individuals $x_{-i} \equiv (x_j)_{j \neq i}$, it holds that $u_i(\phi(\delta u_i, x_{-i})) \geq u_i(\phi(-\delta u_i, x_{-i}))$ for all $i \in N$ and $u_i \in \mathcal{U}_i$. Seeking a contradiction, suppose otherwise. That is, there exist $i^* \in N$, $u_{i^*} \in \mathcal{U}_{i^*}$, and $z_{-i^*} \equiv (z_j)_{j \neq i^*}$ such that $\phi(\delta u_{i^*}, z_{-i^*}) = -\delta u_{i^*}$ and $\phi(-\delta u_{i^*}, z_{-i^*}) = \delta u_{i^*}$. Without loss of generality, we assume that $\delta u_{i^*} = 1$. Let $p \in \Delta(\mathcal{X})$ be such that $p(1, z_{-i^*}) = p(-1, z_{-i^*}) = 1/2$. Then, $p(\phi(x) = x_{i^*}) = 0$ and $p(\phi(x) = x_i) = 1/2$ if $i \neq i^*$. This implies that ϕ is not robust as a mechanism by Lemma A. □

In summary, we can give an equivalent definition of robustness without the assumption of sincere voting and show that a sincere voting strategy is weakly dominant under a robust rule, thus providing justification for the assumption of sincere voting in the main text.

A.2 Robustness in terms of expected utility

Under the assumption of sincere voting justified in Section A.1, we redefine robustness using the expected utility of individuals and provide justification for the use of responsiveness in the original definition of robustness.

Individual $i \in N$ with a utility function $u_i \in \mathcal{U}_i$ votes for the preferred alternative $\delta u_i \in \{-1, 1\}$, so the resulting decision profile is $\delta u = (\delta u_i)_{i \in N} \in \mathcal{X}$. When $\lambda \in \Delta(\mathcal{U})$ is the true probability distribution, individual $i \in N$ is assumed to prefer a voting rule ϕ to another voting rule ϕ' if and only if $E_\lambda[u_i(\phi(\delta u))] \geq E_\lambda[u_i(\phi'(\delta u))]$, where E_λ is the expectation operator with respect to λ. Thus, we consider the following Pareto order in terms of expected utility under λ: we say that ϕ is weakly Pareto-preferred to ϕ' (in terms of expected utility) under λ if $E_\lambda[u_i(\phi(\delta u))] \geq E_\lambda[u_i(\phi'(\delta u))]$ for all $i \in N$. A special case is the Pareto order in terms of responsiveness because the responsiveness equals the expected utility if the utility level is either zero or one. To explain it more formally, denote the set of such utility functions and that of the corresponding profiles of utility functions by

$$\mathcal{U}_i^* = \{u_i \in \mathcal{U}_i : u_i(1), u_i(-1) \in \{0, 1\}\} \text{ and } \mathcal{U}^* \equiv \prod_{i \in N} \mathcal{U}_i^*,$$

respectively. Then, for $p \in \Delta(\mathcal{X})$ and $\lambda_p \in \Delta(\mathcal{U})$ such that $\lambda_p(u) = p(\delta u)$ if $u \in \mathcal{U}^*$ and $\lambda_p(u) = 0$ otherwise, ϕ is Pareto-preferred to ϕ' in terms of expected utility under λ_p if and only if ϕ is Pareto-preferred to ϕ' in terms of responsiveness under p because $E_{\lambda_p}[u_i(\phi(\delta u))] = p(\phi(x) = x_i)$.

We consider an extension of robustness using the Pareto order in terms of expected utility, where the true probability distribution $\lambda \in \Delta(\mathcal{U})$ is assumed to be unknown but known to be contained in $\Lambda \subseteq \Delta(\mathcal{U})$.

Definition B. For $\Lambda \subseteq \Delta(\mathcal{U})$, a voting rule $\phi \in \Phi$ is Λ-*robust* if the inverse rule of ϕ is not weakly Pareto-preferred to ϕ in terms of expected utility under any $\lambda \in \Lambda$.

By the above discussion, $\phi \in \Phi$ is robust if and only if it is Λ^*-robust with

$$\Lambda^* = \{\lambda \in \Delta(\mathcal{U}) : \lambda(u) > 0 \text{ implies } u \in \mathcal{U}^*\} = \{\lambda_p \in \Delta(\mathcal{U}) : p \in \Delta(\mathcal{X})\}.$$

Thus, Λ-robustness with $\Lambda \supset \Lambda^*$ is a stronger requirement than robustness, and Λ-robustness with $\Lambda \subset \Lambda^*$ is a weaker requirement than robustness.

The underlying assumption in Λ^*-robustness is that the utility levels are known to be either zero or one. However, even if $\Lambda \supsetneq \Lambda^*$ and the utility levels are unknown under Λ, Λ-robustness can be equivalent to robustness.

To demonstrate it, consider

$$\Gamma^* = \{\lambda \in \Delta(\mathcal{U}) : \frac{E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i) | \delta u = x]}{E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i) | \delta u = x']} = 1 \text{ for all } x, x' \in \mathcal{X} \text{ and } i \in N\}$$

and assume that the true probability distribution is unknown but known to be contained in Γ^*. Note that $E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i) | \delta u = x]$ is the conditional expected net gain of individual $i \in N$ induced by the switch from an unfavorable alternative $-\delta u_i$ to a favorable alternative δu_i given $\delta u = x$. Thus, Γ^* is the set of all probability distributions over $\mathcal{U}$ such that the conditional expected net gain is the same for all $x \in \mathcal{X}$, and in particular, it holds that $\Lambda^* \subset \Gamma^*$. Note that the conditional expected net gain is unknown because it depends upon the unknown probability distribution λ, which also implies that the utility levels are unknown. However, Γ^*-robustness is equivalent to robustness.

Lemma C. *A voting rule $\phi \in \Phi$ is Γ^*-robust if and only if it is robust.*

Proof. For each $\lambda \in \Delta(\mathcal{U})$, let $p_\lambda \in \Delta(\mathcal{X})$ be the probability distribution of the decision profile δu; that is, $p_\lambda(x) = \lambda(\{u \in \mathcal{U} : \delta u = x\})$ for all $x \in \mathcal{X}$. Then, the expected utility of individual $i \in N$ under $\phi \in \Phi$ is

$$\begin{aligned}
E_\lambda[u_i(\phi(\delta u))] &= \sum_{x:\phi(x)=x_i} E_\lambda[u_i(\delta u_i) | \delta u = x] p_\lambda(x) + \sum_{x:\phi(x)=-x_i} E_\lambda[u_i(-\delta u_i) | \delta u = x] p_\lambda(x) \\
&= \sum_{x:\phi(x)=x_i} E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i) | \delta u = x] p_\lambda(x) + \sum_{x \in \mathcal{X}} E_\lambda[u_i(-\delta u_i) | \delta u = x] p_\lambda(x) \\
&= c_i p_\lambda(\phi(x) = x_i) + E_\lambda[u_i(-\delta u_i)],
\end{aligned}$$

where $c_i = E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i)|\delta u = x] > 0$, which is independent of $x \in \mathcal{X}$. Thus, $E_\lambda[u_i(\phi(\delta u))] \geq E_\lambda[u_i(\phi'(\delta u))]$ if and only if $p_\lambda(\phi(x) = x_i) \geq p_\lambda(\phi'(x) = x_i)$. This implies that Γ^*-robustness is equivalent to robustness. □

B Λ-robustness

We study Λ-robustness introduced in Section A.2 when $\Lambda \neq \Lambda^*$ and $\Lambda \neq \Gamma^*$. First, we consider Λ-robustness with $\Lambda \subset \Lambda^*$ and show that it can be a substantially weaker requirement than robustness in that some deterministic rule can be Pareto-preferred to a Λ-robust rule. Next, we consider Λ-robustness with $\Lambda \supset \Gamma^*$ and demonstrate the following. Even if $\Lambda \supsetneq \Gamma^*$, Λ-robustness can be equivalent to robustness. However, if Λ is too large, Λ-robustness implies the existence of a dictator.

B.1 Λ-robustness with $\Lambda \subset \Lambda^*$

When the true distribution of decision profiles is known to be in $P \subsetneq \Delta(\mathcal{X})$, we can consider the following weaker requirement than robustness (as discussed in Section 6): for each $p \in P$, the inverse rule $-\phi$ is not weakly Pareto-preferred to ϕ. This requirement is equivalent to Λ_P-robustness with $\Lambda_P \equiv \{\lambda_p \in \Lambda^* : p \in P\} \subsetneq \Lambda^*$ discussed in Section A.2, and the mapping $P \mapsto \Lambda_P$ is a bijection from $2^{\Delta(\mathcal{X})}$ to 2^{Λ^*}. We characterize a Λ_P-robust rule with $P \subset \Delta(\mathcal{X})$, or equivalently, a Λ-robust rule with $\Lambda \subset \Lambda^*$.

Because Λ^*-robustness is equivalent to robustness, no other deterministic voting rule is weakly Pareto-preferred to a Λ^*-robust rule under any $p \in \Delta(\mathcal{X})$ by Lemma 4. However, when $\Lambda \subsetneq \Lambda^*$, some deterministic voting rule can be weakly Pareto-preferred to a Λ-robust rule under some $p \in \Delta(\mathcal{X})$. To see this, assume that $n = 9$ and let $\phi \in \Phi$ be a two-thirds rule, i.e., $\phi(x) = 1$ if and only if $\#\{i : x_i = 1\} \geq 6$. Define

$$P = \{p \in \Delta(\mathcal{X}) : p(x) = p(x^\pi) \text{ for all } \pi \in \Pi,\ p(\phi(x) = x_i) > 1/2 \text{ for all } i \in N\},$$

which is the set of symmetric probability distributions such that ϕ is Λ_P-robust. Then, a SMR is strictly Pareto-preferred to ϕ under $p \in P$ with $p(\{x : \#\{i : x_i = 1\} = 5\}) > 0$ because the responsiveness of every individual under a SMR is $\sum_{k=5}^{9} k/9 \times p(\{x : \#\{i : x_i = 1\} = k\})$, and that under a two-thirds rule is $\sum_{k=6}^{9} k/9 \times p(\{x : \#\{i : x_i = 1\} = k\})$.[1]

[1]For example, such p is given by $p(x) = 1 - \varepsilon$ if $\#\{i : x_i = 1\} = 9$ and $p(x) = \varepsilon/\binom{9}{5}$ if $\#\{i : x_i = 1\} = 5$, where

The above discussion suggests that Λ_P-robustness can be a substantially weaker requirement than robustness when $P \subsetneq \Delta(\mathcal{X})$. In particular, a Λ_P-robust rule is not necessarily a WMR. However, a Λ_P-robust rule has a characterization similar to the characterization of a WMR in Lemma 1 if P is a convex hull of a finite set of $\Delta(\mathcal{X})$. The following proposition is an extension of Proposition 1 in the main text. Its proof is essentially the same as that of Proposition 1.

Proposition A. *Suppose that $P \subseteq \Delta(\mathcal{X})$ is a convex hull of a finite set $\{p_j\}_{j\in M} \subsetneq \Delta(\mathcal{X})$. Then, a deterministic voting rule $\phi \in \Phi$ is Λ_P-robust if and only if there exists $w \in \mathbb{R}^n_+$ such that*

$$\sum_{i\in N} w_i E_{p_j}[\phi(x)x_i] > 0 \text{ for all } j \in M. \tag{B.2}$$

Proof. Let $M = \{1, \dots, m\}$ and $L = [l_{ij}]_{n\times m} = \left[E_{p_j}[\phi(x)x_i]\right]_{n\times m}$. Then, (B.2) is rewritten as $\sum_{i\in N} w_i l_{ij} > 0$ for each $j \in M$, or equivalently, $w^\top L \gg 0$. By Lemma 6, this is true if and only if there does not exist $\lambda = (\lambda_j)_{j\in M} > 0$ such that $L\lambda \leq 0$, or equivalently, $\sum_{j\in M} l_{ij}\lambda_j = \sum_{j\in M} \lambda_j E_{p_j}[\phi(x)x_i] \leq 0$ for each $i \in N$. This is true if and only if ϕ is Λ_P-robust because P is a convex hull of $\{p_j\}_{j\in M}$. □

Note that $\Delta(\mathcal{X})$ is the convex hull of the finite set consisting of all degenerate probability distributions assigning probability one to some decision profile $x \in \mathcal{X}$. Thus, if $P = \Delta(\mathcal{X})$, then (B.2) is reduced to (2) in Lemma 1; that is, a Λ_P-robust rule is a WMR. If $P \subsetneq \Delta(\mathcal{X})$, however, (B.2) is a weaker requirement than (2), so a Λ_P-robust rule is not necessarily a WMR.

A similar result holds as long as P is a closed convex set. That is, we can show that a voting rule $\phi \in \Phi$ is Λ_P-robust if and only if there exists $w \in \mathbb{R}^n_+$ such that

$$\sum_{i\in N} w_i E_p[\phi(x)x_i] > 0 \text{ for all } p \in P.$$

Because the proof is similar but more technical, we do not give a proof here (we can provide the proof upon request).

B.2 Λ-robustness with $\Lambda \supset \Gamma^*$

Λ-robustness with $\Lambda \supset \Gamma^*$ is a stronger requirement than robustness. However, if Λ is sufficiently close to Γ^*, Λ-robustness is equivalent to robustness. To see this, for $\varepsilon > 0$, let

$$\Gamma(\varepsilon) = \{\lambda \in \Delta(\mathcal{U}) : \frac{E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i)|\delta u = x]}{E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i)|\delta u = x']} < 1 + \varepsilon \text{ for all } x, x' \in \mathcal{X} \text{ and } i \in N\}.$$

$0 < \varepsilon < 1/2$.

Note that $\Gamma^* \subsetneq \Gamma(\varepsilon)$ and $\Gamma^* = \bigcap_{\varepsilon>0}\Gamma(\varepsilon)$. Because Γ^*-robustness is equivalent to robustness, $\Gamma(\varepsilon)$-robustness is a stronger requirement than robustness. However, if $\varepsilon > 0$ is sufficiently small, $\Gamma(\varepsilon)$-robustness is equivalent to robustness as shown by the next lemma, which generalizes Lemma C.

Proposition B. *Let*

$$\underline{\varepsilon} = \min_{\phi\in\Phi:\phi \text{ is robust}} \min_{p\in\Delta(\mathcal{X})} \max_{i\in N} \frac{2p(\phi(x)=x_i)-1}{1-p(\phi(x)=x_i)},$$

which is strictly positive by the definition of robustness. Then, for any $\Lambda \subset \Delta(\mathcal{U})$ with $\Lambda^* \subseteq \Lambda \subseteq \Gamma(\underline{\varepsilon})$, a voting rule $\phi \in \Phi$ is Λ-robust if and only if it is robust.

Proof. It is enough to show that robustness implies $\Gamma(\underline{\varepsilon})$-robustness. Let $\phi \in \Phi$ be a robust rule. Fix arbitrary $\lambda \in \Gamma(\underline{\varepsilon})$, and let $i \in N$ be the individual who has the maximum responsiveness under λ. Define $\overline{D}_\lambda = \max_{x\in\mathcal{X}} E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i)|\delta u = x]$ and $\underline{D}_\lambda = \min_{x\in\mathcal{X}} E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i)|\delta u = x]$. Note that $\overline{D}_\lambda \geq \underline{D}_\lambda > 0$ and $\overline{D}_\lambda/\underline{D}_\lambda < 1 + \underline{\varepsilon}$ because $\lambda \in \Gamma(\underline{\varepsilon})$. Then,

$$\begin{aligned}
&E_\lambda[u_i(\phi(\delta u)) - u_i(-\phi(\delta u))] \\
&= \sum_{x:\phi(x)=x_i} E_\lambda[u_i(\delta u_i) - u_i(-\delta u_i)|\delta u = x]p_\lambda(x) + \sum_{x:\phi(x)=-x_i} E_\lambda[u_i(-\delta u_i) - u_i(\delta u_i)|\delta u = x]p_\lambda(x) \\
&\geq \underline{D}_\lambda \sum_{x:\phi(x)=x_i} p_\lambda(x) - \overline{D}_\lambda \sum_{x:\phi(x)=-x_i} p_\lambda(x) \\
&> \underline{D}_\lambda\Big(p_\lambda(\phi(x)=x_i) - (1+\underline{\varepsilon})p_\lambda(\phi(x)=-x_i)\Big) \\
&= \underline{D}_\lambda\Big(2p_\lambda(\phi(x)=x_i) - 1 - \underline{\varepsilon}(1-p_\lambda(\phi(x)=x_i))\Big) \geq 0,
\end{aligned}$$

where the last inequality holds by the construction of $\underline{\varepsilon}$. This implies that the inverse rule is not weakly Pareto-preferred to ϕ in terms of expected utility under λ. Because $\lambda \in \Gamma(\underline{\varepsilon})$ is an arbitrary element, ϕ must be $\Gamma(\underline{\varepsilon})$-robust. □

In contrast, if ε is too large, $\Gamma(\varepsilon)$-robustness is equivalent to the existence of a dictator.

Proposition C. *Let $\overline{\varepsilon} \equiv 2^n - 2$. Then, for any $\Lambda \subset \Delta(\mathcal{U})$ with $\Gamma(\overline{\varepsilon}) \subseteq \Lambda$, a voting rule $\phi \in \Phi$ is Λ-robust if and only if there exists a dictator $i \in N$ such that $\phi(x) = x_i$ for all $x \in \mathcal{X}$.*

Proof. It is enough to show that $\Gamma(\overline{\varepsilon})$-robustness implies the existence of a dictator. Let ϕ be $\Gamma(\overline{\varepsilon})$-robust. Seeking a contradiction, suppose that there is no dictator. For each $x \in \mathcal{X}$, let

$u^x \in \mathcal{U}$ be such that

$$(u_i^x(1), u_i^x(-1)) = \begin{cases} (1/(2^n - 1), 0) & \text{if } x_i = 1 = \phi(x), \\ (1, 0) & \text{if } x_i = 1 \neq \phi(x), \\ (0, 1/(2^n - 1)) & \text{if } x_i = -1 = \phi(x), \\ (0, 1) & \text{if } x_i = -1 \neq \phi(x), \end{cases}$$

and let $\lambda \in \Delta(\mathcal{U})$ be the uniform distribution over $\{u^x \in \mathcal{U} : x \in \mathcal{X}\}$. Note that $\lambda \in \Gamma(\overline{\varepsilon})$. Because there is no dictator and $\lambda(u^x) = 1/\#\mathcal{X} = 1/2^n$ for each $x \in \mathcal{X}$, the responsiveness of each individual is at most $1 - 1/2^n$, i.e., $p_\lambda(\phi(x) = x_i) \leq 1 - 1/2^n$ for each $i \in N$. Thus,

$$\begin{aligned} & E_\lambda[u_i(\phi(\delta u)) - u_i(-\phi(\delta u)] \\ &= \sum_{x:\phi(x)=x_i} E_\lambda[u_i(x_i) - u_i(-x_i)|\delta u = x] p_\lambda(x) + \sum_{x:\phi(x)=-x_i} E_\lambda[u_i(-x_i) - u_i(x_i)|\delta u = x] p_\lambda(x) \\ &= p_\lambda(\phi(x) = x_i)/(2^n - 1) - p_\lambda(\phi(x) = -x_i) \leq (1 - 1/2^n)/(2^n - 1) - 1/2^n = 0. \end{aligned}$$

This implies that the inverse rule is weakly Pareto-preferred to ϕ in terms of expected utility under λ and that ϕ is not $\Gamma(\overline{\varepsilon})$-robust. □