ON THE IMPLEMENTATION OF THE MEDIAN*

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

ABSTRACT. In the single-peaked domain, the median rule is strategy-proof but not implementable in (Bayes-)Nash equilibrium by its associated direct mechanism. We define the value-based median mechanism that implements the median rule in (Bayes-)Nash equilibrium in the single-peaked domain under complete and incomplete information. Such a mechanism selects the median of the profile of different values announced by the agents (i.e., ignoring redundant announcements). The value-based median does not depend on agents' beliefs (in line with robust mechanism design). In the case of incomplete information, it induces truthful revelation of preferences (in line with strategy-proofness) for almost all peaks. We present extensions of our results to generalized median rules and finite policy spaces and their limitations.

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

JEL CLASSIFICATION. D71, D78, H41.

1. INTRODUCTION

The median rule is well known to be strategy-proof in the single-peaked domain. Therefore, revealing one's actual peak is always a best response under the associated direct mechanism (henceforth, *median mechanism*). Yet, as observed by Repullo [1985] and later generalized by Saijo et al. [2007], the direct mechanism associated with a rule need not Nash implement such a rule. In many environments of interest, if the direct mechanism associated with a rule implements the rule in dominant strategies, it fails to implement it in (Bayes-)Nash equilibria. There are reasons to investigate implementation in equilibrium even if the rule in question is implementable in dominant strategies. Indeed, there is a literature dealing with the agents' difficulties to recognize and behave according to the fact that truth-telling is always a best response in strategy-proof mechanisms. This literature contains both experimental evidence (see Kagel et al. [1987], Kagel and Levin [1993], Attiyeh et al. [2000], Kawagoe and Mori [2001], Cason et al. [2006]) and recent theoretical developments on *obviously strategy-proofness* (see Li [2017] and Arribillaga et al. [2020]).[1]

---

* We thank Françoise Forges and Hervé Moulin for useful discussions. We also thank the associate editor and two anonymous referees for insightful advice as well as seminar and conference audiences for their comments. Financial support from the ANR-14-CE24-0007-01 (CoCoRICoCoDEC) project and from the ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047 project is acknowledged.
[a] CNRS, CREST & ECOLE POLYTECHNIQUE, 91128 PALAISEAU, FRANCE. *Corresponding author*.
[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: December 21, 2021.
[1] A notable exception is the quasi-linear environments with divisible private or public goods, where there exist surplus-maximizing social choice functions that are implementable both in Nash equilibrium and dominant strategies. Recent work on this area includes Saporiti [2014] secure implementation under partial honesty and its characterization of the augmented representative voter schemes class, a generalization of the generalized median rules.

The classical domain of single-peaked preferences is one environment where it is impossible to reconcile strategy-proofness with implementation in (Bayes-)Nash equilibrium. Saijo et al. [2007] shows that any social choice rule that is implementable in Nash equilibrium and dominant strategies must be either dictatorial or Pareto inefficient. In particular, the median mechanism admits a multiplicity of equilibrium outcomes so that it does not implement the median rule in Bayes-Nash equilibria.[2] This observation leads to the following questions: Is the median mechanism implementable in equilibrium by some other direct mechanism? Does it matter for equilibrium implementation whether agents' preferences are public or private information?

We present a unified treatment of both questions. We study the implementation of the median rule in (Bayes-)Nash equilibrium in both complete and incomplete information settings and show that, with a continuum set of alternatives, the median rule is indeed implementable using a direct mechanism: the value-based median (see a detailed description at the end of the introduction). Therefore, we show that one can design mechanisms whose unique equilibrium outcome coincides with the alternative indicated by the median rule. These positive results provide a novel intuition since, to the extent of our knowledge, little is known about (Bayes-) Nash equilibrium implementation of the median: it satisfies Maskin monotonicity (and hence is Nash implementable), and there are indirect rules that Nash implement it (see Sprumont [1995], Barberà and Jackson [1994], Berga and Moreno [2009] and Núñez and Xefteris [2017]). Notably, not much is known about Bayes-Nash equilibrium implementation of social choice rules in the single-peaked domain when the agents' preferences are private information. We build a direct mechanism whose unique Bayes-Nash equilibrium is truthful and whose outcome coincides with the one prescribed by the median rule with probability one.

Apart from establishing that the median rule is implementable in (Bayes-)Nash equilibrium in its usual setting (finite number of voters and a continuum of alternatives) we also show that we can achieve this in a robust and truthful manner. The former property relates to *robust mechanism design* (see Bergemann and Morris [2005] and Saijo et al. [2007]), and it allows the designer to be ignorant about the agents' beliefs. The latter is a feature typically found in strategy-proof mechanisms. These results seem to be of independent interest as the trade-off between strategy-proofness and (Bayes-)Nash equilibrium implementation has attracted recent attention within the literature. Our findings imply that in the single-peaked domain with a continuum of alternatives, it is possible to combine appealing features of all approaches: i.e., it is possible to design a mechanism that a) does not depend on the exact beliefs that agents hold about the preferences of the other agents, b) involves truthful behavior by every agent, and c) whose equilibrium outcome coincides with the alternative indicated by the median rule.

We also extend our results and explore to which extent they can be used to implement GMRs or when the set of alternatives (and preference types) is finite. In terms of GMRs,[3] we show that any GMR with distinct phantoms is Nash implementable through a direct mechanism in environments of complete information, and we establish the analogous result under incomplete

[2] This extends to any anonymous and strategy-proof rule in this domain. In other words, any direct mechanism associated with any GMR, as characterized by Moulin [1980], generates a plethora of equilibrium outcomes.

[3] A GMR maps each set of points to the median of the union of those points with a fixed collection of exogenous points called *phantoms*.

information. When the set of alternatives is finite (as detailed in the appendix), the implementation of the median rule depends on the information that agents hold about other agents' preferences. With complete information, we can obtain Nash implementation of the median rule using direct and minimally probabilistic mechanisms (i.e., off-equilibrium, the mechanisms induce a lottery over alternatives). Regarding incomplete information, we present an example (based on a result by Triossi [2005]) that shows that the median rule fails Bayesian Monotonicity (a necessary condition for Bayes-Nash implementation, see Jackson [1991] among others).

We now turn to describe the value-based median mechanism and its differences with respect to the median mechanism.

The value-based median mechanism. Consider the median mechanism with an odd number of agents and a continuum of alternatives. 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 value-based median mechanism proceeds in two steps. It first deletes redundant announcements in $(s_1,\ldots,s_n)$ so that if two or more reports coincide, it only keeps one of them. This leads to the vector of pruned announcements $(s'_1,\ldots,s'_{n'})$ with n' distinct values. If n' is odd, the outcome of the value-based median is the median of $(s'_1,\ldots,s'_{n'})$. If n' is even, the outcome is the midpoint between the lower-median $m^-(s'_1,\ldots,s'_{n'})$ and the upper-median $m^+(s'_1,\ldots,s'_{n'})$. Hence, the main difference between the median mechanism and the value-based one is that the latter ignores redundant announcements. To see how this difference solves equilibrium coordination problems, consider the median mechanism and the strategy profile where every agent announces x. Such profile is an equilibrium irrespectively of the agents' peaks because no agent can affect the outcome x by deviating unilaterally. Under the value-based median, this is no longer an equilibrium. Indeed, at such profile, every agent is pivotal since by deviating to any value y with $y \neq x$, the outcome shifts from x to $(x+y)/2$. Ignoring equal reports removes inefficient equilibria and ensures that the unique equilibrium outcome is the median of the true peaks.

After briefly reviewing the literature in Section2, we present the model in Section 3, and the results on complete and incomplete information in Sections 4 and 5. We discuss their extensions and limitations to GMRs and discrete policy spaces in, respectively, Sections 6 and Appendix A.

2. Literature review

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 coincides with the ones of a social choice rule. Suppose 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 social choice rule. In that case, the social choice rule is implemented by this mechanism. Implementation in (Bayes-)Nash equilibrium (or full implementation) requires that all equilibrium outcomes of the mechanism coincide with the outcome selected by the social choice rule.

The literature studying the single-peaked domain can be traced back to Galton [1907]. This literature has mainly considered two different problems: fair division and public good provision. In the literature dealing with fair division, a perfectly divisible resource is divided among a group of agents with single-peaked preferences. In this setting, the uniform rule is strategy-proof and plays a central role. It was introduced by Benassy [1982] and later characterized by Sprumont [1991] (see also Ching [1994]). The most recent approach, by Bochet et al. [2021], gives conditions to ensure that several classes of (non-strategy-proof) direct revelation mechanisms Nash implement the uniform rule.

In turn, this paper belongs to the public good provision literature. Within this literature, voters have single-peaked preferences over the quantity of public good to be provided. In this setting, the median rule is strategy-proof and plays a similar role to the uniform rule in the fair division literature. This rule was first considered by Black [1948] and later characterized by Moulin [1980], as a member of the class of strategy-proof, anonymous, and Pareto efficient peak-only rules, the generalized median rules. These rules select the median of the peaks combined with some fixed weights (or phantoms) whose purpose is to calibrate the final decision. Among the most recent contributions, Arribillaga and Massó [2016, 2017] compare the degree of manipulability of these voting rules in other domains, while Chatterji et al. [2016] underlines the salience of the single-peaked domain as the only one admitting a wide class of well-behaved social choice rules.

As discussed in the Introduction, a limitation of strategy-proofness relevant to the single-peaked domain was underlined by the secure implementation literature (see Cason et al. [2006] and Saijo et al. [2007]). They note that no Pareto efficient mechanism which is non-dictatorial can be both strategy-proof and Nash implementable. The literature has circumvented this impossibility in two main ways.[4] The first approach consists of building direct mechanisms that help agents coordinate at the cost of dispensing of truth-telling incentives. For instance, Yamamura and Kawasaki [2013] show that a wide range of voting mechanisms which includes the average mechanism can Nash implement some generalized median rules in Nash equilibria (their work is hence the symmetric counterpart of Bochet et al. [2021] in the public good provision problem). So far, this approach has not obtained implementation of the median rule since the only studied mechanisms 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 Maskin [1999] and Núñez and Xefteris [2017]) has designed indirect mechanisms that Nash implement any generalized median rule such that its interior phantoms (if any) are all distinct. This allows, in particular, implementing the median rule. Maskin [1999] builds the well-known integer games that apply to any setting. Núñez and Xefteris [2017] allow agents to 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

[4] We focus here on the literature that modifies the mechanism and leaves unchanged the equilibrium notion. Alternatively, one could study the same mechanism under different equilibrium concepts such as partial honesty (see Dutta and Sen [2012]) or the more classic notion of weakly undominated strategies.

dynamic voting games in which agents vote until a qualified majority is reached. Thus, the literature uses mechanisms quite distant from a simple simultaneous direct mechanism to achieve the implementation of the median rule. The current paper fills this gap.

Finally, our discussion in the Appendix that considers a finite set of alternatives is related to the strand of the literature that focuses on implementation with off-equilibrium lotteries. Indeed, the classic implementation approach focuses on deterministic mechanisms such as integer games. Sanver [2006], Bochet [2007] and Benoît and Ok [2008] prove that the set of Nash implementable rules crucially expands if one considers mechanisms that select a single alternative in equilibrium while selecting lotteries or awards off-equilibrium. More recently, Laslier et al. [2021] shows how to construct two-player mechanisms with off-equilibrium threats that ensure Pareto efficiency in any equilibrium. This contrasts with the impossibility of achieving the same goal with deterministic mechanisms as proved by Maskin [1999] and Hurwicz and Schmeidler [1978].

3. 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. We assume that n is odd and deal with the case where n is even in Section 6. 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 Borel-measurable 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 n-dimensional element of the set of profiles $\mathbb{A} = \bigcup_{k=1}^n [0,1]^k$. 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 $a,b \in \mathbb{A}$ with $c(a)+c(b) \leq n$, 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 $a \in \mathbb{A}$,*

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

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

A *social choice rule* is a peak-only function $f : T^n \to A$ that associates an alternative to each peak profile.

Definition 2. *The median rule $f_M : T^n \to A$ associates to each $t \in T^n$ the alternative $f_M(t) := m(t)$.*

A *mechanism* is a function $\theta : S^n \to A$ that assigns a unique element $\theta(s)$ in A to each $s \in S^n$, where S is the strategy space of agent i. In a *direct mechanism*, the set of announcements coincides with the set of peaks so that $S = T$. The *median mechanism* is the direct mechanism $\theta_M : T^n \to A$ with $\theta_M(s) = m(s)$. 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 obtained after removing the minimal number of entries from a so that no two remaining entries coincide 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*

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

The difference between the median and the value-based median mechanism is then clear: the median mechanism selects the median of all announcements whereas the value-based median selects the median of the distinct announcements (with an odd number of them) or the midpoint between the lower-median and upper-median of the distinct announcements (with an even number). This difference generates different truth-telling incentives. Indeed, the former is strategy-proof whereas the latter one fails to be so. To make this difference explicit, we recall the definition of strategy-proofness.

Definition 4. *A direct 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 no agent can do better than announcing her true peak given that any misreport can only shift the outcome further away from her peak.

The value-based mechanism is not strategy-proof because there are announcement profiles of the opponents such that misreporting is profitable as shown by Example 1.

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 truthfully announces her peak. We have $\theta_{VB}(s) = \delta(1/3, 2/3) = 1/2$ because two entries in the strategy profile s coincide. Note that, for any $0 \leq x < 1/3$, $\theta_{VB}(x, s_{-1}) = 1/3$. Thus, $U_1(\theta_{VB}(x, s_{-1}) \mid 1/3) > U_1(1/2 \mid 1/3)$ for any $0 \leq x < 1/3$ since $1/3$ coincides with agent 1's peak. This shows that θ_{VB} is not strategy-proof.

There are two situations in which an agent might find profitable not to reveal her peak in a value-based median. First, agent i with peak t_i may have an incentive to misreport if some other agent already announces t_i. As seen in Example 1, her message will be deleted by the value-based median. Hence, she prefers to announce some value different from t_i to have a higher weight in the final decision. The second situation in which misreporting is profitable is when $cv(t_i, s_{-i})$ is even and t_i coincides with either the lower-median $m^-(v(s))$ or the upper-median $m^+(v(s))$. In this case, the outcome is $\delta(m^-(v(s)), m^+(v(s)))$ and and the agent can do better making any announcement that lies strictly between $\delta(m^-(v(s)), m^+(v(s)))$ and her true peak.

Despite violating strategy-proofness, the value-based median mechanism exhibits good strategic incentives. As the next result shows, for almost all announcements of the other agents every agent i weakly prefers to announce her peak. 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 and for each type t_i, there is some set $\tilde{T}_i(t_i) \subset T^{n-1}$ of zero Lebesgue measure such that*

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

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 with type t_i, let $\tilde{T}_i(t_i) := \{s_{-i} \in T^{n-1} \mid cv(t_i, s_{-i}) < n\}$. The subset $\tilde{T}_i(t_i)$ is closed and has zero Lebesgue measure because is the finite union of lower dimensional hyperplanes. Thus, the set $T^{n-1} \setminus \tilde{T}_i(t_i)$ has Lebesgue measure one and each $s_{-i} \in T^{n-1} \setminus \tilde{T}_i(t_i)$ satisfies $cv(t_i, s_{-i}) = n$. Since n is odd by assumption, $cv(t_i, s_{-i})$ is odd for any $s_{-i} \in T^{n-1} \setminus \tilde{T}_i$. Therefore, $\theta_{VB}(t_i, s_{-i}) = m(t_i, s_{-i})$ for each such s_{-i}.

If $s_i \neq t_i$ and $cv(s_i, s_{-i})$ is even then $\theta_{VB}(s_i, s_{-i}) = \delta\big(m^-(v(s_i, s_{-i})), m^+(v(s_i, s_{-i}))\big)$. Therefore, 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 better response than s_i.

If $s_i \neq t_i$ and $cv(s_i, s_{-i})$ is odd then $\theta_{VB}(s_i, s_{-i}) = m(s_i, s_{-i})$. This again implies that t_i is a better response than s_i and concludes the proof. □

4. Complete Information

In this Section, we show that the mechanism θ_{VB} Nash implements the median rule in a complete information setting. Within this Section, we assume that the median peak is interior (i.e. not in 0 or in 1) but this assumption can be dispensed with.[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 $s'_i \in T$ 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 social choice rule $f : T^n \to A$ 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 $\theta(s) = f(t)$.

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 as illustrated by the next example.

Example 2. Let $N = \{1,2,3\}$ with peak profile $t = (t_1, t_2, t_3)$ such that $t_1 \leq t_2 \leq t_3$ 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 $\theta_M(y,x,x) = x$ for every $y \in [0,1]$. Thus, (x,x,x) is an equilibrium of the mechanism θ_M. In contrast, there is no $x \in [0,1]$ such that (x,x,x) is an equilibrium of the

[5] This is not a substantial limitation and could be circumvented by means of alternative modeling choices (e.g., an unbounded outcome space in the spirit of Moulin [1980]).

value-based median mechanism. Observe that $\theta_{VB}(x,x,x) = x$. By assumption, there is at least some agent j with a peak t_j different from x. Thus, since $\theta_{VB}(t_j,x,x) = \delta(t_j,x)$ is closer to t_j than x, agent j has a profitable deviation. This shows that (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. 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 $v(s) = (v_1,\dots,v_{2k})$ 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 $t_i \leq v_k$, 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 $t_i \geq v_{k+1}$ a similar argument to the case in which $t_i \leq v_k$ 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 $v(s) = (v_1,\dots,v_{k+1},\dots,v_{2k+1})$ 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 $v_k \leq t_i < v_{k+1}$, 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 $v(s) = (v_1,v_2,\dots,v_{k+1},\dots,v_{2k+1})$ and $\theta(s) = m(v(s))$. Assume to the contrary that for some agent i we have $t_i < m(v(s))$ and $s_i > m(v(s))$.

Given that t_i is not announced by any agent under s_{-i}, if $v_k < t_i < v_{k+1}$, then,

$$\theta_{VB}(t_i,s_{-i}) = \begin{cases} t_i & \text{if } s_i \neq s_j \text{ for all } j \neq i \\ \delta(t_i,v_{k+1}) & \text{otherwise} \end{cases}$$

while if $t_i < v_k$, then

$$\theta_{VB}(t_i,s_{-i}) = \begin{cases} v_k & \text{if } s_i \neq s_j \text{ for all } j \neq i \\ \delta(v_k,v_{k+1}) & \text{otherwise} \end{cases}$$

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 $\theta_{VB}(s'_i,s_{-i}) = \delta(v_k,v_{k+1})$. 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 any equilibrium, every agent whose peak does not coincide with the equilibrium outcome makes a unique announcement. Furthermore, each agent's announcement 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 this equilibrium, the agents with peak at 2/5 coordinate to ensure that 2/5 is the equilibrium outcome even though no agent announces 2/5.

Even though some agents may misreport their peaks in equilibrium and equilibria with an even number of different announcement exist, the next theorem shows that the equilibrium outcome is always the median of the peaks.

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. Suppose that there is some equilibrium s with outcome $\theta_{VB}(s) > m(t)$ (a symmetric argument applies when $\theta_{VB}(s) < m(t)$). For each strategy profile $s \in T^n$, we let $N_s^- = \{i \in N \mid t_i < \theta_{VB}(s)\}$, $N_s^+ = \{i \in N \mid t_i > \theta_{VB}(s)\}$ and $N_s^= = \{i \in N \mid t_i = \theta_{VB}(s)\}$ respectively denote the set of agents with a peak lower, higher than and equal to the outcome of profile s. Proposition 2 implies that each agent $i \in N_s^-$ plays $s_i < \theta_{VB}(s)$ and $i \in N_s^+$ plays $s_i > \theta_{VB}(s)$ and each such announcement is different from the others. Moreover, $\theta_{VB}(s) > m(t)$ implies $\#N_s^- > \#N_s^+ + \#N_s^=$. If $cv(s)$ is odd, then $\theta_{VB}(s) = m(v(s))$. Yet, $\#N_s^- > \#N_s^+ + \#N_s^=$ implies that $m(v(s))$ is announced by some agent in $\#N_s^-$. Since any such agent plays $s_i < \theta_{VB}(s)$, we obtain that $\theta_{VB}(s) < \theta_{VB}(s)$, a contradiction. If $cv(s)$ is even, then $\theta_{VB}(s) = \delta\left(m^-(v(s)), m^+(v(s))\right)$. Yet, $\#N_s^- > \#N_s^+ + \#N_s^=$ implies that both $m^-(v(s))$ and $m^+(v(s))$ are announcements of agents in $\#N_s^-$. Since any such agent plays $s_i < \theta_{VB}(s)$, we obtain again that $m^+(v(s)) < \theta_{VB}(s)$, a contradiction. This implies that any equilibrium s satisfies $\theta_{VB}(s) = m(t)$, as wanted.

We now show that, for every $t \in T^n$, there is some equilibrium s of θ_{VB} such that $\theta_{VB}(s) = m(t)$. Relabeling 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. We now build a strategy profile such that the median agent truthfully reveals her peak, the agents with a peak lower than the median peak play to the left of the median peak and the agents with a peak higher than the median peak play to the right of the median peak, while ensuring that every pair of agents with a different peak make different announcement. We then show that this strategy profile is an equilibrium. 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 now 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}, concluding the proof. □

Example 1 (Continued). Consider again the announcement profile $(x, 1/3, 2/3)$ with $0 \leq x < 1/3$. We have $\theta_{VB}(x, 1/3, 2/3) = 1/3$. This profile is a Nash equilibrium of θ_{VB} because (i) the induced outcome coincide with agent 1 and 2's peak, so they do not have a profitable deviation and (ii) agent 3 can only induce outcomes in $[x, 1/3]$ by unilaterally deviating, so she does a profitable deviation either. Moreover, the set of equilibria equals

$$\Big\{(x, 1/3, y) \mid 0 \leq x < 1/3 \text{ and } 1/3 < y \leq 1\Big\} \bigcup \Big\{(1/3, x, y) \mid 0 \leq x < 1/3 \text{ and } 1/3 < y \leq 1\Big\}.$$

As anticipated by the previous theorem, every profile in this set induces outcome 1/3, i.e. the median of the type profile $(1/3, 1/3, 2/3)$.

Remark 1. From Yamamura and Kawasaki [2013] we know that no mechanism which has unrestricted range, is anonymous, own-peak continuous, *and* own-peak strictly monotonic can implement the median. The value-based mechanism θ_{VB} is neither continuous nor monotonic. The main strategic incentives induced by θ_{VB} are described in Proposition 2. Chiefly, by disincentivizing redundant announcements, θ_{VB} ensures that agents whose peaks are strictly smaller (larger) than the outcome make an announcement strictly smaller (larger) than the outcome. As a consequence, the outcome coincides with the median of the peaks. However, lack of continuity is not necessary to obtain such a split. Indeed, we can construct continuous mechanisms that implement the median but, as it happens with θ_{VB}, they will not be own-peak strictly monotonic.[6]

5. Incomplete Information

We show that the value-based median induces a unique Bayes-Nash equilibrium outcome in which every agent reveals her true peak with probability one. As already mentioned in the introduction, the value-based median mechanism is "detail-free" or "non-parametric", that is, it does not depend on the priors held by the agents or the mechanism designer. Thus, we follow the literature on robust implementation and, using the terminology of Saijo et al. [2007] (see also Bergemann and Morris, 2005), we say that the mechanism θ_{VB} *robustly and truthfully*

[6] For instance, suppose that the outcome space is $\mathbb{R}$ and there is an odd number of voters. Consider the continuous mechanism $g(s) = \frac{1}{n}\sum_{i=1}^{n} y_i(s)$, where the function y_i is given by

$$y_i(s) = \begin{cases} m(s) - 1 & \text{if } s_i < m(s) - 1, \\ s_i & \text{if } m(s) - 1 \leq s_i \leq m(s) + 1, \\ m(s) + 1 & \text{if } s_i > m(s) - 1. \end{cases}$$

It can be checked that for any peak profile t, every equilibrium s^* of the mechanism has every agent whose peak is smaller (larger) than $m(t)$ announcing $m(s^*) - 1$ $(m(s^*) + 1)$, an odd number of agents with peak $m(t)$ announcing $m(t)$ and the rest, if any, announcing either $m(s^*) - 1$ or $m(s^*) + 1$ in a way that the median announcement coincides with $m(t)$.

implements the median rule when each agents' priors are probability measures which are absolutely continuous with respect to the Lebesgue measure over the set of peak profiles.

Formally, let Δ be the set of probability measures over the set of peak profiles T^n. The set of belief profiles is $\Delta^n = \Delta \times \cdots \times \Delta$ and agent i's prior belief over T^n is $\mathscr{F}^i$. A *support* for $\mathscr{F}^i$ is any set $C^i \subset T^n$ such that $\mathscr{F}^i(C^i) = 1$. Let $D \subset \Delta^n$ be the set of prior belief profiles $(\mathscr{F}^1, \mathscr{F}^2, \ldots, \mathscr{F}^n)$ that the agents could possibly have. The mechanism designer does not know the agents' true priors; she knows only that they belong to the set D.

We use the *robust and truthful* implementation notion introduced by Saijo et al. [2007].

Definition 7 (Saijo et al. [2007]). *The mechanism θ robustly and truthfully implements f on the domain D if for any collection of priors $(\mathscr{F}^1, \mathscr{F}^2, \ldots, \mathscr{F}^n) \in D$,*

(1) truthful revelation is a Bayes-Nash equilibrium and, for every agent i, there is a support C^i for $\mathscr{F}^i$ such that $\theta(t) = f(t)$ for all $t \in \bigcap_{i=1}^n C^i$.

(2) for any Bayes-Nash equilibrium σ under $(\mathscr{F}^1, \mathscr{F}^2, \ldots, \mathscr{F}^n)$ there is, for every agent i, a support C^i for $\mathscr{F}^i$ such that $\theta(\sigma(t)) = f(t)$ for all $t \in \bigcap_{i=1}^n C^i$.

There are two aspects of this definition that are worth commenting on. The first one is that Saijo et al. [2007]'s definition uses the ex-ante Bayes-Nash equilibrium concept (see expression (2) above). That is, agents optimize their strategy before their type is realized. As we explain in Remark 2, the implementation result also holds if we instead use the interim Bayes-Nash equilibrium concept in which agent's behavior must be optimal conditional on her peak realization. The second one is that, if a mechanism robustly and truthfully implements f_M, then then the equilibrium outcome of the mechanism coincides with the median of the peaks with probability one, but it is not necessarily equal to the median of the peaks *for every* peak profile. As we now easily show, the latter is too strong of a requirement when it comes to the median rule.

Proposition 3. *If mechanism θ robustly and truthfully implements the median rule f_M then there must be a peak profile $t = (t_1, \ldots, t_n)$ such that $\theta(t) \neq f_M(t)$.*

Proof. The median rule f_M is not robustly and truthfully implemented by itself (the constant announcement profile in which every agent announces the same $x \in [0,1]$ for every peak is always a Bayes-Nash equilibrium of the mechanism f_M). Hence, if a mechanism θ robustly and truthfully implement f_M then they must differ for at least one peak profile. □

We define the class of *smooth priors* D^S as the set of all belief profiles $(\mathscr{F}^1, \mathscr{F}^2, \ldots, \mathscr{F}^n) \in \Delta^n$ such that, for every $i = 1, \ldots, n$, agent i's prior belief is absolutely continuous with respect to the Lebesgue measure. We also define the class of *common and independent priors* D^I as the subset of D^S for which (1) agents have a common prior, and (2) peaks are independently distributed. Therefore, agent i's prior $\mathscr{F}^i$ has distribution function of the form $F_1 \times \cdots \times F_n$ where the distribution function F_j represents the common prior about agent j. In what follows we prove that the mechanism θ_{VB} *robustly and truthfully* implements the median rule on the domain D^I. After that, we argue that the result extends to the entire class D^S of smooth priors.

Intuitively, given that the value-based median mechanism is almost strategy-proof (Proposition 1), reporting one's peak truthfully is optimal against all but a subset of announcement profiles of zero Lebesgue measure. Furthermore, truth-telling is the unique optimal action if the agent expects her true peak to be the median of the announcements with positive probability. Therefore, if an agent expects every other agent to be truthful 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 has probability zero. Moreover, every peak that the agent could possibly have is the median of the announcements with positive probability. Therefore, truth-telling is an equilibrium.

The strategy of the proof consists of showing that, in equilibrium, every agent truthfully reveals her peak with probability one. The proof proceeds in two steps. We first show that, if an agent i anticipates that some other agent j will make some announcement s_j with positive probability, then agent i does not play s_j with positive probability. Since the number of agents is odd, we note that the previous observation implies that there are, typically, an odd number of announcements. Building on this fact, we prove that every agent anticipates that her peak is the median of the announcements with positive probability and that, moreover, it coincides with some other agent's announcement with zero probability.

For each agent i, let F_i be an absolutely continuous distribution function over T. The distribution F_i is the common prior held by all agents about agent i's peak. For each i, we let F_{-i} denote the probability distribution on T^{n-1} induced by the profile $F_1,\dots,F_{i-1},F_{i+1}\dots,F_n$. Given $F_1,\dots,F_n$, a mechanism is a Borel-measurable function $\theta : T^n \to \Delta(A)$ that induces an n-agent Bayesian game. Agent i's strategy set is the collection of all Borel-measurable functions $\sigma_i : T \to A$. For each $t_{-i} \in T^{n-1}$, we let $\sigma_{-i}(t_{-i}) := (\sigma_1(t_1),\dots,\sigma_{i-1}(t_{i-1}),\sigma_{i+1}(t_{i+1})\dots,\sigma_n(t_n))$ denote the strategy profile of agents other than agent i.

A strategy profile $\sigma = (\sigma_1,\dots,\sigma_n)$ is a Bayes-Nash equilibrium of the game induced by $\theta : T^n \to A$ if for every agent $i \in N$, and any strategy $\varsigma_i : T \to T$ we have

$$\int_T \left[\int_{T^{n-1}} U_i\big(\theta(\sigma_i(t_i),\sigma_{-i}(t_{-i})) \mid t_i\big) dF_{-i}(t_{-i}) \right] dF_i(t_i) \geq \int_T \left[\int_{T^{n-1}} U_i\big(\theta(t_i',\sigma_{-i}(t_{-i})) \mid t_i\big) dF_{-i}(t_{-i}) \right] dF_i(t_i). \quad (2)$$

In the proof, it is sometimes convenient to directly work the probability distribution over announcements induced by a strategy and the profile of common priors $F_1,\dots,F_n$. Consider every agent i. The distribution function F_i and the strategy σ_i induce a distribution function $G_i^{\sigma_i}$ over announcements; so that the probability that agent i makes an announcement smaller than or equal to a is given by $G_i^{\sigma_i}(a) := \int_{\sigma_i^{-1}([0,a])} dF_i$. Given a strategy profile σ, we write $G^\sigma := \prod_i G_i^{\sigma_i}$ and $G_{-i}^{\sigma_{-i}} := \prod_{j\neq i} G_j^{\sigma_j}$. The carrier $\mathscr{C}(G_i^{\sigma_i})$ of $G_i^{\sigma_i}$ is the collection of points $x \in A$ such that for every open subset $\mathscr{O}$ satisfying $x \in \mathscr{O}$ we have $\int_{\mathscr{O}} dG_i^\sigma > 0$. When there is no possibility of confusion, we simply write G_i and G instead of $G_i^{\sigma_i}$ and G^σ.

As the next result shows, the outcome induced by truthful play is the unique equilibrium outcome of the value-based median mechanism.

Proposition 4. *The value-based median mechanism θ_{VB} robustly and truthfully implements the median rule on the domain of common and independent priors D^I.*

Proof. Let $\sigma = (\sigma_i)_{i\in N}$ be an equilibrium. Given that priors are absolutely continuous with respect to the Lebesgue measure, it is enough to prove that every agent i reports truthfully for almost every peak in $[0,1]$. We start by putting aside a countable number of her peaks. For each $i \in N$, let Φ_i^σ denote the set of points in $[0,1]$ that receive positive probability according to the probability distribution over announcements G_j for some $j \neq i$. That is, points $x \in [0,1]$ such that

$$G_j(t) - \lim_{\varepsilon\to 0^+} G_j(t-\varepsilon) > 0 \text{ for some } j \neq i \tag{3}$$

We also need to set aside the set of agent i's peaks that would coincide with outcomes that receive positive probability according to G_{-i} if agent i sent a redundant announcement. To that end, we define the superset of Φ_i^σ that contains all its midpoints

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

As anticipated, the set $\tilde{\Phi}_i^\sigma$ is at most countable.

Claim 1. No agent i makes an announcement in Φ_i^σ with positive probability, that is, the set of peaks $\{t \in T : \sigma_i(t) \in \Phi_i^\sigma\}$ has probability zero according to F_i.

Proof of the claim. Take agent i's peak $t \notin \tilde{\Phi}_i^\sigma$ and assume $x := \sigma_i(t)$ satisfies $x \in \Phi_i^\sigma$. Take any announcement profile s_{-i} such that $s_j \neq t$ for every $j \neq i$ and $s_k = x$ for some $k \neq i$. If $\theta_{VB}(x, s_{-i}) \neq t$, following the same arguments used in the proof of Proposition 2 (a) we obtain

$$U_i(\theta_{VB}(t,s_{-i}) \mid t) > U_i(\theta_{VB}(x,s_{-i}) \mid t),$$

Furthermore, the set of announcement profiles s_{-i} such that some agent $j \neq i$ announces x receives strictly positive probability according to G_{-i} because $x \in \Phi_i^\sigma$. Likewise, the probability that some agent $j \neq i$ announces t, or that $\theta_{VB}(x,s_{-i}) = t$ is zero according to G_{-i} because $t \notin \tilde{\Phi}_i^\sigma$. The combination of these two facts implies

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

Since the previous inequality is true for any type $t \notin \tilde{\Phi}_i^\sigma$ and the set $\tilde{\Phi}_i^\sigma$ is at most countable, if F_i assigns positive probability to $\{t \in T : \sigma_i(t) \in \Phi_i^\sigma\}$ we obtain

$$\int_T \left[\int_{T^{n-1}} U_i(\theta_{VB}(t,s_{-i}) \mid t) dG_{-i}(s_{-i}) \right] dF_i(t_i) >$$
$$\int_T \left[\int_{T^{n-1}} U_i(\theta_{VB}(\sigma_i(t),s_{-i}) \mid t) dG_{-i}(s_{-i}) \right] dF_i(t_i).$$

But σ_i is an equilibrium strategy, so the set of peaks above must receive zero probability according to F_i. □

The previous claim implies that if σ is an equilibrium, the probability that two agents end up submitting the same announcement is zero. Thus, an agent's optimal strategy depends only on those announcement profiles in which all announcements are different. Furthermore, Claim 1 also implies that there is a product set $Q = \prod_{i=1}^n Q_i \subset [0,1]^n$ with:

(1) for each $i \in N$, $Q_i \subset \mathscr{C}(G_i)$,
(2) for each $i \in N$, the subset Q_i is either a closed interval or a singleton, and

(3) for each pair $i, j \in N$ we have $Q_i \cap Q_j = \emptyset$.

Pick some arbitrary $q \in Q$. Condition (3) implies there is an agent, say agent k, such that $q_k = m(q)$. 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 $U_k(\theta_{VB}(s_k, q_{-k}) \mid t_k)$ is $s_k = t_k$. Therefore, if agent k has peak $t_k \in (a_k, b_k) \setminus \tilde{\Phi}_i^\sigma$, then no other agent makes the same announcement with positive probability, so $\int_{T^{n-1}} U_k(\theta_{VB}(x, s_{-k}) \mid t_k) dG_{-k}(s_{-k})$ is maximal only when $x = t_k$. Since the set $\tilde{\Phi}_i^\sigma$ is countable, and agent k's strategy σ_k is optimal, it must prescribe a truthful announcement for almost every peak in (a_k, b_k).

But this implies that for almost every peak in (a_k, b_k) agent $i \neq k$ 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 plays truthfully under σ for almost every peak in (a_k, b_k).

We let $\mathscr{I}$ denote the collection of open intervals (a, b) such that every interval satisfies (i) $(a_k, b_k) \subseteq (a, b)$ and (ii) for every agent i the strategy σ_i prescribes truthful revelation for almost every peak in (a, b). The set $\mathscr{I}$ can be partially ordered by $\subseteq$. Furthermore, every totally ordered subset of $\mathscr{I}$ has an upper bound.[7] Therefore, by Zorn's Lemma, the set $\mathscr{I}$ contains a maximal element that we denote by (a^*, b^*).

To show $a^* = 0$ we assume to the contrary $a^* > 0$. Let ε be arbitrarily small. If every agent made an announcement smaller than $a^* - \varepsilon$ with positive probability, then every agent i with a peak in $(a^* - \varepsilon, b^*) \setminus \tilde{\Phi}_i^\sigma$ would maximize her interim utility by making a truthful announcement. However, (a^*, b^*) is the largest interval of its family. Thus, there is at least one agent, say, agent j whose carrier $\mathscr{C}(G_j)$ is a subset of $[a^*, 1]$.

We consider such an agent j, any peak $t \in [0, a^*) \setminus \tilde{\Phi}_j^\sigma$, and any announcement $t' \in (a^*, 1]$. We must have

$$\int_{T^{n-1}} U_j(\theta_{VB}(t, s_{-i}) \mid t) dG_{-j}(s_{-j}) > \int_{T^{n-1}} U_j(\theta_{VB}(t', s_{-i}) \mid t) dG_{-j}(s_{-j})$$

because every other agent is being truthful with probability one on (a^*, b^*), so that, with positive probability, t' is either the median of the announcements or the highest announcement (if $t' > b^*$) and, correspondingly, the truthful announcement t would induce an outcome strictly closer to agent j's type. That is, no $t' \in (a^*, 1]$ maximizes agent j's interim utility when her peak is in $[0, a^*) \setminus \tilde{\Phi}_j^\sigma$, since $\mathscr{C}(G_j) \subset [a^*, 1]$, we must have $\sigma_j(t) = a^*$ for almost every $t \in [0, a^*) \setminus \tilde{\Phi}_j^\sigma$.

For the same reason as agent j, any other agent $j' \neq j$ with peak in $[0, a^*) \setminus \tilde{\Phi}_{j'}^\sigma$ strictly prefers a truthful announcement to any announcement in $(a^*, 1]$. But, in turn, due to Claim 1, agent $j' \neq j$ announces a^* with probability zero. Hence, there exists some $\xi > 0$ such that $\mathscr{C}(G_{j'}) \cap [0, a^* - \xi) \neq \emptyset$. That is, every agent $j' \neq j$ announces some value smaller than $a^* - \xi$ with

[7] If a subset of $\mathscr{I}$ is totally ordered, then either $I \succeq J$, $J \succeq I$ or both. That means that one can express a totally ordered subset as a sequence of intervals $I_1, I_2, \ldots, I_n$ with $I_i \subseteq I_{i+1}$ for each $i = 1, \ldots, n$. Since each I_i satisfies $I_i \subseteq (0, 1)$, the subset admits an upper-bound.

positive probability and some value larger than a^* also with positive probability. Therefore, the unique maximizer of agent j's expected utility conditional on her type being in the open set $(a^* - \xi, a^*) \setminus \tilde{\Phi}_j^\sigma$ is a truthful announcement. This contradicts $\mathscr{C}(G_j) \subset [a^*, 1]$, therefore, $a^* = 0$. For an analogous reason, $b^* = 1$.

Thus we conclude that σ prescribes, for every agent, a truthful announcement for almost every peak in $[0,1]$ as we wanted. Moreover, every announcement is unique with probability one according to agents' beliefs. That is, θ_{VB} robustly and truthfully implements the median rule. □

Notice that the only property of beliefs that we used in the previous proof is that agents agree on which subsets of the peak profile receive positive probability. This holds if for every $i = 1, \ldots, n$, agent i's prior belief is absolutely continuous with respect to the Lebesgue measure. Therefore, we have actually proven the following theorem.

Theorem 2. *The value-based median mechanism θ_{VB} robustly and truthfully implements the median rule on the domain of smooth priors D^S.*

To our knowledge, the above result is the first one showing that the median rule is implementable in equilibrium 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 social choice rules in the single-peaked domain.

Remark 2. Notice that a similar result can be obtained if one uses the notion of interim Bayes-Nash equilibrium instead. Indeed, it can be easily seen that truth-telling is an interim Bayes-Nash equilibrium. Moreover, since every interim Bayes-Nash equilibrium must also be an ex-ante Bayes-Nash equilibrium, it trivially follows that the value-based median robustly and truthfully implements the median rule also if the equilibrium concept is interim Bayes-Nash equilibrium.

Remark 3. In contrast with the situation in the complete information environment (see Remark 1), discontinuity of θ_{VB} is a necessary property of the mechanism behind the implementation result in Theorem 2. Indeed, from Proposition 3 we know that a mechanism that robustly and truthfully implements the median rule must differ from the median in at least one peak profile. But if the mechanism is continuous, this actually implies that it must differ from the median in an open set of peak profiles, and if that is the case, such a continuous mechanism cannot robustly and truthfully implement the median rule on D^I.[8]

6. Generalized Median Rules

A GMR selects the median of a profile of peaks *jointly* 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. Similarly, any k-th order statistic of a peak profile (which selects the k-th largest peak) can be expressed in a similar way.

[8] This observation is in line with papers such as Kim [2017] and Ehlers et al. [2020] that show how discontinuous mechanisms can help expand the space of implementable rules in some environments.

Before formally defining GMRs, we enlarge the sets A and $\mathbb{A}$ to incorporate 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 $\mathbb{A}^{\varepsilon} := \bigcup_{k=1}^{\infty} [-\varepsilon, 1+\varepsilon]^k$. 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 social choice rule $f_y : T^n \to A$ 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. So far we have assumed that the number of agents is odd. Using phantoms, we can easily extend the median to the case in which the number of agents is even. If n is even, then $y = (-\varepsilon)$ 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 $y = (1 + \varepsilon)$ 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 generates 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 equal announcements as one.

Definition 9. *The generalized value-based median mechanism associated to the permissible phantom profile y is the direct mechanism $\theta_y : T^n \to A^{\varepsilon}$ in which 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 truthful 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, 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 strictly 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 proof of Theorem 1, such an outcome coincides 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 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. □

We also obtain an implementation result for GMRs under the incomplete information environment described in Section 5. 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' strategy consists of truthfully reporting their peak as well. Hence, truthful revelation is an equilibrium of the mechanism. The proof of the following result is very similar to that of Proposition 4 and is relegated to the appendix.

Theorem 4. *The mechanism θ_y robustly and truthfully implements the GMR f_y on the domain D^S.*

Remark 4. For a general *GMR*, the only reason why an agent may find it advantageous to misreport her peak is that a phantom coincides with her peak, something that is impossible if f_y selects the k-th order statistic since all phantoms are either smaller than 0 or larger than 1. Indeed, if no phantom is in the interval $[0,1]$, as it is the case for the social optima in examples 4 and 5, then sincerity is the *unique* interim Bayes-Nash equilibrium. Under such an equilibrium, the median peak is implemented with probability one according to any agent's prior.

Appendix A. Discrete policy space

A.1. *Complete information*

When the number of distinct announcements is even, the value-based median mechanism selects the midpoint between the lower and the upper median. If the policy space is discrete, the midpoint between two alternatives need not belong to the policy space. Hence, we need modify the value-based median mechanism so that, when the number of distinct announcements is even, it randomizes between the lower and the upper median.

In this Section we assume that the set of alternatives A and the set of peaks T coincide with the set of integers $\mathbb{Z}$.[9] For each pair of alternatives $x,y \in A$, the lottery that selects x with probability 1/2 and y with probability 1/2 is denoted $\ell(x,y)$. We write $\Delta(A)$ for the set of alternatives jointly with the set of lotteries between pairs of alternatives in A. Furthermore, an agent with type t_i evaluates lotteries through expected utility so that

[9] We require the policy space to be unbounded because if the policy space is bounded and the median is sufficiently close to one of the most extreme values, agents may not have enough distinct values at one of the sides of the median to be able to induce it as an equilibrium outcome using the random value-based mechanism. It seems that allowing for a finite number of alternatives would require indirect mechanisms to implement the median.

$U_i(\ell(x,y) \mid t_i) = \frac{1}{2}U_i(x \mid t_i) + \frac{1}{2}U_i(y \mid t_i)$ (by slight abuse of notation we use U to denote both utility and expected utility). Single-peakedness of preferences implies that if $x < y \le t_i$ or $t_i \le y < x$, then $U_i(y \mid t_i) > U_i(\ell(x,y) \mid t_i) > U_i(x \mid t_i)$.

The *random value-based mechanism* endows each agent i with message space T so that the set of message profiles is T^n. The random value-based median mechanism $\theta_{RVB} : T^n \to \Delta(A)$ is given by

$$\theta_{RVB}(s) := \begin{cases} m(v(s)) & \text{if } v(s) \text{ is odd,} \\ \ell\left(m^-(v(s)), m^+(v(s))\right) & \text{if } v(s) \text{ is even.} \end{cases} \tag{5}$$

That is, when the number of unique announcements is even, the outcome is the median of the unique values of the announcements. When the number of unique announcements is even, the outcome is the fair lottery between the two medians.[10]

The same arguments as those used in Section 4 apply to the current setting. Indeed, consider an equilibrium $s \in T$ and an agent with peak t_i. If $t_i \neq \theta_{RVB}(s)$, agent i's best response consists of announcing an alternative s_i which is different from any other announcement s_j, since otherwise her announcement will be ignored. Furthermore, her best response is located to the same side of the equilibrium as her peak, so that if $t_i < \theta_{RVB}(s)$, then $s_i < \theta_{RVB}(s)$ and if $t_i > \theta_{RVB}(s)$, then $s_i > \theta_{RVB}(s)$. This implies that the statement of Proposition 2 also applies to the random value-based median.

This, in turn, implies the following result where Nash implementation is an adaptation of the definition of implementation to random mechanisms.

Theorem 5. *The random value-based median mechanism θ_{RVB} Nash implements the median rule.*

This theorem shows that Theorem 1 does not crucially depend on having a continuous policy space. The key ingredient is that each agent is always able of submitting a unique announcement (i.e. different from the rest of announcements) while still being able to affect the outcome, independently of the rest of the votes. Note that the equilibrium outcome is deterministic (some alternative in A) whereas off-equilibrium lotteries may arise.

A.2. *Incomplete information*

We now discuss Bayesian implementation with a discrete set of alternatives and peaks. For simplicity, we do not provide definitions of the main objects since there are the obvious adaptations of the ones in the continuum setting. As usual in the literature, there is a common knowledge prior $P(\cdot)$ over the set T^n. Therefore, the marginal distribution $p_i(\cdot)$ of each agent i is derived from p according to $p_i(t_i, t_{-i}) = \sum_{t_{-i} \in T^{n-1}} P(t_i, t_{-i})$ for all $t_i \in T$.

The interim expected utility of an agent i with type t_i when f is the social choice rule is denoted by:

$$V_i(f \mid t_i) = \sum_{t_{-i} \in T^{n-1}} U_i(f(t_i, t_{-i}) \mid t_i) p(t_{-i} \mid t_i),$$

where $U_i : A \times T \to \mathbb{R}$ is the usual utility function.

[10] Note that this altered mechanism also implements the median rule when the policy space is a continuous interval.

Within this setting, the literature on Bayesian implementation has focused on the notion of Bayesian Monotonicity (see Palfrey and Srivastava [1989] and Jackson [1991] for classic treatments, see also Korpela [2014] for a recent discussion). If a social choice rule is implementable in interim Bayes-Nash equilibrium, then it satisfies Bayesian Monotonicity. The definition of Bayesian Monotonicity hinges on the notion of deception. A deception of agent i is any function $\alpha_i : T \to T$. The set of all possible deceptions of agent i is denoted by D_i, and any profile $\alpha \in \prod_{i=1}^n \alpha_i$ of agent deceptions is called simply a deception that belongs to D, the set of deception profiles. The identity mapping $id : T \to T$ is denoted by $\hat{\alpha}_i$ and $\hat{\alpha} = (\hat{\alpha}_i)_{i \in N}$ denotes the identity profile.

Definition 10. *A social choice rule $f : T^n \to A$ is Bayesian Monotonic if for all $\alpha \in D$, if $f \equiv f(\alpha)$ does not hold, then there exists $i \in N$, $t_i \in T$ and a social choice rule $g : T^n \to A$ such that:*

(i): $V_i(g(\alpha) \mid t_i) > V_i(f(\alpha) \mid t_i)$, *and*
(ii): $V_i(f(\hat{\alpha}) \mid t_i') \geq V_i(g(\alpha_i(t_i), \hat{\alpha}_{-i}) \mid t_i')$ *for any* $t_i' \in T$.

For the sake of completeness, we present an example, adapted from Triossi [2005], that proves that the median rule is not Bayesian Monotonic for some open set of priors, which underlines our case for the value-based median mechanisms: the median rule is not Bayesian implementable with a finite set of alternatives. While the example uses the median rule, the same result applies for GMRs. The result extends to mechanisms that allow for simple lotteries between two adjacent alternatives such as the random value-based mechanism of Section 7. The argument is done for three alternatives for simplicity but can be extended to a setting with an unbounded and discrete number of alternatives.

Example 6. Let $n = 3$ and $A = \{x, y, z\}$ with $0 < x < y < z < 1$. There are three possible types, so that $T = \{x, y, z\}$, with for each $i = 1, 2, 3$, the utility function $U_i : A \times T \to \mathbb{R}$ such that:

$$U_i(x \mid x) = 2,\ U_i(y \mid x) = 1,\ U_i(z \mid x) = 0,$$
$$U_i(x \mid y) = 0,\ U_i(y \mid y) = 2,\ U_i(z \mid y) = 1,$$
$$U_i(x \mid z) = 0,\ U_i(y \mid z) = 1,\ U_i(z \mid z) = 2.$$

Agent's preferences are single-peaked for each possible type. The prior distribution $p : T^n \to \mathbb{R}$ is such that for each triple $(a, b, c) \in T^n$, $p(a, b, c) = p_a p_b p_c$ with $p_x + p_y + p_z = 1$ and $p_y^2 > 2/3$.

Assume by contradiction that the median rule f_M is Bayesian Monotonic. We denote by $f_M(\alpha)$ the outcome associated by the median rule f_M to the constant deception α, with $\alpha = (\alpha_i)_{i \in N}$ and $\alpha_i(t) = y$ for any $t \in T$ and for any $i = 1, 2, 3$. Then, the outcome $f_M(\alpha)$ equals y since for each $(a, b, c) \in T^n$, $\alpha(a, b, c) = (y, y, y)$ and $f_M(y, y, y) = y$. Moreover, note that $f_M(x, x, y) = x$ so that there is some state in which $f_M \neq f_M(\alpha)$. Thus, the definition of Bayesian Monotonicity implies that, for some $t \in T$ and some SCR g, **(i)** $V_i(g(\alpha) \mid t) > V_i(f_M(\alpha) \mid t)$, and **(ii)** $V_i(f_M(\hat{\alpha}) \mid t') \geq V_i(g(\alpha_i(t), \hat{\alpha}_{-i}) \mid t')$ for any $t' \in T$.

Since the expected utility of an agent with type y is maximal at $f_M(\alpha)$ (since she obtains her most preferred alternative with certainty), an agent with type y cannot satisfy **(i)**. Then, for **(i)** to hold, we need $t \in \{x, z\}$ and $g(\alpha) \in \{x, z\}$.

Assume first that $g(\alpha) = x$. Consider an agent with type z. It follows that

$$V_i(g(\alpha) \mid z) = 0 < 1 = V_i(f_M(\alpha) \mid z),$$

so that (**i**) is violated when $t = z$. Thus, for (**i**) to hold, the agent's type satisfies $t = x$. Consider then an agent with type x that reports $\alpha_i(x) = y$ and assume that her opponents report truthfullly. If her two opponents have type y, the agent's payoff under g equals 2, since $g(\alpha_i(x), y, y) = x$. This occurs with probability p_y^2. If her two opponents do not have type y, the agent's payoff under g is at least 0, since 0 is her worst utility payoff. Thus, the expected utility $V_i(g(\alpha(x), \hat{\alpha}_{-i}) \mid x)$ satisfies:

$$V_i(g(\alpha(x)), \hat{\alpha}_{-i}) \mid x) \geq p_y^2(2) + (1 - p_y^2)(0). \tag{6}$$

Similarly, consider an agent with type x that reports truthfully and assume that her opponents report truthfully. If her two opponents have type y (which occurs with probability p_y^2), the median is y so that the agent's payoff equals 1. If her two opponents do not have type y (with probability $1 - p_y^2$), the agent's payoff is at most 2. Thus, the expected utility $V_i(f_M(\hat{\alpha}) \mid x)$ satisfies:

$$V_i(f_M(\hat{\alpha}) \mid x) \leq p_y^2(1) + (1 - p_y^2)(2).$$

Moreover, $p_y^2 > 2/3$ implies that $p_y^2(1) + (1 - p_y^2)(2) < p_y^2(2) + (1 - p_y^2)(0)$ which leads to:

$$V_i(f_M(\hat{\alpha}) \mid x) < p_y^2(2) + (1 - p_y^2)(0). \tag{7}$$

Yet, equations (6) and (7) together imply that $V_i(f_M(\hat{\alpha}) \mid x) < V_i(g(\alpha(x), \hat{\alpha}_{-i}) \mid x)$ which contradicts (**ii**) , and concludes the argument when $g(\alpha) = x$.

Now, assume that $g(\alpha) = z$. Again, for (**i**) to hold, $t = z$, as otherwise there is a contradiction. By a similar logic to the case with $g(\alpha) = x$, one can show that

$$V_i(f_M(\hat{\alpha}) \mid z) \leq p_y^2(1) + (1 - p_y^2)(2) < p_y^2(2) + (1 - p_y^2)(0) \leq V_i(g(\alpha(z), \hat{\alpha}_{-i}) \mid z),$$

where the second inequality holds for $p_y^2 > 2/3$. The previous inequality contradicts (**ii**) . Therefore, we have shown that the median rule is not Bayesian Monotonic and hence is not interim Bayes-Nash implementable.

APPENDIX B. PROOF OF THEOREM 4

The proof of Theorem 4 is based on the proof of Proposition 4. The only difference is that we need to take into account the fact that some agents can have peaks that coincide with the phantoms introduced by *GMR*s. As in the proof of Proposition 4 and to simplify notation, we prove the result using the domain of priors D^I, but the same arguments go through for the more general domain of smooth priors D^S.

Proof of Theorem 4. Let h be the length of the permissible profile of phantoms y. We argue that every Bayes-Nash equilibrium σ satisfies $(\sigma_1(t_1), \ldots, \sigma_n(t_n)) = (t_1, \ldots, t_n)$ for almost every peak profile $t \in T^n$.

Therefore, let σ be a Bayes-Nash equilibrium and let Φ_i^σ be the set of points $x \in [0,1]$ such that

$$G_j(t) - \lim_{\varepsilon \to 0^+} G_j(t-\varepsilon) > 0 \text{ for some } j \neq i \tag{8}$$

Define $\hat{\Phi}^\sigma := \Phi_i^\sigma \cup \{y_1\} \cup \cdots \cup \{y_h\}$ and

$$\tilde{\Phi}_i^\sigma := \hat{\Phi}_i^\sigma \bigcup \left\{x \in A \mid x = \delta(z,w) \text{ for some } z,w \in \hat{\Phi}_i^\sigma\right\}. \tag{9}$$

Claim 2. No agent i sends an announcement in $\hat{\Phi}_i^\sigma$ apart from (possibly) a zero measure set of her peaks.

Proof of the claim. To the contrary, assume that the set of peaks $\sigma_i^{-1}(x)$ for which agent i announces $x \in \hat{\Phi}_i^\sigma$ has positive measure. Then the set $\sigma_i^{-1}(x) \setminus \tilde{\Phi}_i^\sigma$ also has positive measure. Using the same arguments as in the proof of Claim 1 we know that the expected payoff of any peak $t \in \sigma_i^{-1}(x) \setminus \tilde{\Phi}_i^\sigma$ is strictly larger if she announces t than if she announces x. But, since $\sigma_i^{-1}(x)$ has positive measure, this contradicts the fact that σ_i maximizes the ex-ante payoff of agent i, which is impossible. Therefore, $\sigma_i^{-1}(x)$ has measure zero and no agent i sends an announcement in Φ_i^σ apart from (possibly) a zero measure set of her peaks. □

Therefore, we can choose a collection of sets $\{Q_i\}_{i=1}^{n+h}$ such that

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

As in the proof of Theorem 2, we use the product set Q to find an interval of peaks such than some agent (and then every agent) plays truthfully for almost every peak in that interval. However, now there are two cases: either the median $m(q)$ coincides with a phantom or it does not. Take some $q \in Q$, if the median $m(q)$ is not a phantom, then $m(q) \in Q_k$ for some agent k and we can select the interval $(a_k,b_k) \cap [0,1]$ where (a_k,b_k) is defined as in the proof of Theorem 2 (we take the intersection with $[0,1]$ to guarantee that the selected interval is a subset of the set of peaks). If the median $m(q)$ is a phantom then let agent k be agent 1. If every element in Q_1 is smaller than $m(q)$ then define

$$a_k := m(q) \text{ and } b_k := \min\left\{x \in \bigcup_i Q_i : x > m(q)\right\}.$$

If otherwise every element in Q_1 is larger than $m(q)$, let

$$a_k := \max\left\{x \in \bigcup_i Q_i : x < m(q)\right\} \text{ and } b_k := m(q).$$

Consider the interval (a_k,b_k). Using the same arguments as in Theorem 2, we know that agent k's maximizing strategy must prescribe truth-telling for almost every peak in (a_k,b_k). This conclusion extends to every other agent. From this point forward the proof of Theorem 2 applies almost verbatim. □

References

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.

R.P. Arribillaga, J. Massó, and A. Neme. On obvious strategy-proofness and single-peakedness. *Journal of Economic Theory*, pages 1049–92, 2020.

G. Attiyeh, R. Franciosi, and R. M. Isaac. Experiments with the pivot process for providing public goods. *Public choice*, 102(1-2):93–112, 2000.

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. Benassy. *The Econonomics of Market Disequilibrium*. 1982.

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.

O. Bochet. Nash implementation with lottery mechanisms. *Social Choice and Welfare*, 28(1): 111–125, 2007.

O. Bochet, T. Sakai, and W. Thomson. Preference manipulations lead to the uniform rule. 2021.

T.N. Cason, T. Saijo, T. Sjostrom, 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.

S. Ching. An alternative characterization of the uniform rule. *Social Choice and Welfare*, 11 (2):131–136, 1994.

B. Dutta and A. Sen. Nash Implementation with Partially Honest Individuals. *Games and Economic Behavior*, 74:154–169, 2012.

L. Ehlers, D. Majumdar, D. Mishra, and A. Sen. Continuity and incentive compatibility in cardinal mechanisms. *Journal of Mathematical Economics*, 88:31–41, 2020.

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.

L. Hurwicz and D. Schmeidler. Construction of Outcome Functions guaranteeing Existence and Pareto Optimality of Nash Equilibria. *Econometrica*, 46:1447–1474, 1978.

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

J. H. Kagel and D. Levin. Independent private value auctions: Bidder behaviour in first-, second-and third-price auctions with varying numbers of bidders. *The Economic Journal*, 103(419):868–879, 1993.

J. H. Kagel, R.M. Harstad, and D. Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55:1275–1304, 1987.

T. Kawagoe and T. Mori. Can the pivotal mechanism induce truth-telling? An experimental study. *Public Choice*, 108(3-4):331–354, 2001.

S. Kim. Ordinal versus cardinal voting rules: A mechanism design approach. *Games and Economic Behavior*, 104:350–371, 2017.

J.-F. Laslier, M. Núñez, and M.R. Sanver. A solution to the two-person implementation problem. *Journal of Economic Theory*, 194:105261, 2021.

S. Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287, 2017.

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

M.R. Sanver. Nash implementing non-monotonic social choice rules by awards. *Economic Theory*, 28(2):453–460, 2006.

A. Saporiti. Securely implementable social choice rules with partially honest agents. *Journal of Economic Theory*, 154:216–228, 2014.

Y. Sprumont. The division problem with single-peaked preferences: a characterization of the uniform allocation rule. *Econometrica*, 59:509–519, 1991.

Y. Sprumont. Strategy-proof collective choice in economic and political environments. *Canadian Journal of Economics*, 28:68–107, 1995.

M. Triossi. Pure public goods and incomplete information: Multiplicity and implementation. 2005.

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.