Information Disclosure under Strategy-proof Social Choice Functions

Salvador Barberà
Antonio Nicolò

This version: March 2019
(May 2016)

Barcelona GSE Working Paper Series

Working Paper nº 904

Information Disclosure under Strategy-proof Social Choice Functions*

Salvador Barberà[†]

MOVE, Universitat Autònoma de Barcelona and Barcelona GSE

Antonio Nicolò[‡]

University of Padova and University of Manchester

March 22, 2019

Abstract

We consider collective decision problems where some agents have private information about alternatives and others don't. Voting takes place under strategy-proof rules. Prior to voting, informed agents may or may not disclose their private information, thus eventually influencing the preferences of those initially uninformed. We provide general conditions on the voting rules guaranteeing that informed agents will always be induced to disclose what they know. In particular, we apply our general result to environments where agents' preferences are restricted to be single-peaked, and

*We thank for their comments and suggestions, Jordi Massò, Francois Maniquet, Hervé Moulin, Matthew O. Jackson, Wonki Jo, and Navin Kartik; S. Barberà acknowledges financial support through grants ECO2014-53051-P and SGR2014-515, and Severo Ochoa Programme for Centers of Excellence in R&D (SEV-2015-0563).

†salvador.barbera@uab.cat. Universitat Autònoma de Barcelona; Facultat d'Economia i Empresa; Edifici B; 08193 Bellaterra (BARCELONA); Spain.

‡corresponding author; antonio.nicolo@unipd.it. Department of Economics, University of Padua, via del Santo 33, 35124 Padova - Italy.

characterize the strategy-proof rules that ensure information disclosure in this setting.

Keywords: strategy-proofness, information disclosure, voting rules, single-peaked preferences.

JEL Codes: D70, D71, D82.

1 Introduction

A set of agents must choose by vote their collective course of action. Their preferences regarding potential outcomes will typically differ. Moreover, the perception of the same outcome by each voter may depend on what she knows and eventually learns about it. If information that is relevant to the appraisal of an outcome by one agent is in the hands of another, the latter can condition the vote of the first, by inducing a change in her preferences. Examples abound. In some cases, the set of better informed agents is predictable. Members of a House committee will know more than other representatives about what happened behind closed doors. Faculty in a hiring committee of an academic institution get better chances to learn about candidates than other colleagues. Yet, in other cases, who becomes better informed may depend on varying factors, and it may be impossible to know a priori who will become an expert. In either case, it is important to understand under what circumstances will informed members of society be inclined to share their knowledge with others.

To analyze the problem of strategic disclosure of information, Milgrom and Roberts (1986) introduced the class of persuasion games, in which all better informed agents share the same knowledge, can withhold information but not lie, and players interact only once, so that issues of reputation do not arise. Theirs and subsequent papers (Lipman and Seppi, 1995; Dewatripont and Tirole, (1999); Bhattacharya and Mukherjee, 2013; Gentzkow and Kamenica, 2017) have studied the case where a single decision maker must elicit the private information she receives from competing experts. By contrast, we consider settings where the decision maker is a set of individuals, some of them experts and others not, who must jointly adopt decisions, and study which voting rules provide experts with incentives to disclose their private information to other members of society. The question has been addressed by Schulte (2010) and

Jackson and Tan (2013) for the special situation where the group has only two choices. We tackle the general case where society faces any number of alternatives, focusing on the choice of strategy-proof social choice voting rules.

The fact that voting rules shape deliberation has been well-noted in important papers by Austen-Smith and Feddersen (2005, 2006), building on insights from earlier work by Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998). Our work contrasts with theirs in that we study persuasion games rather than cheap-talk games, and, more importantly, we consider the general class of strategy-proof social choice functions, rather concentrating on quota rules alone. This extension becomes substantial when society can choose among more than two alternatives.

In our setting all experts are fully informed, and therefore a trivial but satisfactory way to attain full disclosure is to grant dictatorial power to one of them. Yet, our remark hints at the fact that providing incentives to disclose information may come at the price of granting experts some special privileges. Following that hint, we provide a sufficient condition on strategy-proof social choice functions guaranteeing that in societies using such rules the outcome of voting will always be the same that the one that would obtain if all agents were fully informed at voting time. This condition requires that experts should be collectively granted what we call coalitional veto power, in the following sense: if the social choice function selects two different outcomes at two profiles that only differ in the preferences of non-experts, there must exist two experts who rank these outcomes in opposite directions, according to their unchanged preferences. In other words: if a rule attributes coalitional veto power to experts, changes in the outcome driven by changes in non experts' preferences cannot harm all the experts at a time. In addition to its sufficiency, we also prove that the condition becomes necessary in those cases where the impact of information disclosure on non experts' preferences is ex-ante unpredictable. These results are related to classical models in political economy depicting strategic information transmission as a mutually beneficial exchange, where a legislative committee member (Gilligan and Krehbiel, 1989, 1990; Krishna and Morgan, 2001) or a lobbyist (Austen-Smith and Wright, 1992) offers helpful policy expertise in exchange for some influence over the outcome. They also confirm the intuition that emerges in several previous works that creating competition between experts is beneficial to disclosure.

After our general results we present an application to the case in which agents have single-peaked preferences, and consider two different scenarios. We first characterize the class of strategy-proof and unanimous social choice functions that always induce information disclosure under the assumption that they can be defined after observing what set of agents will be the experts. After that, we consider the scenario in which the identity of experts is still unknown at the time when the voting rule must be chosen. In this case, our general results imply that no unanimous strategy-proof rule can ensure that all agents will be ex-post fully informed. But then we can prove that majority voting is a second best solution to the disclosure problem in a precise way.

The paper proceeds as follows. In Section 2 we introduce a motivating example and in Section 3 the model. Section 4 presents general results regarding necessary and sufficient conditions for information disclosure. Sections 5 applies these results to the single-peaked preference domain. All proofs of propositions are in the Appendix.

2 A Motivating Example

Five agents (1,2,3,4,5) face the choice between six ordered alternatives $(a_1, a_2, a_3, a_4, a_5, a_6)$. All agents have single peaked preferences relative to that order. Agents 1 and

2 are informed agents (experts), have their peaks in a_3 and a_6 respectively, and both have one piece of information. If no expert discloses that information, the peak of all uninformed agents is a_2, and change to a_1 if at least one of the experts discloses. If society decides according to the median rule, the outcome under no disclosure is a_2, and changes to a_1 in case information is disclosed. Clearly, no expert would want to disclose.

Next, consider the modified rule under which the median is selected as long as it lays between the peak alternatives of experts, and otherwise selects the peak of the expert that is closest to the median. The outcome of the rule in the case we considered is a_3 both before and after disclosure, and disclosure is acceptable for both experts. This is achieved due to the power of experts to veto those alternatives that do not lie between their tops. Moreover, notice that if the peak of uninformed voters 3 and 4 is a_5 under no information, and goes to a_4 after disclosure, while the top of voter 5 is at a_2 and remains the same after information disclosure, the alternative chosen by the modified rule is a_5 if information is not disclosed and a_4 if it is disclosed. The modified rule creates a conflict of interests between experts 1 and 2 leading 1 to disclose.

The modified rule is strategy-proof (every agent has incentive to truthfully report his preferences irrespective of what the other agents report), and unanimous, (if a peak is the preferred one by every agent, then it is the selected alternative). It induces information disclosure thanks to a combination of two ingredients: endowing experts with some veto power against the uninformed, yet creating internal conflict among them. We now extend the ideas suggested by this example to our general framework.

3 The model

A finite set of agents $A = \{1, \ldots, n\}$ faces a set X of two or more alternatives and must choose one of them. Let $\tilde{\mathcal{R}}$ be the set of all complete, reflexive, and transitive binary relations on X and $\mathcal{R}_i \subseteq \tilde{\mathcal{R}}$ be the set of those preferences that are allowed for agent i. $R_i \in \mathcal{R}_i$ will denote agent $i's$ preferences and $R \equiv R_1 \times \ldots R_n$ a preference profile. Let P_i be the strict part of R_i.

$\mathcal{I}$ is a finite set of elementary pieces of verifiable hard information and $I \in 2^{\mathcal{I}}$ a generic subset of information. Different information may be available to different agents, and the preferences of each one depend on the information he or she holds. We formalize this dependence through the notion of an agent's type. Agent $i's$ type $\theta_i : 2^{\mathcal{I}} \to \mathcal{R}_i$ is a function which assigns a preference $\theta_i(I) \in \mathcal{R}_i$ to agent i for each set of information I that she is aware of; Θ_i denotes the set of types for agent i and $\theta \in \Theta_1 \times \ldots \times \Theta_n \equiv \Theta$ stands for a full profile of types.

The set of agents is partitioned into a set of experts E and a set of non-experts N, with $N = A \backslash E$. Without loss of generality let $E = \{1, \ldots, l\}$ and $N = \{l+1, \ldots, n\}$.

We denote $\theta_E \in \Theta_1 \times \ldots \times \Theta_l$ a profile of types for experts and $\theta_N \in \Theta_{l+1} \times \ldots \times \Theta_n$ a profile for non-experts.

We assume that every expert knows the full set of information $\mathcal{I}$, and this information is not available to the rest of agents.[1]

The collective decision process involves two stages. In the first stage, experts decide what information $I \in 2^{\mathcal{I}}$ they want to disclose and do it publicly. In the second stage, all agents vote according to their best interest, given the information they hold.

[1]This assumption is not too restrictive. We can assume otherwise that non-experts may have some information but lack the technology to disclose it effectively, and adjust the type of that agent to only react when she acquires additional information.

Non-experts' preferences depend on the overall amount of information disclosed at the first stage, and not on the identity of who has disclosed the information.

Formally, m_i is a message of expert $i \in E$, $M \equiv 2^{\mathcal{I}}$ stands for the set of messages available to every $i \in E$. Profiles of messages are denoted by $m = (m_1, \dots, m_l) \in M_E = M^l$; disclosure decisions by experts are described by disclosure functions $g : M_E \to 2^S$ and $g(m) = \bigcup_{i \in E} m_i$ is the amount of information disclosed if message profile m is chosen.

Once the first stage is terminated and information $g(m)$ has been publicly disclosed, each non-expert $j \in N$ is endowed with a preference $\theta_j(g(m)) \in \mathcal{R}_j$ and each expert $i \in E$ with preference $\theta_i(\mathcal{I}) \in \mathcal{R}_i$. The type function encompasses any (Bayesian) updating inference process that a non-expert may carry out when observing what information has been disclosed.

We now formalize how decisions are reached in the second stage.

Definition 1 *A social choice function $f : \mathcal{R}_1 \times \ldots \times \mathcal{R}_n \to X$ is a map from each preference profile reported by agents to one alternative.*

We concentrate on strategy-proof social choice functions.

Definition 2 *Given a social choice function $f : \mathcal{R}_1 \times \ldots \times \mathcal{R}_n \to X$, we say that agent $i \in N$ can manipulate at profile R via R_i', if $f(R_{-i}, R_i') P_i f(R_{-i}, R_i)$. A social choice function $f : \mathcal{R}_1 \times \ldots \times \mathcal{R}_n \to X$ is strategy-proof if no agent can manipulate f at any profile R.*

Our interest in strategy-proof rules is normative, but admittedly also pragmatic. Indeed, it allows us to eliminate strategic behavior from the second stage of the game, and thus concentrate all attention on the only strategic consideration left, that of information disclosure. We shall also assume that agents operating under these rules are partially honest (see for instance Dutta and Sen, 2011; Kartik, Tercieux, and Holdeny, 2014): if truthfully reporting their preferences is a weakly dominant strategy, then agents declare their true preferences.

Given a type-profile θ, to each message profile m is associated a preference profile $(\theta_E(\mathcal{I}), \theta_N(g(m)))$. Since the voting rule is strategy-proof and agents are partially honest, it follows that to each m is associated a unique alternative $x = f(\theta_E(\mathcal{I}), \theta_N(g(m)))$.

Given a type-profile $\theta \in \Theta$, a θ-game is the corresponding simultaneous move game $(E, \theta_E(\mathcal{I}), M_E, f(\theta_E(g(\mathcal{I}), \theta_N(g(\cdot)))$. A Nash equilibrium of a θ-game is a profile of messages m^* such that

$$u_i\left(f(\theta_E(\mathcal{I}), \theta_N(g(m^*)))\right) \geq u_i\left(f(\theta_E(\mathcal{I}), \theta_N(g(m_i, m^*_{-i})))\right),$$

for all $i \in E$, $m_i \in M_i$.

We model the choice of social choice functions as if a designer had to select one of them, and was interested in methods that would achieve the social decision corresponding to societies where all agents were fully informed. We assume that, at the time of the decision, the designer knows the set of admissible preferences $\mathcal{R}_i$ for each agent i, but she does not necessarily observe the sets $\mathcal{I}$ and Θ_i, or which message profile m will be played by the experts.

The following definition describes the first best objective the designer may try to attain.

Definition 3 *A strategy-proof voting rule f ensures full information disclosure if for every $\theta \in \Theta$ and for every Nash equilibrium m^* of the corresponding θ-game, $f(\theta_E(\mathcal{I}), \theta_N(g(m^*)) = f(\theta(\mathcal{I}))$.*

Full information disclosure requires that every equilibrium outcome coincides with the outcome of the social choice function when all agents have fully informed preferences. Hence, it may be the case that there are equilibria in which not all information is transmitted, because the missing pieces are irrelevant for the outcome.

Even if full disclosure could not be guaranteed, one might be able to evaluate what social choice functions would perform better in promoting disclosure. As we shall see, the need for a second best approach arises when the designer does not know who are the experts at the time of selecting a rule.

We model the designer's lack of information about what agents will be experts by assuming that nature draws their set E out of a family $\mathcal{W}$ (with cardinality larger than one) of non-empty coalitions of A. The designer knows $\mathcal{W}$, but it must choose the social choice function f before Nature's choice. In that context, we propose the following terms of comparison

Definition 4 *The strategy-proof voting rule f ensures **better information disclosure** than the strategy-proof voting rule f', if (i) for every $M \in \mathcal{W}$ such that f' ensures full information disclosure when M is the set of experts, then f also ensures it and (ii) there is some $T \in \mathcal{W}$ such that f ensures full information disclosure when T is the set of experts, while f' does not.*

This definition provides a partial ordering over social choice functions with respect to the incentives they provide to disclose information, and may be used to compare the performance of social choice rules in settings where it is impossible to achieve the goal of full information disclosure. We shall use it in Section 5.

4 Main Results

We first provide a straightforward necessary and sufficient condition for a strategy-proof social choice function to ensure full information disclosure.

Proposition 1 *A strategy-proof social choice function* $f : \mathcal{R}_1 \times \ldots \times \mathcal{R}_n \rightarrow X$ *ensures full information disclosure if and only if for every* $\theta = (\theta_E, \theta_N) \in \Theta$ *and every* m^* *either* $f(\theta_E(\mathcal{I}), \theta_N(g(m^*)) = f(R(\mathcal{I}))$*, or there exist* $I \subseteq \mathcal{I}$*,* $i \in E$ *and* $\theta_i \in \Theta_i$*, such that* $I \supset g(m^*)$ *and* $f(\theta_E(\mathcal{I}), \theta_N(I))P_i f(\theta_E(\mathcal{I}), \theta_N(g(m^*)))$*, where* $P_i = \theta_i(\mathcal{I})$*.*

The above condition is simply a restatement of our desideratum, because it implies that, if information is not fully disclosed, there must be an expert who has incentive to disclose more information. This condition resembles that of Proposition 4 in Milgrom and Roberts (1986), but it is worth stressing that our requirement is different: we do not require experts to prefer the full information decision to any other, but only to be interested in contributing some additional information, not necessarily all of it.

Admittedly, it is hard to check whether this condition holds, because it requires to know the set of types Θ_i for each i, which ultimately means to know how every piece of information affects individual preferences. But it opens the door to express a less obvious, easy to check sufficient condition, which only requires to know the set of of individual preferences of every agent. This condition is based on the power attributed to experts by the social choice function.

Definition 5 *A social choice function* $f : \mathcal{R}_1 \times \ldots \times \mathcal{R}_n \rightarrow X$ *attributes coalitional veto power to a set* $M \subseteq N$*, if for all pairs* $R, R' \in \mathcal{R}$ *with* $R_i = R'_i$ *for all* $i \in M$*, either* $f(R) = f(R')$ *or there is a pair* $i, j \in M$ *such that* $f(R')P_i f(R)$ *and* $f(R)P_j f(R')$*.*

When a set M is endowed with coalitional veto power, changes in the social outcome that are driven by changes in the preferences of agents outside M cannot harm all agents in M at the same time. Hence, agents in M collectively have the power "to veto" changes in the social outcome that they all dislike. We can now state the following.

Proposition 2 *If a strategy-proof social choice function* $f : \mathcal{R}_1 \times \ldots \times \mathcal{R}_n \to X$ *attributes coalitional veto power to the set of experts* E*, then it ensures full information disclosure.*

Here is the intuition for the result. Suppose that only some information, or none, is disclosed by the experts, and that disclosing some further information would modify the preferences of non-experts in a way that changed the social outcome. Then, coalitional veto power guarantees that there is an expert who has incentives to disclose this additional information.

Attributing coalitional veto power to the experts only requires a conflict of interest among experts when their preferences are unchanged, and yet the social outcome varies due to shifts in the preferences and votes of non-experts. Proposition 2 shows that endowing experts with such power is sufficient to guarantee full information disclosure, but it is not necessary, unless further conditions are imposed on the set of types.

Definition 6 *The set* Θ_i *is rich in the domain* $\mathcal{R}_i$ *if every function* $\theta_i : 2^{\mathcal{I}} \to \mathcal{R}_i$ *belongs to* Θ_i*.*

Informally we can say that if an agent has a rich set of possible types, the way in which she may react to additional information is unpredictable for the designer. Under this additional assumption, we obtain a full characterization result.

Proposition 3 *Suppose the type set* Θ_i *is rich for all* $i \in A$*. A strategy-proof voting rule* $f : \tilde{\mathcal{D}}^n \to X$ *ensures full information disclosure only if it attributes coalitional veto power to the set of experts* E*.*

In the rest of the paper we restrict the attention to the case in which the type set Θ_i is rich for all $i \in A$, for which our coalitional veto power condition is necessary and sufficient. Notice that any constant social choice function is strategy-proof and trivially ensures full information disclosure in the universal preference domain. To avoid conclusions that refer to such uninteresting functions, we shall consider from now on to social choice functions that satisfy the additional, mild and compelling property of unanimity. Unanimity requires that if all agents agree at some profile that an alternative $x \in X$ is their preferred one, then it has to be chosen by the social choice function. We have already mentioned that two distinct cases arise, depending on whether the planner knows or does not know who will be the future experts at the time of deciding what rule to adopt. In the first case, social choice functions can be defined by endowing the set of experts with coalitional veto power, thus ensuring full information disclosure. In the latter case, however, it is not possible to achieve full information disclosure at large, as expressed by the following corollary of Proposition 3.

Corollary 1 *Suppose Θ_i is rich for every agent $i \in A$, and that the designer does not know who are the experts. No unanimous and strategy-proof voting rule $f : \tilde{\mathcal{D}}^n \to X$ ensures full information disclosure.*

The intuition of this result is clear. Consider for simplicity that there are only two alternatives $a, b \in X$. By strategy-proofness and unanimity if x is the preferred alternative at P_i for all $i \in A$, then $f(P) = x$ for both $x \in \{a, b\}$. So both a and b are in the range of f. Now consider a (strict) preference P_i for agent i such that a is the preferred alternative and a preference P_i' such that b is the preferred alternative. Let $\bar{\theta}_i \in \Theta_i$ be such that $\bar{\theta}_i(I) = P_i'$ for all $I \subset \mathcal{I}$ and $\bar{\theta}_i(\mathcal{I}) = P_i$ and $\hat{\theta}_i \in \Theta_i$ be such that $\hat{\theta}_i(I) = P_i$ for all $I \subset \mathcal{I}$ and $\hat{\theta}_i(\mathcal{I}) = P_i'$. Consider any arbitrary $M \subset A$. Consider the type profile $(\bar{\theta}_M, \hat{\theta}_{-M})$. Suppose first that M is the set of experts. By strategy-proofness and unanimity $f(\bar{\theta}_M(\mathcal{I}), \hat{\theta}_{-M}(I)) = a$ for all $I \subset \mathcal{I}$, and therefore the social choice function f ensures full information disclosure only if $f(\bar{\theta}_M(\mathcal{I}), \hat{\theta}_{-M}(\mathcal{I})) = a$. Similarly, suppose that $A \backslash M$ is the set of experts. By strategy-proofness and unanimity $f(\bar{\theta}_M(I), \hat{\theta}_{-M}(\mathcal{I})) = b$ for all $I \subset \mathcal{I}$, and therefore the social choice function f ensures full information disclosure only if $f(\bar{\theta}_M(\mathcal{I}), \hat{\theta}_{-M}(\mathcal{I})) = b$ which is a contradiction.

5 An application: Single-Peaked Domain

Consider a set X of ordered alternatives, which may be identified with an interval in the real line, or with a finite integer interval $[a, b]$. For each $i \in A$, R_i is single-peaked over X if there exists a unique $B(R_i) \in X$ (agent i's peak), and xP_iy for all $x, y \in X$ such that $y < x \leq B(R_i)$ or $B(R_i) \geq x > y$. Let $\hat{\mathcal{D}}$ denote the set of all single-peaked preferences and $|A|$ be odd.

The class of unanimous and strategy-proof social choice functions in the domain of all single-peaked alternatives was proved to be that of generalized median voter (Moulin 1980), which can be described through the use of left coalition systems (Barberà, Gul and Stacchetti 1993).

Definition 7 *A left coalition system on $X = [a, b]$ is a correspondence W assigning to every $x \in X$ a non-empty collection of non-empty coalitions $W(x)$, satisfying the following requirements:*

1. if $c \in W(x)$ and $c \subset c'$, then $c' \in W(x)$;

2. if $x' > x$ and $c \in W(x)$, then $c \in W(x')$; and

3. $W(b) = 2^N$.

Definition 8 *Given a left coalition system W on X, its associated generalized median agent rule $f : \hat{\mathcal{D}}^n \to X$ is defined so that, for all profiles R,*

$$f(R) = x \mathit{iff} \{i | B(R_i) \leq x\} \in W(x)$$

and

$$\{i|B(R_i) \leq y\} \notin W(y) \text{for all } y < x.$$

Proposition 4 ***(Moulin 1980)*** *A unanimous social choice function whose domain is the set of all single-peaked preferences on X is strategy-proof if and only if it is a generalized median agent rule.*

We first look at the case in which the designer knows in advance who are the set of experts and can use this information to endow experts of coalitional veto power.

The following proposition follows from our main result.

Proposition 5 *Let the type set Θ_i be rich in $\hat{\mathcal{D}}$ for each $i \in A$. Then, a social choice function $f : \hat{\mathcal{D}}^n \to X$ is strategy-proof, unanimous and ensures full information disclosure if and only if it is a generalized median agent rule with an associated left coalition system such that for each alternative $x \in X$, (i) there exists $c \in G(x)$ such that $c \subseteq E$; and (ii) for all $c \in G(x)$, $c \cap E \neq \varnothing$.*

Full information disclosure restricts the set of left coalition systems in two ways: first for every alternative x, it must contain a coalition $c(x)$ formed by the experts alone. Second, no coalition can contain only non-experts. Unanimity, in turn, requires that the empty coalition cannot belong to the left coalition system associated with any alternative.

Proposition 3 tells us that coalition veto power has to be granted to experts. In our case coalitional veto power narrows the set of generalized median agent rules guaranteeing information disclosure in a very intuitive way, as expressed in the following corollary.

Corollary 2 *A generalized median voter rule that attributes coalitional veto power to the set of experts E always selects alternatives that lie in between the minimum and the maximum peak of experts.*

The formal proof of this corollary is left to the reader. To see how this simple condition works in favor of information disclosure, consider any pair of profiles for which the rule selects two different alternatives. Since both of them fall between the maximum and the minimum peak of the experts, at least two of them will have opposite views on the two outcomes. Whenever the change from one of the two profiles to the other can be prompted by information disclosure, some expert will be inclined to reveal additional information. This was, indeed, the main intuition behind our initial motivational example. When the number of experts is large enough, attributing them veto power becomes compatible with anonymity[2], thus avoiding the bias in favor of experts that, in general, is needed to induce information disclosure.

Corollary 3 *Suppose n is odd. If $|E| \geq \frac{n+1}{2}$, then the median voter rule ensures full information disclosure.*

The median voter rule, when the number of experts is sufficiently large, is an example of an anonymous, efficient, strategy-proof rule that ensures full disclosure information.

We are now ready to consider the case in which the designer does not know the identity of the experts. In this case it seems natural to focus on anonymous social choice functions.

Proposition 6 *Suppose agents have single-peaked preferences and the type set Θ_i is rich for every agent $i \in A$. The median voter rule ensures better information disclosure than any other anonymous, unanimous, strategy-proof voting rule.*

Corollary 3 gives the intuition for this result.The median voter rule minimises the cardinality of the set of experts that is needed to ensure full information

[2] A social choice function $f : \mathcal{R}_1 \times \ldots \times \mathcal{R}_n \to X$, is anonymous if for all R, R', $f(R) = f(R')$ whenever R is a permutation of R'.

disclosure. If the cardinality is exactly (n-1)/2, the median voter rule is the only anonymous, unanimous, strategy-proof social choice function that ensures full information disclosure.

6 Appendix

Proof of Proposition 1. Sufficiency. Consider any $\theta \in \Theta$. Suppose that the conditions stated in Proposition 1 hold but there exists an equilibrium m^* such that $f(\theta_E(\mathcal{I}), \theta_N(g(m^*)) \neq f(R(\mathcal{I}))$. Notice that $g(m^*) \subset \mathcal{I}$. By assumption, there exists $i \in E$, $\theta_i(\mathcal{I}) = P_i$ and $m_i = I \supset g(m^*)$ such that $f(\theta_E(\mathcal{I}), \theta_N(I))P_i f(\theta_E(\mathcal{I}), \theta_N(g(m^*)))$. This contradicts that m^* is an equilibrium because expert $i \in E$ has a profitable deviation $m_i = I$ with $g(m^*_{-i}, m_i) = I$.

Necessity. Suppose there exists $\theta \in \Theta$ and $I \subset \mathcal{I}$,such that $f(\theta_E(\mathcal{I}), \theta_N(I)) \neq f(\theta(\mathcal{I}))$ and for all $I' \supset I$ and for all $i \in E$, $f(\theta_E(\mathcal{I}), \theta_N(I))P_i f(\theta_E(\mathcal{I}), \theta_N(I'))$, where $P_i = \theta_i(\mathcal{I})$. Consider the strategy profile m such that $m_i = I$ for all $i \in E$ and notice that $g(m) = I$, it follows that m is a Nash equilibrium of the θ-game, contradicting that f ensures full information disclosure.

Proof of Proposition 2. Suppose f attributes coalitional veto power to the set E. Fix any $\theta \in \Theta$. Let $\bar{m}$ be a strategy profile such that $\bar{m}_i = \mathcal{I}$ for all $i \in E$. It is immediate to check that $\bar{m}$ is a Nash equilibrium of the game and $f(\theta_E(\mathcal{I}), \theta_N(g(\bar{m})) = f(\theta(\mathcal{I}))$. To prove that there are no other Nash equilibrium outcomes suppose m be a Nash equilibrium involving partial or no disclosure and $f(\theta_E(\mathcal{I}), \theta_N(g(m))) \neq f(\theta(\mathcal{I}))$. By assumption there is an expert $i \in E$ such that $f(\theta(\mathcal{I}))P_i f(\theta_E(\mathcal{I}), \theta_N(g(m)))$. Expert i can profitably deviate by announcing $m'_i = \mathcal{I}$: in fact $g(m'_i, m_{-i}) = \mathcal{I}$ and therefore $f(\theta_E(\mathcal{I}), \theta_N(g(m'_i, m_{-i})))P_i f(\theta_E(\mathcal{I}), \theta_N(g(m)))$.

Proof of Proposition 3. If f does not attribute coalitional veto power to the set E, then there exists a pair of preference profiles (R_E, R_N), (R_E, R'_N) such that $f(R_E, R_N) \neq f(R_E, R'_N)$ and $f(R_E, R_N) R_i f(R_E, R'_N)$ for all $i \in E$. Consider a type $\theta_j \in \Theta_j$ such that $\theta_j(I) = R_j$ for all $I \subset \mathcal{I}$ and and $\theta_j(\mathcal{I}) = R'_j$. This type exists because Θ_j is rich for every $j \in A$. Consider a strategy profile m^* such that for all $i \in E$, $m^*_i = \hat{I}$ for some $\hat{I} \subset \mathcal{I}$. The strategy profile m^* is a Nash equilibrium of the θ-game. In fact any deviation $m'_i \neq \mathcal{I}$ is irrelevant beca.use it does not modify non-experts' preferences and, consequently, the final outcome, while the deviation $m'_i = \mathcal{I}$ is not profitable for any expert i, because $\theta_j(\mathcal{I}) = R'_j$ for all $j \in N$, and $f(R_E, R_N) R_i f(R_E, R'_N)$ for all $i \in E$.

Proof of Proposition 5. Sufficiency. Consider any generalized median voter rule with a left coalition system W such that for all $x \in X$, $i)$ there exists $c_x \subseteq E$ and $ii)$ $c \in W(x)$ only if there exists $i \in E \cap c$. This social choice function is strategy-proof and efficient and therefore unanimous because at every profile it selects an allocation that is in between the minimum and the maximum peak of the agents. Fix an arbitrary $R \in \tilde{\mathcal{D}}^n$. Let $l \in E$ be such that for all $j \in E$, $B(R_l) \leq B(R_j)$ and let $r \in E$ be such that for all $j \in E$, $B(R_r) \geq B(R_j)$. By (ii) $f(R) \geq B(R_l)$ and by (i) $f(R) \leq B(R_r)$. Consider now any $R' \in \hat{\mathcal{D}}^n$ such that for all $j \in E$, $R_j = R'_j$. For the same arguments as above $B(R_l) \leq f(R') \leq B(R_r)$. Suppose $f(R') \neq f(R)$, and without loss of generality, suppose that $f(R') < f(R)$. It follows that $f(R') P_l f(R)$ and $f(R) P_r f(R')$ and coalition veto power is satisfied.

Necessity. Consider any rule f that is strategy-proof and satisfies unanimity, hence f has to be a generalized median agent rule. By Proposition **3** the voting rule must satisfy coalitional veto power relative to E. Let W be its associated left coalition system and $X = [a, b]$. Suppose first that there exists $x < b$ such that for each coalition $c \in W(x)$, a member of c is a non-expert. Consider $R \in \hat{\mathcal{D}}^n$ such that for all $i \in E$, $B(R_i) = x$ and for all $j \in N$, $B(R_j) =$

b; by definition of $8 f(R) > x$. Let R_j' be a preference such that $B(R_j') = x$. By unanimity $f(R_E, R_N') = x$. Consider a profile θ such that for all $i \in A$,$\theta_i(I) = R_i'$ for all $I \subset \mathcal{I}$ and $\theta_i(\mathcal{I}) = R_i$. It follows that there exists a Nash equilibrium of the θ-game such that for all $i \in E$,$m_i^* = \hat{I}$ for some $\hat{I} \subset \mathcal{I}$ and $f(\theta_E(\mathcal{I}), \theta_N(g(m^*))) \neq f(\theta(\mathcal{I}))$. and the social choice function does not ensure full information disclosure. Suppose now that there exists $x < b$ and $c \in W(x)$ such that $E \cap c = \varnothing$. Let $R_i', \bar{R}_i \in \hat{\mathcal{D}}$ be a pair of preference such that $B(R_i') = x$ and $B(\bar{R}_i) = b$. Consider the preference profile $(\bar{R}_E, R_N')$: by Definition $8 f(\bar{R}_E, R_N') \leq x$. Let $\bar{R} = (\bar{R}_1, ..., \bar{R}_n)$. By unanimity $f(\bar{R}) = b$. Consider a profile θ such that for all $j \in N$,$\theta_j(I) = \bar{R}_j$ for all $I \subset \mathcal{I}$ and $\theta_j(\mathcal{I}) = R_j'$ and for all $i \in E$,$\theta_i(\mathcal{I}) = \bar{R}_i$. It follows that there exists a Nash equilibrium of the θ-game such that for all $i \in E$,$m_i^* = \hat{I}$ for some $\hat{I} \subset \mathcal{I}$ and $f(\theta_E(\mathcal{I}), \theta_N(g(m^*))) = f(\bar{R}) \neq f(\bar{R}_E, R_N') = f(\theta(\mathcal{I}))$. and the social choice function does not ensure full information disclosure.

Proof of Proposition 6. By Moulin (1980) we know that a social choice function is anonymous, unanimous and strategy-proof if and only if there are $n+1$ points $p_1, \ldots, p_{n+1}$ in $[a, b]$ such that for all profiles $f(P) = med(p_1, \ldots, p_{n+1}, B(P_1), \ldots B(P_n))$. By Corollary 2 when $|E| < \frac{n+1}{2}$, it follows immediately that no anonymous strategy-proof social choice function may ensure full information disclosure, and by Corollary 3 we know that the median rule ensures information disclosure when $|E| \geq \frac{n+1}{2}$. We prove that every anonymous, unanimous and strategy-proof voting rule different than the median rule does not ensure full information disclosure when $|E| = \frac{n+1}{2}$. Consider an arbitrary set of agents M with $|M| = \frac{n+1}{2}$. Consider any anonymous, unanimous and strategy-proof voting rule f such that for some $z \in X$, $c \in C(z)$ if and only if $|c| \geq k$ with $k \neq \frac{n+1}{2}$. Suppose first $k < \frac{n+1}{2}$. Consider a pair of preference profiles R^0, R^1 with $R_M^1 = R_M^0 = R_M$, $B(R_i) = z$ for all $i \in M$, $B(R_j^0) = z$,

$B(R_j^1) = y$ for all $j \notin M$ with $y < z$. Consider a type profile θ such that for all $i \in M, \theta_i(\mathcal{I}) = R^0$ and for all $j \notin M$ $\theta_j(\varnothing) = R_j^0$ and $\theta_j(I) = R_j^1$ for all $I \in 2^{\mathcal{I}} \backslash \varnothing$. Consider a θ-game such that $M = E$. It is immediate to check that there exists a Nash equilibrium m^* such that $m_i^* = \varnothing$ for all $i \in E$ and $f(\theta_E, \theta_N(g(m^*))) = z \neq f(\theta(\mathcal{I})) = y$. The proof for the case $k > \frac{n+1}{2}$ is analogous: consider a pair of preference profiles $\bar{R}^0, \bar{R}^1$ with $\bar{R}_M^1 = \bar{R}_M^0 = \bar{R}_M$, $B(\bar{R}_i) = y$ for all $i \in M$, $B(\bar{R}_j^0) = y$, $B(\bar{R}_j^1) = z$ for all $j \notin M$ and a type-profile θ such that for all $j \notin M$ $\theta_j(\varnothing) = \bar{R}_j^0$ and $\theta_j(I) = \bar{R}_j^1$ for all $I \in 2^{\mathcal{I}} \backslash \varnothing$.

References

[1] David Austen-Smith and John R. Wright. (1992). "Competitive lobbying for a legislator's vote", *Social Choice and Welfare*, 9, 3, 229–257.

[2] David Austen-Smith and Jeffrey. S. Banks. (1996). "Information Aggregation, Rationality, and the Condorcet Jury Theorem", *American Political Science Review* 90, 34-45.

[3] David Austen-Smith and Timothy J. Feddersen. (2005). "Deliberation and Voting Rules." In Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks, eds. David Austen-Smith and John Duggan. Heidelberg: Springer.

[4] David Austen-Smith and Timothy J. Feddersen. (2006) "Deliberation, Preference Uncertainty, and Voting Rules", *American Political Science Review*, 100, 2, 211-217.

[5] Salvador Barberà, Faruk Gul, and Ennio Stacchetti. (1993). "Generalized Median agent Schemes and Committees", *Journal of Economic Theory* 61, 262-289.

[6] Sourav Bhattacharya and Arijit Mukherjee. (2013). "Strategic Information Revelation when Experts Compete to Influence," *RAND Journal of Economics*, 44,3,522-544.

[7] Mathias Dewatripont and Jean Tirole. (1999). "Advocates", *Journal of Political Economy,* 107(1),1-39.

[8] Bashkar Dutta and Arunava Sen. (2011). "Nash Implementation with Partially Honest Individuals", *Games and Economic Behavior*, 74, 154-169.

[9] Timothy Feddersen and Wolfgang Pesendorfer. (1998). "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting", *American Political Science Review*, 92, 1, 23-35.

[10] Matthew Gentzkow Emir Kamenica, (2017) "Competition in Persuasion", *Review of Economic Studies*, 84, 1, 1, 300-322.

[11] Thomas W. Gilligan and Keith Krehbiel. (1989). "Asymmetric Information and Legislative Rules with a Heterogeneous Committee", *American Journal of Political Science,* 33, 2 459-490.

[12] Thomas W. Gilligan and Keith Krehbiel. (1990). "Organization of Informative Committees by a Rational Legislature", *American Journal of Political Science,* 34, 2, 531-564.

[13] Matthew O. Jackson and Xu Tan. (2013). "Deliberation, Disclosure of Information, and Voting ", *Journal of Economic Theory*, 148, 1, 2-30.

[14] Navin Kartik, Olivier Tercieux, and Richard Holdeny. (2014). "Simple Mechanisms and Preferences for Honesty", *Games and Economic Behavior*, 83, 284-290.

[15] Vijay Krishna and John Morgan. (2001). "Asymmetric Information and Legislative Rules: Some Amendments". *American Political Science Review*, 95, 2, 435-452.

[16] Lipman, B. L. and D.J. Seppi. (1995). "Robust Inference in Communication Games with Partial Provability", *Journal of Economic Theory*, 66, 370-405.

[17] Paul Milgrom and John Roberts. (1986). "Relying on the information of interested parties", *Rand Journal of Economics*, 17, 18-32.

[18] Hervé Moulin. (1980). "On Strategyproofness and Single-peakedness", *Public Choice* 35, 4, 437-455.

[19] Elisabeth Schulte. (2010). "Information aggregation and preference heterogeneity in committees", *Theory and Decision,* 69, 97-118.