IMPLEMENTING THE MEDIAN*

MATÍAS NÚÑEZ[a], CARLOS PIMIENTA[b], AND DIMITRIOS XEFTERIS[c]

ABSTRACT. In the single-peaked domain, the median rules (Moulin [1980]) are of special interest. They are, essentially, the unique strategy-proof rules as well as the unique Nash implementable ones under complete information. We show that, under mild assumptions on admissible priors, they are also Bayes-Nash implementable by the means of "detail-free" mechanisms. That is, mechanisms that do not rely on the mechanism designer having detailed information about the priors that the agents hold. Furthermore, detail-free implementation of the median rules does not clash with truthful behavior. The provided mechanism is such that, in every equilibrium, all agents reveal their true peak with probability one.

KEYWORDS. Nash Implementation, Bayesian Implementation, Robust Implementation, Detail-free, Median rule, Strategy-proofness, Single-Peaked Preferences, Condorcet Winner.

JEL CLASSIFICATION. C9, D71, D78, H41.

1. INTRODUCTION

As it is well known, the domain of single-peaked preferences admits a unique set of implementation relevant welfare optima: the generalized median rules (see Moulin [1980]). Each generalized median rule (henceforth, GMR) is strategy-proof and, correspondingly, it is implemented in dominant strategies by its associated direct mechanism. Even though it seems intuitive that agents reveal their true type whenever doing so is a dominant strategy, as Saijo et al. [2007] argue, strategy-proof mechanisms have some drawbacks. One of them is that many strategy-proof mechanisms may have a large set of equilibrium outcomes that are undesirable from the point of view of the welfare optimum. This is, in particular, the case for the direct revelation mechanisms associated to the GMRs.[1] Since "one cannot simply assume that agents will play a vector of dominant strategies when alternative Nash or Bayesian equilibria exist" [Repullo, 1985] one should, whenever possible, turn to designing mechanisms that admit a unique equilibrium outcome. Barberà and Jackson [1994], Sprumont [1995] and Berga and Moreno [2009] show that, in environments with complete information, there are mechanisms that Nash implement all GMRs. However, due to their more realistic assumptions, we are also interested in studying implementation of GMRs in environments in which agents only have

* We thank Françoise Forges and Hervé Moulin for useful discussions. We also thank seminar and conference audiences for their comments.

[a] UNIVERSITÉ PARIS-DAUPHINE, PSL RESEARCH UNIVERSITY, CNRS, LAMSADE, 75016 PARIS, FRANCE.
[b] SCHOOL OF ECONOMICS, UNSW BUSINESS, THE UNIVERSITY OF NEW SOUTH WALES, SYDNEY, AUSTRALIA.
[c] UNIVERSITY OF CYPRUS, DEPARTMENT OF ECONOMICS, P.O. BOX 20537, 1678 NICOSIA, CYPRUS.

Date: June 24, 2018.

[1] For example, consider a set of agents with single-peaked preferences over the outcome space $[0,1]$ and the problem of implementing the median of their peaks using the associated direct revelation mechanism that selects the median of their announcements. As long as there are three or more agents, any outcome is an equilibrium outcome of the mechanism and it can be induced by the announcement profile in which every agent announces such an outcome (because any profile created after a unilateral deviation has the same median.) See Section 1.1 for a short review of Nash implementation in the single-peaked domain.

incomplete information about the peak profile of their opponents. Moreover, it is desirable to consider an environment in which the mechanism designer is uninformed about the underlying distribution of the voters' preferences so that she has to design a mechanism that is "detail-free" or, in other words, that does not rely on having available that specific information.

Existing results that assume complete information provide invaluable insights—and also practical solutions to the few problems for which such an assumption is plausible—but, as noted by Wilson [1987], "only by repeated weakening of common knowledge assumptions will the theory approximate reality." Hence, it is imperative to address whether GMRs can be implemented when the subjects' preferences are their private information (in the spirit of Bayesian implementation, such as Jackson [1991]) and the mechanism designer does not have information about the agents' preferences or the beliefs that they hold about the preferences of the rest of the agents (in the spirit of detail-free or robust approaches to mechanism design, such as Bergemann and Morris [2005] or Saijo et al. [2007]).

In this paper, we adopt a mainly positive approach.[2] We proceed by construction and build detail-free direct mechanisms that implement GMRs in Bayes-Nash equilibria under mild assumptions on the set of admissible priors. These mechanisms, the *value-based medians* (henceforth, VB-medians), are quite similar to the direct revelation games of the rules that they implement. The direct revelation game of a GMR selects the median of the combined profile made of (at most) $n-1$ exogenous points and the n peaks announced by the agents. In turn, a VB-median mechanism selects the median of the combined profile made of the those exogenous points and an endogenous selection of the announcements. For instance, consider the canonical median mechanism with an odd number of agents. This mechanism selects the announcement $m(s_1,\ldots,s_n)$ that divides the sample of announcements $(s_1,\ldots,s_n)$ in two exact halves. The associated VB-median proceeds in two steps. It first deletes duplicates in the vector of announcements $(s_1,\ldots,s_n)$. This leads to the vector of pruned announcements $(s'_1,\ldots,s'_m)$ with $m \le n$. If m is odd, the outcome of the VB-median is the median of $(s'_1,\ldots,s'_m)$. If m is even, the median of pruned announcements is the midpoint between the lower-median $m^-(s'_1,\ldots,s'_n)$ and the upper-median $m^+(s'_1,\ldots,s'_n)$.

We show that this detail-free mechanism Bayes-Nash implements the median when each individual's beliefs regarding the preferences of the rest of the agents is smooth (i.e., an atomless distributions with full support), and that a straightforward generalization implements other GMRs. While the restriction to smooth priors is limiting, they arguably constitute the unique most relevant class in the standard setting where GMRs were introduced by Moulin [1980]. The proposed mechanism also induces truthful behavior. In every equilibrium, every agent reports her true peak with probability one. To our knowledge, this is the first time that the GMRs are shown to be Bayes-Nash implementable in their natural setting—i.e., considering a continuous policy space and a finite number of agents—and, also, by truth-revealing mechanisms that do not require parametric alterations depending on environmental particularities. In other papers in the literature, the alternative space is considered to be discrete and, thus,

[2] This approach is quite common in the implementation literature, see Maskin [1999] and Abreu and Sen [1991].

any unanimous rule that satisfies no-veto power[3] in a single-peaked context (such as a GMR) fails Bayesian monotonicity in several reasonable classes of priors [Jackson, 1991, Duggan, 1995]. That is, existing approaches have little to say regarding the problem tackled in this paper.

Apart from establishing the above main results, the VB-median mechanisms are also the first simple simultaneous direct mechanisms in the literature that Nash implement GMRs under complete information. This is arguably of interest on its own, since, to our knowledge, there is only one other simple simultaneous mechanism (the approval mechanism, see Núñez and Xefteris [2017]) that Nash-implements the median rule (i.e., the Condorcet winner alternative) and this mechanism is indirect since the strategy set does not coincide with the peak space. Intuitively, the VB-median mechanisms generate unique equilibrium outcomes because they make every agent always decisive. This contrasts with the standard median mechanism in which, for some instances, no individual agent can affect the outcome. For example, if every agent is expected to announce the same outcome, under the standard median mechanism, a unilateral deviation cannot affect the median announcement. On the other hand, under the VB-median mechanism, identical announcements are ignored, which means that every announcement always affects the outcome crucially, thus, making inefficient coordination impossible.

After briefly reviewing the literature, we present the model in Section 2, the results on complete and incomplete information in Sections 3 and 4, and their extensions to GMRs in Section 5.

1.1. *Review of the Literature*

This paper is at the intersection of two rich branches of the literature: Implementation and Social Choice. Implementation focuses on designing mechanisms whose equilibrium outcome coincide with the outcome of a welfare optimum or social choice rule. If a mechanism has the property that, independently of the model's primitives, the set of equilibrium outcomes coincides with the set of outcomes identified by the welfare optimum, then the welfare optimum is said to be implemented by the given mechanism. Among the different notions of implementation, dominant strategy and (Bayes-)Nash equilibrium are, arguably, the most classic ones. Implementation in dominant strategies requires that each agent has a dominant strategy and that if agents use their dominant strategies, the desired welfare optimum is attained. Implementation in (Bayes-)Nash equilibrium (or full implementation) requires that all equilibrium outcomes of the mechanism be socially optimal.

The literature studying the single-peaked domain can be traced back to Galton [1907] and to the pioneering contributions of Black [1948] and Moulin [1980]. This early literature underlines a distinct feature of this domain: strategy-proof, anonymous, and efficient peak-only rules exist and are simple to describe. These mechanisms coincide with the direct revelation games of the generalized median rules, that is, of the procedures that select as an outcome the median of the peaks combined with some fixed weights (or phantoms) whose purpose is to calibrate the

[3]A rule that satisfies no-veto power selects an alternative if it is ranked at the top by all but at most one agent. See Benoît and Ok [2008] for a discussion.

final decision. Among the most recent contributions, Arribillaga and Massó [2016, 2017] compare the degree of manipulability of these voting schemes in other domains, while Chatterji et al. [2016] underline the salience of the single-peaked domain as the only one admitting a wide class of well-behaved social choice rules. As discussed above, a negative feature of the single-peaked domain was underlined by the secure implementation literature (Cason et al. [2006], Saijo et al. [2007]). They note that no mechanism is able to implement a generalized median rule both in dominant strategies and in Nash equilibrium. In other words, no direct mechanism implements a generalized median rule simultaneously in dominant strategies and in Nash equilibrium (i.e., no direct mechanism obtains full implementation).

As far as complete information is concerned, the literature has circumvented this impossibility of secure implementation in primarily two ways. The first approach consists of building direct mechanisms that help agents coordinate at the cost of removing incentives for truth-telling (see, e.g., Yamamura and Kawasaki [2013]). So far, this approach has not obtained implementation of the median rule since the only studied mechanisms (such as simple averaging of agents' reports) Nash implement generalized median rules with *exactly* $n-1$ interior and distinct phantoms. Moreover, the designed mechanisms do not generate incentives for truth-telling and for almost all admissible preference profiles, almost all agents largely misreport their preferences. The second approach (see Núñez and Xefteris [2017]) has designed indirect mechanisms that Nash implement any generalized median such that its interior phantoms (if any) are all distinct. This allows, in particular, implementing the median rule. To achieve this, agents approve of intervals of alternatives and the mechanism selects the median of these intervals. Finally, in an incomplete information setting, Gershkov et al. [2017] succeeds in partially implementing the median by building dynamic voting games in which agents vote until a qualified majority is reached. Thus, the literature is so far using mechanisms quite distant from a simple simultaneous direct mechanism to achieve the implementation of the median rule. The current paper fills this gap.

2. Setting

We consider a set of *agents* $N := \{1,\dots,n\}$ with $n \geq 3$ and a set of *alternatives* $A := [0,1]$ with typical elements x and y. For the time being, we assume that n is odd and deal with the case where n is even in Section 5. Each agent i has single peaked preferences with set of possible peaks $T := [0,1]$. Agent i's preferences are represented by means of the utility function $U_i : A \times T \to \mathbb{R}$. If agent i has peak t_i then $U_i(x \mid t_i) < U_i(y \mid t_i)$ whenever $x < y \leq t_i$ or $t_i \leq y < x$. We let $T^n := \prod_{i=1}^{n} T$ be the set of peak profiles. For any pair of alternatives x, $y \in A$, their midpoint is $\delta(x,y) := (x+y)/2$.

An announcement profile is an element of the set of all finite profiles $\mathbb{A} = \bigcup_{m=1}^{\infty}[0,1]^m$. For each $a \in \mathbb{A}$ we let $c(a)$ denote its length. A profile $a \in \mathbb{A}$ is odd if $c(a)$ is odd and it is even otherwise. For any two profiles $a,b \in \mathbb{A}$, we let (a,b) denote the profile of length $c(a)+c(b)$ resulting from appending the profile b to the profile a. We define the lower-median, upper-median, and median of an announcement profile as follows.

Definition 1. *For any profile $a \in \mathbb{A}$ with $c(a) = m$,*

- *the lower-median* $m^-(a)$ *is the smallest* a_k *for which* $\#\left\{\ell \mid a_\ell \leq a_k\right\} \geq \frac{m}{2}$, *and*
- *the upper-median* $m^+(a)$ *is the largest* $a_{k'}$ *for which* $\#\left\{\ell \mid a_\ell \geq a_{k'}\right\} \leq \frac{m}{2}$.

If a *is odd, then* $m^-(a) = m^+(a)$ *and such a common value is the median* $m(a)$ *of the profile* a.

We are interested in implementing *generalized median rules* (GMRs) which are, arguably, the only interesting Nash implementable rules in the single-peaked domain. In this domain, a welfare optimum is a GMR if and only if it has convex range and is both Pareto efficient and Nash-implementable.[4] We begin focusing on the *median rule* and extend the analysis to GMRs in Section 5.

Definition 2. *The median rule* $f_M : T^n \to A$ *associates to each* $t \in T^n$ *the alternative* $f_M(t) := m(t)$. *Its associated direct mechanism is the median mechanism* $\theta_M : T^n \to A$ *and it satisfies* $f_M = \theta_M$.

To define the *value-based median mechanism*, we introduce the function $v : \mathbb{A} \to \mathbb{A}$ that assigns to every $a \in \mathbb{A}$ the ordered profile $v(a)$ of *distinct* values in a. That is, $v(a) \in \mathbb{A}$ is the profile that is obtained after removing the minimum number of entries from a so that no two remaining entries are the same and then ordering them from smallest to largest. To simplify, the length of $v(a)$ is denoted $cv(a)$ rather than $c(v(a))$.

Definition 3. *The value-based median mechanism* $\theta_{VB} : T^n \to A$ *associates to each* $s \in T^n$ *the alternative*

$$\theta_{VB}(s) := \begin{cases} m(v(s)) & \text{if } v(s) \text{ is odd,} \\ \delta\left(m^-(v(s)), m^+(v(s))\right) & \text{if } v(s) \text{ is even.} \end{cases} \tag{1}$$

Thus, instead of selecting the median of the announcements, the value based median selects the median of the distinct announcements (if there are an odd number of them) or the mid-point between the lower-median and upper-median of the distinct announcements (if they are an even number). This difference generates different truthtelling incentives. To make them explicit, we recall the definition of strategy-proofness and the fact that the median mechanism θ_M satisfies it [Moulin, 1980].

Definition 4. *A mechanism* $\theta : T^n \to A$ *is strategy-proof if, for any agent* i, *revealing her true peak* t_i *is a best response to any* $s_{-i} \in T^{n-1}$, *that is*

$$U_i(\theta(t_i, s_{-i}) \mid t_i) \geq U_i(\theta(s_i, s_{-i}) \mid t_i) \text{ for any } s_i \in T \text{ and any } s_{-i} \in T^{n-1}.$$

The median mechanism is strategy-proof because an agent cannot do better that announcing her true peak given that any misreport can only shift the outcome further away from her peak. This is not the case for the value-based mechanism θ_{VB}.

Example 1. Let $N = \{1,2,3\}$ with $t = (1/3, 1/3, 2/3)$. Consider the profile $s = (1/3, 1/3, 2/3)$ in which every agent sincerely announces her peak. We have $\theta_{VB}(s) = \delta(1/3, 2/3) = 1/2$ because $t_1 = t_2 = 1/3$. In turn, for any $0 \leq x < 1/3$, $\theta_{VB}(x, s_{-1}) = 1/3$ which coincides with agent 1's peak.

[4] GMRs are "peak-only" welfare optima, that is, they only depend on the agents' peak profile and not on the shape of their utility functions. Sprumont [1995] shows that any strategy-proof welfare optimum whose range is an interval is a peak-only welfare optimum.

Hence, for such values of x, we obtain $U_1(\theta_{VB}(x, s_{-1}) \mid 1/3) > U_1(1/2 \mid 1/3)$, which shows that θ_{VB} is not strategy-proof.

Even if the mechanism θ_{VB} is not strategy-proof, we show below that it still gives truth-revealing incentives to every agent *for almost every* announcement profile of the other agents. This motivates the following weakening of strategy-proofness.

Definition 5. *A mechanism $\theta : T^n \to A$ is almost strategy-proof if, for any agent i, revealing her true peak t_i is a best response to almost all $s_{-i} \in T^{n-1}$. That is, if for each agent i there is some set $\tilde{T}_i \subset T^{n-1}$ of zero Lebesgue measure such that*

$$U_i(\theta(t_i, s_{-i}) \mid t_i) \geq U_i(\theta(s_i, s_{-i}) \mid t_i) \text{ for any } s_i \in T \text{ and any } s_{-i} \in T^{n-1} \setminus \tilde{T}_i.$$

Under the mechanism θ_{VB}, every agent i weakly prefers t_i to any other announcement for almost all announcement profiles of the other agents. In particular, agent i with peak t_i may have an incentive to misreport it only if some other agent already announces t_i (as in Example 1) or if, given the profile s_{-i}, we have that $cv(t_i, s_{-i})$ is even and that t_i is either the lower-median or upper-median of the entire announcement profile (see Definition 3). The next proposition establishes that the set of announcement profiles for which either of these alternatives holds has zero Lebesgue measure.

Proposition 1. *The value-based median mechanism is almost strategy-proof.*

Proof. For each agent i, let $\tilde{T}_i := \{s_{-i} \in T^{n-1} \mid cv(t_i, s_{-i}) < n\}$. The closed subset $\tilde{T}_i$ has zero Lebesgue measure because is the finite union of lower dimensional hyperplanes.

Take any $s_{-i} \in T^{n-1} \setminus \tilde{T}_i$. Since $cv(t_i, s_{-i})$ is odd, $\theta_{VB}(t_i, s_{-i}) = m(t_i, s_{-i})$. Consider an alternative message $s_i \neq t_i$. If $cv(s_i, s_{-i})$ is even, then $\theta_{VB}(s_i, s_{-i}) = \delta\big(m^-(v(s_i, s_{-i})), m^+(v(s_i, s_{-i}))\big)$, but agent i can induce her preferred point on the interval $\big[m^-\big(v(s_i, s_{-i})\big), m^+\big(v(s_i, s_{-i})\big)\big]$ by announcing t_i, so that t_i is a best response. If otherwise $cv(s_i, s_{-i})$ is odd, then $\theta_{VB}(s_i, s_{-i}) = m(s_i, s_{-i})$ is no closer to t_i than $\theta_{VB}(t_i, s_{-i})$, which again implies that t_i is a best response. □

3. Complete Information

In this section, we show that the mechanism θ_{VB} Nash implements the median rule in a complete information setting. We assume that all peaks are interior, so that no agent has a peak at either zero or one.[5] The strategy profile $s \in T^n$ is a Nash equilibrium of the game induced by the direct mechanism $\theta : T^n \to A$ at the profile $t \in T^n$ if for every agent $i \in N$ and every announcement $s'_i \in S_i$ we have $U_i(\theta(s_i, s_{-i}) \mid t_i) \geq U_i(\theta(s'_i, s_{-i}) \mid t_i)$. We denote by $N^\theta(t)$ the set of Nash equilibria of the game induced by the mechanism θ at the peak profile t.

Definition 6. *The mechanism $\theta : T^n \to A$ Nash implements the welfare optimum $f : T^n \to A$ if for each $t \in T^n$*

(1) there exists an equilibrium $s \in N^\theta(t)$ satisfying $\theta(s) = f(t)$
(2) for any $s \in N^\theta(t)$ we have $\theta(s) = f(t)$.

[5] This assumption guarantees that the median peak is an interior point of the set of announcements. This is not a substantial limitation and could be circumvented by means of alternative modelling choices (e.g. an unbounded outcome space in the spirit of Moulin [1980]). Alternatively, we can also expand the message space to $[-\varepsilon, 1+\epsilon]$ for some arbitrary $\varepsilon > 0$, thus, allowing agents to submit messages outside the interval $[0, 1]$.

As emphasized by Repullo [1985], the direct mechanism associated to a rule does not necessarily Nash implement such a rule. A prominent example of this fact is the median rule. The median mechanism θ_M admits a continuum of Nash equilibrium outcomes because all agents coordinating on the same announcement is always a Nash equilibrium. This is in contrast with the value-based mechanism θ_{VB}.

Example 2. Consider three agents with peak profile $t = (t_1, t_2, t_3)$ such that $t_1 \le t_2 \le t_3$ with at least one strict inequality. Take any $x \in A$ and consider the announcement profile (x,x,x). We obtain $\theta_M(x,x,x) = x$ and no unilateral deviation modifies the alternative implemented by the median mechanism because $\theta_M(y,x,x) = x$ for every $y \in [0,1]$. Thus, (x,x,x) is an equilibrium of the mechanism θ_M. In contrast, (x,x,x) is not an equilibrium of the value-based median mechanism θ_{VB}. Indeed, $\theta_{VB}(x,x,x) = x$ and note that we assumed that at least some agent j has a peak t_j different from x. Thus, if agent j deviates to t_j, the mechanism θ_{VB} induces $\delta(t_j, x)$ which is closet to t_j than x. That is, (x,x,x) is not an equilibrium of θ_{VB}.

Since the value-based median ignores redundant announcements, counting several equal announcements as just one, agents have an incentive to make their announcements unique. The next proposition formalizes this intuition.

Proposition 2. *Let $s \in T^n$ be an equilibrium of the mechanism θ_{VB} at the peak profile $t \in T^n$. Then,*

(1) if $t_i \neq \theta_{VB}(s)$ then agent i's announcement is unique, i.e. $s_i \neq s_j$ for every $j \in N \setminus \{i\}$, and
(2) if $t_i < \theta_{VB}(s)$ then $s_i < \theta_{VB}(s)$ whereas if $t_i > \theta_{VB}(s)$ then $s_i > \theta_{VB}(s)$.

Proof. We prove each part in turn. To prove (1), take some equilibrium $s \in T^n$ of θ_{VB} at the peak profile $t \in T^n$. Assume first that $cv(s)$ is even so that $v(s) = (v_1, \ldots, v_{2k})$ for some integer k and, therefore, $\theta_{VB}(s) = \delta(v_k, v_{k+1})$. Since n is odd and $cv(s)$ is even, there are $i, j \in N$ with $s_i = s_j$. We claim that if $t_i \neq \delta(v_k, v_{k+1})$ then agent i has a profitable deviation. If $t_i \in (v_k, v_{k+1})$, agent i can profitably deviate announcing her true peak t_i which induces $\theta_{VB}(t_i, s_{-i}) = t_i$. If $t_i \le v_k$, agent i can induce $\theta_{VB}(x, s_{-i}) = v_k$ by playing some x with $x < v_k$ and $x \neq s_j$ for any $j \neq i$. Finally, if $t_i \ge v_{k+1}$ a similar argument to the case in which $t_i \le v_k$ applies. Hence, in any equilibrium s with an even number of announcements, the strategy of any agent i with $t_i \neq \delta(v_k, v_{k+1})$ is unique.

Let $s \in T^n$ be an equilibrium of θ_{VB} with $cv(s)$ odd. Since $cv(s)$ is odd, there is some integer k such that $v(s) = (v_1, \ldots, v_{k+1}, \ldots, v_{2k+1})$ so that $\theta_{VB}(s) = m(v(s)) = v_{k+1}$. Let i be an agent with $t_i \neq v_{k+1}$ and assume to the contrary that $s_i = s_j$ for some $j \in N$. Again, agent i has a profitable deviation. If $v_k \le t_i < v_{k+1}$, then $\theta_{VB}(t_i, s_{-i}) = \delta(t_i, v_{k+1})$ so that t_i is a profitable deviation for agent i. If $t_i < v_k$, then any unique announcement s'_i with $s'_i < v_k$ induces outcome $\theta_{VB}(s'_i, s_{-i}) = \delta(v_k, v_{k+1})$ which agent i strictly prefers to v_{k+1}, so that any such s'_i is a profitable deviation for agent i. If $v_{k+1} < t_i$ a similar argument applies and shows that s is not an equilibrium as assumed.

To show (2), consider an equilibrium $s \in T^n$ of the mechanism θ_{VB}. If $cv(s)$ is odd, there is some integer k with $v(s) = (v_1, v_2, \ldots, v_{k+1}, \ldots, v_{2k+1})$ and $\theta(s) = m(v(s))$. Assume to the contrary that for some agent i we have $t_i < m(v(s))$ and $s_i > m(v(s))$. If t_i is not announced any agent under s_{-i} and $v_k < t_i < v_{k+1}$, then $\theta(t_i, s_{-i}) = t_i$, while if $t_i < v_k$, then $\theta(t_i, s_{-i}) = v_k$. If t_i is announced by some agent under s_{-i} then, for any deviation s'_i such that $s'_i < t_i$ and $s'_i \neq s_j$ for every $j \neq i$, we obtain $\theta(s'_i, s_{-i}) = \delta(v_k, v_{k+1})$. In all of these cases, agent i has a profitable deviation so that s is not an equilibrium. A similar contradiction arises when $cv(s)$ is even and when $t_i > m(v(s))$ and $s_i < m(v(s))$, concluding the proof. □

Thus, in equilibrium, every agent whose peak does not coincide with the equilibrium outcome makes a unique announcement that, furthermore, lies on the same side of the outcome as her peak. Since we assumed that the number of agents is odd, this means that, the number of different announcements is "typically" odd. The next example shows that equilibria with an even number of different announcements are possible, but they have a very special structure.

Example 3. Let $N = \{1, 2, 3, 4, 5\}$ with $t = (1/5, 2/5, 2/5, 2/5, 3/5)$ so that $m(t) = 2/5$. The profile $s = (0, 3/10, 3/10, 5/10, 6/10)$ satisfies $\theta_{VB}(s) = \delta(3/10, 5/10) = 2/5 = m(t)$. Agents 2, 3 and 4 do not have a profitable deviation because the outcome coincides with their peak. Agent 1 and 5 are best responding by making an announcement that is to the same side of the outcome as their peak. Therefore, s is an equilibrium with an even number of different announcement in which no agent announces 2/5. Regardless, those agents with peak at 2/5 coordinate in such a way to make it the equilibrium outcome.

We are ready to conclude with the main result of this section.

Theorem 1. *The value-based median θ_{VB} Nash implements the median rule.*

Proof. We first show that $m(t)$ is the unique equilibrium outcome for every peak profile t. Let s be an equilibrium of θ_{VB}. From Proposition 2, every agent i with peak $t_i < \theta_{VB}(s)$ plays $s_i < \theta_{VB}(s)$ while every agent with peak $t_i > \theta_{VB}(s)$ plays $s_i > \theta_{VB}(s)$. Furthermore, they send unique announcements in equilibrium so that the difference between the number of agents whose peak is strictly larger and strictly smaller than $\theta_{VB}(s)$ is no larger than the number of agents whose peak is $\theta_{VB}(s)$. There is at least one agent with peak $\theta_{VB}(s)$ as otherwise s would not be an equilibrium. Indeed, if no agent has peak $\theta_{VB}(s)$ then $cv(s) = n$ because every announcement would be different from each other. This would imply $\theta_{VB}(s) = m(s)$ and the agent sending announcement $m(s)$ would have a different peak and would not be optimizing. But these observations about the number of peaks that are larger, smaller, and equal to $\theta_{VB}(s)$, together with the fact that they must add up to n readily imply that $\theta_{VB}(s) = m(t)$, as we wanted.

We now show that, for every $t \in T^n$, there is some equilibrium s of θ_{VB} such that $\theta_{VB}(s) = m(t)$. Relabelling if necessary, let $t = (t_1, \ldots, t_n)$ be such that $t_1 \leq \ldots \leq t_k \leq t_{k+1} \leq t_{k+2} \leq \ldots t_n$, with $n = 2k+1$ for some k. Select a strategy profile s satisfying (1) $s_{k+1} = t_{k+1}$, (2) $0 < s_i < t_{k+1}$ for any $i = 1, \ldots, k$, (3) $1 > s_i > t_{k+1}$ for any $i = k+2, \ldots, n$, and (4) for any $1 \leq i, j \leq n$, if $t_i < t_j$ then $s_i < s_j$. For any such s, $v(s) = s$ and $\theta_{VB}(s) = t_{k+1}$. We must prove that s is an equilibrium. The agent with peak t_{k+1} does not have an incentive to deviate by construction. Every agent with peak $t_i < t_{k+1}$ announces $s_i < t_{k+1}$ and has three possible deviations. If she deviates to $s'_i = s_j$ for some $j \in N$, induces outcome $\theta_{VB}(s'_i, s_{-i}) = \delta(t_{k+1}, s_{k+2})$ so that it is not a profitable deviation. If she plays some unique announcement $s'_i < t_{k+1}$, then $\theta_{VB}(s'_i, s_{-i}) = \theta_{VB}(s_i, s_{-i})$. Finally, if she plays some unique announcement $s'_i > t_{k+1}$, then the outcome $\theta_{VB}(s'_i, s_{-i})$ is either s'_i (if $s'_i < s_{k+2}$) or s_{k+2} (if $s'_i > s_{k+2}$); and neither of these deviations are profitable for agent i since $t_i < t_{k+1}$. An analogous argument applies if the agent has peak $t_i > t_{k+1}$. Therefore, s is an equilibrium of the mechanism θ_{VB}. □

4. Incomplete Information

We now extend the arguments of the previous section to an incomplete information environment. We show that the value-based median mechanism induces a unique equilibrium outcome given that every agent reveals her peak truthfully with probability one. As already mentioned in the introduction, the VB-median mechanism is "detail-free" or "nonparametric", that is, it does not depend on the priors held by the agents or the mechanism designer. Using the terminology of Saijo et al. [2007] (see also Bergemann and Morris, 2005), the mechanism θ_{VB} *robustly and truthfully* implements the median rule when each agents' priors are restricted to be a probability measure that is absolutely continuous with respect to the Lebesgue measure over the set of peak profiles.

Nonetheless, to simplify notation, we make two additional assumptions on beliefs: (1) agents have a common prior, and (2) peaks are independently distributed. But we remark that the arguments of the proof do not depend on either of them. Intuitively, given that the VB-median mechanism is almost strategy proof (Proposition 1), being truthful in the interim stage is optimal against all but a subset of announcement profiles of zero Lebesgue measure. Furthermore, it is the unique optimal action if it is expected that the true peak is the median of the announcements with positive probability. Therefore, if an agent expects every other agent to be truthful (so that they play according to the identity function) and she has an absolutely continuous prior over the set of peak profiles, the subset of announcement profiles of the opponents for which being truthful is not optimal always receives probability zero. And, moreover, every peak that the agent may have is the median of the announcements with positive probability. Therefore, truthtelling is an equilibrium.

In turn, the strategy of the proof to show that, in equilibrium, every agent truthfully reveals her peak with probability one proceeds in two steps. We first show that no agent makes an announcement with positive probability (i.e., for a positive probability subset of her peaks) if she anticipates that some other agent will also make that same announcement with positive probability. And, second, we note that that implies that there are, typically, an odd number of announcements, fact that we use to prove that every agent must anticipate that her peak will be the median of the announcements with positive probability. In both steps, the relevant piece of information about beliefs is understanding which subsets of peak profiles have positive or zero probability. As long as beliefs over peak profiles are absolutely continuous with respect to the Lebesgue measure, those subsets are the same and common to all agents.

Thus, for every agent i, let F_i be an absolutely continuous distribution function over T. For each i, we let F_{-i} denote the probability distribution on T^{n-1} induced by the family of distribution functions $F_1, \ldots, F_{i-1}, F_{i+1} \ldots, F_n$. Given $F_1, \ldots, F_n$, a mechanism characterized by the Borel-measurable function $\theta : T^n \to \Delta(A)$ induces an n-agent Bayesian game. Agent i's strategy set is the collection of all Borel-measurable functions $\sigma_i : T \to T$. We write $\sigma_{-i}(t_{-i}) = (\sigma_1(t_1), \dots, \sigma_{i-1}(t_{i-1}), \sigma_{i+1}(t_{i+1}) \dots, \sigma_n(t_n))$.

A strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is an equilibrium of the Bayesian game induced by the mechanism $\theta : T^n \to A$ if no agent has an incentive to unilaterally deviate, that is, if for every agent $i \in N$, and any strategy $\varsigma_i : T \to T$ we have

$$\int_T \left[\int_{T^{n-1}} U_i\big(\theta(\sigma_i(t_i), \sigma_{-i}(t_{-i})) \mid t_i\big) dF_{-i}(t_{-i}) \right] dF_i(t_i) \geq \int_T \left[\int_{T^{n-1}} U_i\big(\theta(\varsigma_i(t_i), \sigma_{-i}(t_{-i})) \mid t_i\big) dF_{-i}(t_{-i}) \right] dF_i(t_i).$$

For every agent i, the distribution function F_i and the strategy σ_i induce a distribution function over announcements G_i defined by $G_i(a) := \int_{\sigma_i^{-1}([0,a])} dF_i$. We write $G := \prod_i G_i$ and $G_{-i} := \prod_{j \neq i} G_j$. The support $\mathscr{S}(G_i)$ of G_i is defined as the collection of points $x \in A$ such that for every open subset $\mathscr{O}$ satisfying $x \in \mathscr{O}$ we have $\int_{\mathscr{O}} dG_i > 0$. Strategy $\varsigma_i : T \to T$ of agent i is a best response against σ_{-i} if it maximizes

$$\int_T \left[\int_{T^{n-1}} U_i(\theta(\varsigma_i(t_i), s_{-i}) \mid t_i) dG_{-i}(s_{-i}) \right] dF_i(t_i). \tag{2}$$

Before proving our main result, we formally recall what it means to implement a welfare optimum in this incomplete information environment (see Jackson [1991], Duggan [1997] and Saijo et al. [2007] among others).

Definition 7. *The mechanism $\theta : T^n \to A$ Bayes-Nash implements the welfare optimum $f : T^n \to A$ in an incomplete information environment if*

(1) the Bayesian game induced by $\theta : T^n \to A$ has at least one equilibrium, and
(2) every equilibrium satisfies $\theta(\sigma_1(t_1), \dots, \sigma_n(t_n)) = f(t)$ for almost every peak profile $t \in T^n$.

Recall that, in this incomplete information setting, the median mechanism admits a large multiplicity of equilibrium outcomes (every point in A is an equilibrium outcome) and hence it fails to implement the median rule.

Theorem 2. *The value-based median mechanism θ_{VB} Bayes-Nash implements the median rule. Furthermore, in every equilibrium, an agent's strategy may only differ from truthful revelation in a zero measure set of peaks.*

Proof. Given Proposition 1, the strategy profile in which every agent reveals her true peak is an equilibrium of θ_{VB}. Therefore, for almost every peak profile $t \in T^n$, the outcome induced by θ_{VB} under such an equilibrium is $m(t)$.

We now prove that, in every equilibrium, every agent plays truthfully apart from, possibly, a set of peaks with zero measure.[6] Let $\sigma = (\sigma_i)_{i \in N}$ be an equilibrium. For every agent i consider the set of isolated announcements in $[0,1]$ that are sent with positive probability by agents

[6] By assumption, each distribution function F_i is absolutely continuous. Therefore, if ς_i is a best response against σ_{-i} and ς_i' coincides with ς_i almost everywhere on T, the value of (2) does not change and shows that ς_i' is also a best response against σ_{-i}. Hence, it is enough to show that, in any equilibrium, every agent is playing truthfully apart from a set of peaks of measure zero.

other than i under σ. That is, let Φ_i^σ be the set of points $x \in [0,1]$ such that

$$\int_{\sigma_j^{-1}(x)} dF_j > 0 \text{ for some } j \neq i. \tag{3}$$

Similarly, we define the superset of Φ_i^σ that that contains all its midpoints, that is

$$\tilde{\Phi}_i^\sigma = \Phi_i^\sigma \bigcup \left\{ x \in A \mid x = \delta(y,z) \text{ for some } y, x \in \Phi_i^\sigma \right\}. \tag{4}$$

We claim that, in equilibrium, no agent sends an announcement that coincides with some other agent's announcement with positive probability. To the contrary, assume that the set of peaks $\sigma_i^{-1}(a_i)$ for which agent i announces $a_i \in \Phi_i^\sigma$ has positive positive probability. Agent i has a profitable deviation to the strategy σ_i' obtained from σ_i by replacing the announcement a_i that is sent when agent i has a peak in $\sigma_i^{-1}(a_i) \setminus \tilde{\Phi}_i^\sigma$ by the corresponding truthful announcement. After such a deviation, if $t_i \in \sigma_i^{-1}(a_i) \setminus \tilde{\Phi}_i^\sigma$, the truthful announcement t_i is unique with probability one, while a_i is redundant with positive probability. Using the same arguments as in Proposition 2, it follows that $U_i(\theta_{VB}(t_i, s_{-i}) \mid t_i) > U_i(\theta_{VB}(a_i, s_{-i}) \mid t_i)$ for every s_{-i} such that $s_j = a_i$ for some $j \neq i$ and $t_i \neq \theta_{VB}(a_i, s_{-i})$. Therefore,

$$\int_{T^{n-1}} U_i(\theta_{VB}(t_i, s_{-i}) \mid t_i) dG_{-i}(s_{-i}) > \int_{T^{n-1}} U_i(\theta_{VB}(a_i, s_{-i}) \mid t_i) dG_{-i}(s_{-i})$$

for every $t_i \in \sigma_i^{-1}(a_i) \setminus \tilde{\Phi}_i^\sigma$ because some agent other than i announces a_i with positive probability and we only have $t_i = \theta_{VB}(a_i, s_{-i})$ with probability zero according to G_{-i}. Since the set $\tilde{\Phi}_i^\sigma$ is at most countable, agent i's peak t_i belongs to $\sigma_i^{-1}(a_i) \setminus \tilde{\Phi}_i^\sigma$ with positive probability. That is, the strategy σ_i' is indeed a profitable deviation.

Thus, there is a product set $Q = \prod_{i=1}^n Q_i \in \mathbb{A}$ that receives positive probability under the induced distribution function G and satisfies the following two properties:

(1) for every $i \in N$, the subset Q_i is either a closed interval or a singleton, and
(2) for any two $i, j \in N$ we have $Q_i \cap Q_j = \emptyset$.

Condition (2) implies that there is an agent, say agent k, such that for any arbitrary $q \in Q$ we have $q_k = m(q)$. Consider the following two bounds,

$$a_k := \max \left\{ x \in \bigcup_{i \neq k} Q_i : x < \min Q_k \right\} \text{ and } b_k := \min \left\{ x \in \bigcup_{i \neq k} Q_i : x > \max Q_k \right\}.$$

Given our choice of Q we have $a_k < b_k$. If agent k has peak $t_k \in (a_k, b_k)$ then the unique announcement s_k that maximizes $U_k(s_k, q_{-k} \mid t_k)$ is $s_k = t_k$. Therefore, if agent k has peak $t_k \in (a_k, b_k) \setminus \Phi_i^\sigma$ the unique announcement s_k that maximizes $\int_{T^{n-1}} U_k(s_k, s_{-k} \mid t_k) dG_{-k}(s_{-k})$ is also $s_k = t_k$. Hence, since the set Φ_i^σ is at most countable, agent k's unique best response for almost every peak she might have in (a_k, b_k) is to reveal her true peak. But this implies that for almost every peak in (a_k, b_k) that agent $i \neq k$ may have, agent i also plays truthfully under σ. This is because, with positive probability under σ, each $j \notin \{i, k\}$ makes an announcement in Q_j and agent k makes an announcement in (a_k, b_k). Hence, if agent i's truthful announcement lies in (a_k, b_k), then it is the median of the different announcements with positive probability. We conclude that, every agent with a peak almost everywhere in (a_k, b_k) plays truthfully under σ.

The set of open intervals $(\hat{a}, \hat{b})$ with $(a_k, b_k) \subset (\hat{a}, \hat{b}) \subset [0,1]$ such that every agent with a peak almost everywhere in $(\hat{a}, \hat{b})$ truthfully reports her peak under σ can be partially ordered by set inclusion. Furthermore, every totally ordered subset of such intervals has an upper bound. Therefore, by Zorn's Lemma, there is a largest interval (a^*, b^*) for which such a property is true. Suppose that $a^* > 0$. Because (a^*, b^*) is the largest interval of its family, there is at least one agent, say, agent j whose support $\mathscr{S}(G_j)$ is a subset of $[a^*, 1]$. If agent j has peak $t_j \in [0, a^*) \setminus \Phi_j^\sigma$, then she strictly prefers a truthful announcement to any announcement in $(a^*, 1]$ because the latter would induce a worse outcome with strictly positive probability. Since $\mathscr{S}(G_j) \subset [a^*, 1]$, agent j announces a^* with positive probability. This implies that any other agent $j' \neq j$ with peak in $[0, a^*) \setminus \Phi_{j'}^\sigma$ strictly prefers a truthful announcement to announcing a^*, which is redundant with positive probability, and also to any announcement in $(a^*, 1]$, for the same reason as agent j above. Hence, for every agent $j' \neq j$ we have $\mathscr{S}(G_{j'}) \cap [0, a^*) \neq \varnothing$ and, in such a case, the optimal response for a positive measure set of peaks of agent j in $[0, a^*)$ is a truthful announcement. This contradicts $\mathscr{S}(G_j) \subset [a^*, 1]$, therefore, $a^* = 0$. For an analogous reason, $b^* = 1$.

We conclude that, in equilibrium, every agent makes a truthful announcement for almost every peak in $[0, 1]$. Therefore, every equilibrium σ satisfies $\theta(\sigma_1(t_1), \ldots, \sigma_n(t_n)) = m(t)$ for almost every peak profile $t \in T^n$. □

Remark 1. A similar proof shows that the mechanism θ_{VB} has a unique interim Bayesian equilibrium and that, in such an equilibrium, truthtelling is the unique optimal action for every agent at every information set.

To our knowledge, the above result is the first one showing that the median rule is Bayes-Nash implementable in general environments of incomplete information. This confirms the privileged empirical relevance that the median rule and its variants (see the next section) have within the class of welfare optima in the single-peaked domain.

5. Generalized Median Rules

Each generalized median rule (GMR) selects the median of a profile of peaks *together* with a collection of fixed elements (called *phantoms*). For instance, the lowest peak in the peak profile t of length n is the median of the profile obtained after adding $n-1$ points to the profile t that are all smaller than any peak in t. Of course, any k-th order statistic of a peak profile (which selects the k-th largest peak) can be expressed in a similar way.

Before formally defining GMRs, we enlarge the sets X and A to allow for points that are smaller (larger) than any conceivable peak. Fix some arbitrary $\varepsilon > 0$ and define the expanded set of alternatives $A^\varepsilon := [-\varepsilon, 1+\varepsilon]$ and the set of finite profiles $\mathbb{A}^\varepsilon := \bigcup_{h=1}^\infty [-\varepsilon, 1+\varepsilon]^h$. Extend the functions v, c, m, m^-, and m^+ to $\mathbb{A}^\varepsilon$ in the obvious way.

We say that a profile $y \in \mathbb{A}^\varepsilon$ of length h is a *permissible* profile of phantoms if (i) $h < n$, (ii) $n + h$ is odd, and (iii) every entry in y is different.

Definition 8. *Given a permissible profile of phantoms y, the welfare optimum $f_y : T^n \to A$ is the generalized median rule associated to y if for each $t \in T^n$ we have*

$$f_y(t) = m(t, y).$$

Example 4. Let n be even. Then $y = (0)$ is a permissible profile of phantoms and we can say that $f_y(t)$ is the median rule given that, for every $t \in T^n$, we have $f_y(t) = m(0,t) = m^-(t)$. (Alternatively, if $y = (1)$ the rule $f_y(t)$ selects the upper-median as the median of an even profile.)

Example 5. Take some integer $k \leq n/2$. The social optimum that selects the k-th lowest (respectively, highest) peak can be expressed as a GMR by selecting a permissible profile of phantoms y of length $n-2k+1$ and such that every one of its entries lies in the subinterval $[-\varepsilon, 0)$ (respectively, in the subinterval $[1, 1+\varepsilon]$).

Similarly to the median mechanism θ_M, using f_y as the outcome function of a direct mechanism produces a game with a large multiplicity of equilibrium outcomes due to coordination problems. To discourage agents from coordinating on the same announcement we construct a value-based mechanism θ_y that counts several equal announcements as one.

Definition 9. *The generalized value-based median mechanism associated to the permissible phantom profile y is the direct random mechanism $\theta_y : T^n \to \Delta(A^\varepsilon)$ defined so that, for each strategy profile $s \in T^n$,*

$$\theta_y(s) := \begin{cases} m(v(s,y)) & \text{if } v(s,y) \text{ is odd , and} \\ \delta\Big(m^-(v(s,y)), m^+(v(s,y))\Big) & \text{if } v(s,y) \text{ is even.} \end{cases}$$

Notice that, under this mechanism, phantoms are treated as if they were announcements made by sincere agents with known peaks.

Theorem 3. *The mechanism θ_y implements the GMR f_y under complete information.*

Proof. Let h be the length of the permissible profile of phantoms y. We first argue that if s is an equilibrium of the complete information mechanism at the peak profile t, then it satisfies $\theta_y(s) = m(t,y)$. Because y is a permissible profile of phantoms, $n+h$ is odd. Hence, the same arguments used in Proposition 2 prove that, if s is an equilibrium, no agent makes an announcement either equal to some other agent's announcement or equal to a phantom. Furthermore, they also show that every agent with peak different from $\theta_y(s)$ makes an announcement that is to the same side of the outcome as her peak. Thus, the difference between the number of entries in (t,y) that are strictly larger and smaller than $\theta_y(s)$ is no larger than the number of entries in (t,y) equal to $\theta_y(s)$. Since at least one entry in (t,y) is equal to the outcome induced by s, for analogous reasons as in the Theorem 1, such an outcome must coincide with $m(t,y)$ as we wanted.

To show that the complete information mechanism has at least one equilibrium s such that $\theta_y(s) = m(s,y)$, construct a strategy profile along the lines of the second part of the proof of Theorem 1. The only difference here is that elements in y are not agent announcements and, therefore, have a fixed position. However, this difference does not impose any relevant restriction and arguments analogous to those in that proof show that the mechanism has, indeed, at least one equilibrium that induces the desired outcome. □

In the same vein, we also obtain an implementation result for GMRs under the incomplete information environment described in Section 4.

Theorem 4. *The mechanism θ_y implements the GMR f_y under incomplete information. Moreover, in every equilibrium, every agent reveals her true peak with probability one.*

Proof. Proposition 1 implies that if agent i's peak t_i does not coincide with a phantom then she cannot do better than reporting her true peak if her opponents truthfully report their peak as well and, therefore, they announce t_i with probability zero. Hence, truthful revelation is an equilibrium of the mechanism.

Let h be the length of the permissible profile of phantoms y. With the same notation as in the proof of Theorem 2 we argue that every equilibrium σ satisfies $\theta_y(\sigma_1(t_1),\ldots,\sigma_n(t_n)) = m(t,y)$ for almost every peak profile $t \in T^n$.

Therefore, let σ be an equilibrium and let the induced probability distribution G and the set Φ_i^σ be as in the proof of Theorem 2. We can use analogous arguments to the ones used in the proof of Theorem 2 to show that no agent i makes an announcement in $\Phi_i^\sigma \cup \{y_1\} \cup \cdots \cup \{y_h\}$ with positive probability. Therefore, we can choose a collection of sets $\{Q_i\}_{i=1}^{n+h}$ such that $\{Q_i\}_{i=1}^{n}$ receives positive probability under G and:

(1) for $i = 1,\ldots,n$, the subset Q_i is either a closed interval or a singleton,
(2) for $i = n+1,\ldots,n+h$ we have $Q_i = \{y_{i-n}\}$, and
(3) for any two $i, j \in N$ we have $Q_i \cap Q_j = \emptyset$.

The product set Q can be used to select an agent k such that for any $q \in Q$ we have $m(q) = q_k$. Instead of the interval (a_k, b_k) constructed in the proof of Theorem 2 we need to use the interval $(a_k, b_k) \cap [0,1]$ to make sure that it is a subset of the set of peaks T. From this point forward the proof of Theorem 2 applies almost verbatim to the current case. □

Remark 2. The only reason why an agent with known peak may find it advantageous to misreport it is if there is a phantom that coincides with her peak. Therefore, any *GMR* with no interior phantoms (i.e., any k-th order statistic) can also be implemented using interim Bayesian equilibrium. Thus, under such an assumption, the value-based mechanism has a unique and sincere equilibrium.

References

D. Abreu and A. Sen. Virtual implementation in Nash equilibrium. *Econometrica*, 59(4):997–1021, 1991.

R. P. Arribillaga and J. Massó. Comparing generalized median voter schemes according to their manipulability. *Theoretical Economics*, 11(2):547–586, 2016.

R. P. Arribillaga and J. Massó. Comparing voting by committees according to their manipulability. *American Economic Journal: Microeconomics*, 9(4):74–107, 2017.

S. Barberà and M. Jackson. A characterization of strategy-proof social choice functions for economies with pure public goods. *Social Choice and Welfare*, 11(3):241–252, 1994.

J.-P. Benoît and E.A. Ok. Nash implementation without no-veto power. *Games and Economic Behavior*, 64(1):51–67, 2008.

D. Berga and B. Moreno. Strategic requirements with indifference: single-peaked versus single-plateaued preferences. *Social Choice and Welfare*, 32(2):275–298, 2009.

D. Bergemann and S. Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005.

D. Black. On the rationale of group-decision making. *Journal of Political Economy*, 56:23–34, 1948.

T.N. Cason, T. Saijo, T. Sjöström, and T. Yamato. Secure implementation experiments: Do strategy-proof mechanisms really work? *Games and Economic Behavior*, 57:206–235, 2006.

S. Chatterji, A. Sen, and H. Zeng. A characterization of single-peaked preferences via random social choice functions. *Theoretical Economics*, 11(2):711–733, 2016.

J. Duggan. *Bayesian implementation*. PhD thesis, California Institute of Technology, 1995.

J. Duggan. Virtual Bayesian implementation. *Econometrica*, pages 1175–1199, 1997.

F. Galton. One Vote, One Value. *Nature*, 75:414, 1907.

A. Gershkov, B. Moldovanu, and X. Shi. Optimal voting rules. *Review of Economic Studies*, 84: 688–717, 2017.

M. Jackson. Bayesian implementation. *Econometrica*, 59:461–477, 1991.

E. Maskin. Nash Equilibrium and Welfare Optimality. *Review of Economic Studies*, 66:23–38, 1999.

H. Moulin. On Strategy-proofness and Single Peakedness. *Public Choice*, 35:437–455, 1980.

M. Núñez and D. Xefteris. Implementation via approval mechanisms. *Journal of Economic Theory*, 170:169–181, 2017.

R. Repullo. Implementation in dominant strategies under complete and in- complete information. *Review of Economics Studies*, 52:223–229, 1985.

T. Saijo, T. Sjöström, and T. Yamato. Secure implementation. *Theoretical Economics*, 2:203–229, 2007.

Y. Sprumont. Strategyproof collective choice in economic and political environments. *Canadian Journal of Economics*, pages 68–107, 1995.

R. Wilson. Game-theoretic analysis of trading processes. In Truman Fassett Bewley, editor, *Advances in Economic Theory: Fifth World Congress*, chapter 2, pages 33–70. Cambridge University Press, Cambridge UK, 1987.

H. Yamamura and R. Kawasaki. Generalized Average Rules as stable Nash mechanisms to implement generalized median rules. *Social Choice and Welfare*, 40:815–832, 2013.