Truth-revealing voting rules for large populations

Matías Núñez and Marcus Pivato

Abstract

Deterministic voting rules are notoriously susceptible to strategic voting. We propose a new solution to this problem for large electorates. For any deterministic voting rule, we can design a stochastic rule that asymptotically approximates it in the following sense: for a sufficiently large population of voters, the stochastic voting rule (*i*) incentivizes every voter to reveal her true preferences and (*ii*) produces the same outcome as the deterministic rule, with very high probability.

KEYWORDS: Large Elections, Truth-telling, Incentives.

JEL CODES: D71.

1 Introduction

Strategic voting is a pervasive problem in social choice theory. The Gibbard-Satterthwaite Theorem (1973, 1975) says that any nontrivial deterministic voting rule is susceptible to strategic voting.[1] However, as Gibbard (1977) noted, we can prevent strategic voting if we incorporate some randomness into the voting rule. To see this, consider the *random dictatorship* rule: each voter is asked to report a preference order, and one of these preference orders is then selected at random. It is easy to see how this rule prevents strategic voting: each voter anticipates that, if her vote is selected, she should report her true preferences, whereas in any other case, her vote is simply irrelevant. It is then a dominant strategy for her to reveal her true preferences. However, the random dictatorship is undesirable, because one voter can impose an outcome even if it is the worst outcome for every other voter. In effect, this rule removes the incentives for misrepresentation at the cost of ignoring social preferences. In contrast, deterministic rules aim to reflect the "will of the group". We propose a compromise: a class of randomized mechanisms that ensure, for large electorates, that the "will of the group" is represented, while at the same time voters do not have incentives to misrepresent their opinions. How do these randomized mechanisms work?

To illustrate, consider the Borda rule: each voter ranks each of the different alternatives, each rank is worth a certain number of points, and the winner is the alternative with the most points (invoking some tie-breaker rule in the event of a tie). To incentivize sincere voting, we build a "stochastic" Borda rule as follows. Each voter is asked to declare a complete preference order over the alternatives. The outcome is now determined through a lottery (independent of the voters' announcements). With probability $1-q$, we select the winner using the (deterministic) Borda rule. However, with probability q, we use the following random device instead:

1. First randomly choose one of the voters n and any pair of alternatives a and b.

2. If n prefers a to b, then select a. Otherwise, select b.

Consider now the behavior of a rational voter in this stochastic voting rule. Clearly, when confronted with the random device, she has a unique dominant strategy: reveal her true ordinal preferences. On the other hand, under the deterministic Borda rule, she will have an incentive to misrepresent her true preferences only when her vote is *pivotal*, meaning that it could modify the outcome of the election. But if the probability of such a pivotal event is small enough relative to q, then the *expected* utility gain from misrepresenting her preferences becomes negligible in comparison with the expected utility *loss* of misrepresenting her preferences when confronted with the random device. Hence, we can adequately calibrate the probability q to ensure that truth-revelation is her strictly dominant strategy. In fact, we will let q depend on the size of the electorate; we will then give a sufficient condition which makes pivotal events unlikely enough that truthful voting becomes strictly dominant while q converges towards zero as the electorate grows large. In other words, the bigger the electorate, the more probable (*i*) that a voter reveals her true preferences and (*ii*) that the actual outcome coincides with the sincere one under

[1] This assumes an unrestricted domain of preferences. Under reasonable restrictions such as single-peakedness, incentive-compatible voting rules do exist (see Moulin (1980) among others).

Plurality rule. Thus, with this stochastic voting rule, the "true" Plurality winner will be selected, with very high probability.

The basic idea behind our stochastic Borda rule applies to most of the well-known voting rules, including ordinal voting rules, cardinal voting rules (e.g. evaluative voting), and approval voting. In each case, we will introduce a random device which is activated with some relatively small probability. Although unlikely, these random "checks" are enough to incentivize sincerity as long as the pivot probabilities are not too high. In each case, we will study the asymptotic performance of the rule as the population of voters becomes large, and show that, with very high probability, our stochastic rules will select the same alternative that would win under the original deterministic voting rule, if the voters had been sincere.

The rest of this paper is organized as follows. Section 2 reviews previous literature proposing similar truth-eliciting mechanisms. Section 3 introduces some basic notation and terminology about voting rules and stochastic voting rules. Section 4 defines a *culture* to be (roughly) the set of all beliefs that any voter could reasonably have about the voting behaviour of the other voters. It also introduces our key hypothesis, *regularity*, which states (roughly) that every voter believes that a nearly-tied vote is very unlikely. Section 5 deals with ordinal voting rules, Section 6 deals with approval voting, and Section 7 deals with cardinal voting rules. Finally, Section 8 applies the previous results to obtain an implementation in Bayesian Nash equilibrium. All proofs are in the Appendix.

2 Prior literature

In public finance, it is well-known that tax evasion can be reduced through random audits. The key insight in this paper is analogous: subjecting voters to "random checks", even with a tiny probability, can severely curtail strategic voting in large elections. If the voter misrepresents her preferences and her vote gets checked, then she will surely end up with a worse outcome than if she had voted honestly. A similar idea is present in the *virtual implementation* literature. Classical implementation theory observes that many social choice rules are not Nash implementable, because they violate Maskin monotonicity. Virtual implementation overcomes this difficulty by using random mechanisms, which are arbitrarily close to the original ones in probability (see Jackson (2001) for a review). This implementation concept was introduced by Matsushima (1988) and Abreu and Sen (1991) and achieves remarkable results. For example, if the voters have complete information about one another, then *any* social choice rule can be virtually implemented in Nash equilibrium (Abreu and Sen, 1991) or iterated undominated strategies (Abreu and Matsushima, 1992). Even with incomplete information, a very large class of social choice rules can be virtually implemented in Bayesian Nash equilibrium (Serrano and Vohra, 2005), or even robustly virtually implemented (Artemov et al., 2013). The basic idea of virtual implementation is that it is sufficient to obtain a very high *probability* of selecting a socially optimal outcome, rather than certainty; the present paper shares this idea. But we propose a very simple device (using "random checks") to incentivize sincere voting, whereas most papers in implementation use rather abstract mechanisms to obtain their strong conclusions.

Our contribution also builds on the recent literature studying asymptotic restrictions on strategic voting as the voting population becomes large. For instance, Ehlers et al. (2004) and Renault and Trannoy (2007, 2011) focus on the average voting rule in large populations, while Laslier and Weibull (2013) consider the Condorcet Jury Theorem. McLennan (2011) states an impossibility result in Poisson games, and Carroll (2013) and Azevedo and Budish (2015) are concerned with the extent of manipulation in large environments.

Among these papers, Laslier and Weibull (2013) is the most closely related to our work. It aims to incentivize truth-revelation in an epistemic voting model. In this model, voters receive independent noisy signals of some unknown binary random variable. If each voter votes honestly (according to her private signal), then the Condorcet Jury Theorem says that simple majority vote should identify the true value, with very high probability. However, a strategic voter will condition her voting strategy on the event that she is a pivotal voter, and since such an event may reveal information which is contrary to her private signal, she may find it optimal to misrepresent her private information. Like us, Laslier and Weibull are interested in the asymptotic behaviour of the voting rule for large populations, and incentivize honesty by offering each voter a small probability of being a random dictator. This probability becomes increasingly small as the population gets larger, but it is still large enough that it outweighs the aforementioned incentive for dishonesty; thus, everyone votes sincerely in the unique (Bayesian) Nash equilibrium. In Section 8, we will also obtain a Bayesian Nash implementation result. However, unlike Laslier and Weibull (and the aforementioned literature on virtual implementation), we will also consider weaker notions of implementation which do *not* depend on any particular equilibrium concept from game theory (in Sections 5, 6 and 7). In particular, we do not need to assume that each voter has correct beliefs about the preferences and/or strategic behaviour of the other voters. Since we intend our model to apply to social choice in very *large* populations, we think that this is a much more realistic assumption about the informational environment of the voters.

For similar reasons, McLennan (2011), Carroll (2013), and Azevedo and Budish (2015) consider models in which each voter treats the actions of the other voters as independent, identically distributed (i.i.d.) random variables. In such an environment, McLennan (2011) proves a version of Gibbard (1977)'s impossibility theorem, which states that the only anonymous, strategy-proof, Pareto optimal stochastic voting rule is the random dictatorship. On a more optimistic note, Carroll (2013) argues that, even if voters *can* gain by strategic misrepresentation, most of them will act sincerely if the expected gains are too small.[2] He thus proposes to quantitively measure the "susceptibility to manipulation" of a voting rule as the maximum gain in expected utility a voter could obtain by misrepresentation; he then computes the large-population asymptotics of this measure for several common voting rules. Meanwhile, Azevedo and Budish (2015) define the "large-market limit" of a mechanism by taking the limit as the population goes to infinity of the mechanism's behaviour, as seen by a single agent who regards all other agents as independent identically distributed (i.i.d.) random variables distributed according to some probability distribution μ. Azevedo and Budish say the mechanism is "strategy-proof in the large" if this infinite-population limit is strategy-proof for all μ with full support. They present this as a unifying framework for several classic results in the mechanism design literature.

The present paper differs from the contributions by McLennan (2011), Carroll (2013), and Azevedo and Budish (2015) in that we allow a voter's beliefs to take the form of *any* probability distribution over the other voters, not just an i.i.d. distribution. In other words, we allow a voter to believe that the actions of the other voters are correlated. We further differ from McLennan (2011) in that we only require the asymptotic *probability* of strategic voting to become small, whereas he seeks to exclude strategic voting altogether, and thereby gets an impossibility result. On the other hand, in contrast to Azevedo and Budish (2015), we work with *finite* (but large) populations, rather than only considering the infinite-population limit.

While the previous papers are in economic theory, a recent branch of the literature on computer science has focused on similar ideas. Among these papers, Procaccia (2010) designs strategyproof randomized voting rules that are close, in a standard approximation sense, to prominent score-based (deterministic) voting rules. His method for eliciting preferences is quite different from ours: it chooses a voter at random and then selects an alternative with probability proportional to the alternative's score in the voter's vector. Birrell and Pass (2011) focus on a concept of "ϵ-strategyproofness" for large populations, and prove that any deterministic voting can be approximated by some ϵ-strategyproof randomized voting rule. Chierichetti and Kleinberg (2014) study similar issues in an incomplete information setting *à la* Condorcet Jury Theorem, in which voters have common value preferences and are uncertain about the true state of the world. Finally, Leung et al. (2015) study a related problem under a bounded rationality approach assuming that the voters have coarse i.i.d beliefs over the preferences of the rest of the voters. Our contribution is hence different from theirs: first, the device that induces truth-telling in our rules is particularly simple and applies to a very large class of rules and voter beliefs; furthermore, our main results deal with implementation, and hence are equilibrium results, whereas the previously mentioned results do not take into account strategic considerations.

3 Voting rules

Let $\mathbb{N}$ denote the set of natural numbers. Let $\mathcal{V}$ be the set of messages which could be sent by each voter; $\mathcal{V}$ could be finite or infinite. For any $N \in \mathbb{N}$, an *N-voter profile* is an element $\mathbf{v} = (v_n)_{n=1}^N$ of the Cartesian power $\mathcal{V}^N$. Let $\mathcal{A}$ be a finite set of alternatives with $|\mathcal{A}| \geq 3$. Let $\widetilde{\mathcal{A}}$ denote the set of all $\mathcal{A}$-valued random variables. (Formally, an element of $\widetilde{\mathcal{A}}$ is a measurable function $\widetilde{a} : \Omega \longrightarrow \mathcal{A}$, where Ω is some probability space.) We will refer to elements of $\widetilde{\mathcal{A}}$ as *random alternatives*. We will consider two different sorts of

[2]Partial honesty, recently proposed in implementation by Dutta and Sen (2012), has a similar idea. A partially honest player is one who has a strict preference for revealing her true type over lying when truthtelling does not lead to a worse outcome than that which she obtains when lying.

voting rules: deterministic and stochastic. A deterministic voting rule assigns a unique alternative for every voter profile $\mathbf{v}$, whereas a stochastic rule assigns a lottery over the different alternatives.

Formally, a *deterministic voting rule* is a sequence $F := (F_N)_{N=1}^{\infty}$ where, for all $N \in \mathbb{N}$, $F_N : \mathcal{V}^N \longrightarrow \mathcal{A}$ is a function which assigns a unique alternative to every profile. Each F_N is assumed to be *anonymous*: if $\sigma : [1 \dots N] \longrightarrow [1 \dots N]$ is a permutation, and we define $\mathbf{v}' := (v'_n)_{n=1}^N$ by setting $v'_n := v_{\sigma(n)}$ for all $n \in [1 \dots N]$, then $F_N(\mathbf{v}') = F_N(\mathbf{v})$. For simplicity, we will just call F a *voting rule*.[3] We impose no structure on $\mathcal{V}$; thus, most of the standard voting rules are allowed in our model, including the following classes.

Ordinal Rules: Let $\mathcal{P}$ be the set of all possible preference orders over $\mathcal{A}$. If $\mathcal{V} = \mathcal{P}$, then we say F is an *ordinal* voting rule. Most of the voting rules considered in the literature are ordinal voting rules, including the plurality rule, antiplurality rule, Borda rule, single transferable vote, etc. This class also includes rules where each voter declares her "ideal point" on a linear domain, and preferences are assumed to be single-peaked; this includes the median rule and the average rule (Renault and Trannoy, 2007, 2011).

Cardinal Rules: Let $\mathcal{U} := \{u : \mathcal{A} \longrightarrow [0,1];\ \min_{a \in \mathcal{A}} u(a) = 0 \text{ and } \max_{a \in \mathcal{A}} u(a) = 1\}$. We interpret the elements of $\mathcal{U}$ as normalized, nonconstant von Neumann-Morgenstern (vNM) utility functions on $\mathcal{A}$. Any nonconstant vNM utility function has a unique representative in $\mathcal{U}$. If $\mathcal{V} = \mathcal{U}$, then we say that F is a *cardinal* voting rule. Examples include the evaluative voting rule (Dhillon and Mertens, 1999; Núñez and Laslier, 2014), the Nash rule (which maximizes the product of the utilities), and the relative egalitarian rule of Kalai and Smorodinsky (1975), which maximizes the minimum utility.

Scoring Rules: A *scoring rule* F is a voting rule where each voter $n \in [1 \dots N]$ assigns a "score" $s_a^n \in [0,1]$ to each alternative $a \in \mathcal{A}$. Let $S_a := \sum_{n=1}^N s_a^n$ be the total score for alternative a; the alternative with the highest total score wins. A well-known example is the Approval Voting rule (Brams and Fishburn, 1983; Laslier and Sanver, 2010). But this class also includes many ordinal rules (e.g. plurality, Borda) and cardinal rules (e.g. evaluative voting).

A *stochastic voting rule* is a system $\widetilde{F} := (\widetilde{F}_N)_{N=1}^{\infty}$ where, for all $N \in \mathbb{N}$, $\widetilde{F}_N : \mathcal{V}^N \longrightarrow \widetilde{\mathcal{A}}$ is a function which assigns a random alternative to every profile.[4] Again, we assume that $\widetilde{F}_N$ is anonymous —i.e. invariant under all permutations of $[1 \dots N]$. Given a (deterministic) voting rule F, we might say that a stochastic voting rule $\widetilde{F}$ is a good "approximation" of F if $\widetilde{F}$ is very likely to agree with F when the number of voters gets large enough. The next definition formalizes this idea.

Definition. For any $N \in \mathbb{N}$, let $P_N(F, \widetilde{F}) := \inf_{\mathbf{v} \in \mathcal{V}^N} \text{Prob}\left[\widetilde{F}_N(\mathbf{v}) = F_N(\mathbf{v})\right]$.

We say that $\widetilde{F}$ is *asymptotically equal to* F if $\lim_{N \to \infty} P_N(F, \widetilde{F}) \ = \ 1$.

Thus, if $\widetilde{F}$ and F are asymptotically equal, then in a sufficiently large population, the outcome of $\widetilde{F}$ will be the same as the outcome of F, with very high probability, independently of the actual profile which occurs.

4 Cultures

Our objective is to design, for every voting rule F, a stochastic rule $\widetilde{F}$ which induces truthful voting when the population is large, and which also delivers the same outcome as F. To do this, we need some assumptions about what information is available to each agent in the model.

Let $\Delta(\mathcal{A})$ be the set of probability distributions over $\mathcal{A}$. We assume that each voter has preferences over $\Delta(\mathcal{A})$, given by some vNM utility function in the set $\mathcal{U}$ (as defined in Section 3), which is known only to her. The mechanism designer does not know the true vNM utility functions of the voters. Let ρ be a probability distribution on $\mathcal{U}$, describing the designer's beliefs about the voters: we suppose that

[3] Strictly speaking, F should be called a *variable-population, anonymous, deterministic voting rule.*

[4] Equivalently, if $\Delta(\mathcal{A})$ is the set of all probability distributions over $\mathcal{A}$, we could represent $\widetilde{F}_N$ as a function from $\mathcal{V}^N$ into $\Delta(\mathcal{A})$. But the random variable representation is more convenient.

the designer regards each voter's utility function as a random variable with distribution ρ.[5] The designer wishes to design a mechanism such that, for each voter, there is a high probability that this voter will find it her optimal strategy to report her true preferences, where this probability is computed using ρ.

Each voter does not know the preferences or voting behaviour of the other voters, but she has some beliefs about them, given by a probability distribution on the set $\mathcal{V}^N$ of profiles. These beliefs might not be correct, and different voters might have different beliefs; we do not assume any relationship between a given voter's beliefs and reality, or between the beliefs of different voters, or between the beliefs of the voters and those of the mechanism designer. However, we assume that the beliefs of all voters in a population of size N are drawn from some common set $\mathcal{B}_N$; this is the set of all beliefs which any "reasonable" person could have, given publicly available information. The key assumption in our model can be expressed informally as follows: *If N is very large, then every belief in $\mathcal{B}_N$ assigns an extremely low probability to a tie or near-tie occuring.*

Formally, for each $N \in \mathbb{N}$, let $\Delta(\mathcal{V}^N)$ be the set of probability distributions over $\mathcal{V}^N$. Each voter's beliefs are represented by some element $\beta \in \Delta(\mathcal{V}^N)$.[6] Let $\mathcal{B}_N \subseteq \Delta(\mathcal{V}^N)$ be the set of all possible beliefs which any voter could have about an N-voter profile. The sequence $\boldsymbol{\mathcal{B}} := (\mathcal{B}_N)_{N=1}^{\infty}$ is called a *culture*.

Two profiles $\mathbf{v}, \mathbf{v}' \in \mathcal{V}^N$ are *adjacent* if there exists some $m \in [1 \ldots N]$ such that $v_n = v'_n$ for all $n \in [1 \ldots N] \setminus \{m\}$. In other words, the two profiles only differ for voter m. A profile $\mathbf{v} \in \mathcal{V}^N$ is *nearly tied* for F_N if there is some adjacent profile $\mathbf{v}'$ such that $F_N(\mathbf{v}) \neq F_N(\mathbf{v}')$. In other words, a single voter could change the outcome, by changing her vote. Because the rule F_N is anonymous, this means that *any* single voter could change the outcome, by changing her vote.

Definition. For any belief $\beta \in \mathcal{B}_N$, let $\tau(\beta, F_N)$ be the probability (according to β) that the profile will be nearly tied for F_N. Let $\tau(\mathcal{B}_N, F_N) := \sup_{\beta \in \mathcal{B}_N} \tau(\beta, F_N)$. The culture $\boldsymbol{\mathcal{B}}$ is *regular* for the rule F if $\lim_{N\to\infty} N \cdot \tau(\mathcal{B}_N, F_N) \ = \ 0$.

Informally, $\tau(\mathcal{B}_N, F_N)$ is the highest probability that any voter in a population of size N could assign to the possibility of a nearly-tied profile (in the culture $\boldsymbol{\mathcal{B}}$). Intuitively, we would expect that $\lim_{N\to\infty} \tau(\mathcal{B}_N, F_N) = 0$, reflecting the idea that, in large societies, everyone believes that nearly-tied profiles will be extremely rare. *Regularity* is a slightly stronger condition: it requires that $\tau(\mathcal{B}_N, F_N) \xrightarrow[N\to\infty]{} 0$ "faster than $1/N$". Note that whether or not a culture is regular depends on the precise voting rule being used.

The well-known Impartial Culture (IC) model is obviously not regular for any neutral voting rule.[7] But IC is popular in the literature only because it is simple to define and easy to analyse, not because it is remotely plausible as a model of a real society. Indeed, the realism of IC has been questioned by Tsetlin et al. (2003) and Lehtinen and Kuorikoski (2007), among others. So although our main results do not apply to IC, we do not consider this to be a real shortcoming. We will now present examples of other, more realistic models which *do* satisfy the regularity hypothesis.

Regularity for scoring rules. Let F be any scoring rule, and for any alternative $a \in \mathcal{A}$, recall that S_a denotes its total score. If we regard the profile of votes as a β-random variable (for some $\beta \in \mathcal{B}_N$), then S_a is also a random variable. The profile is nearly tied for F only if the top two alternatives a and b satisfy $|S_a - S_b| \leq 1$, so that a single voter could tip the balance between a to b by changing her scores. (Recall that the scores assigned to each candidate are in $[0, 1]$, by assumption.) For any $\beta \in \mathcal{B}_N$ and any $a, b \in \mathcal{A}$, let $\beta_{a,b}$ denote the probability density function which β induces over the possible values of $S_a - S_b$. Heuristically, the culture $\boldsymbol{\mathcal{B}}$ will be regular for F if, as $N\to\infty$, the event "$|S_a - S_b| \leq 1$" receives a very small probability from $\beta_{a,b}$, for every $\beta \in \mathcal{B}_N$ and any $a, b \in \mathcal{A}$, as illustrated by the next example.

Example 1. Fix some $\mu_0, \varsigma > 0$. For all $N \in \mathbb{N}$, suppose that every belief $\beta \in \mathcal{B}_N$ is such that, for any two alternatives $a, b \in \mathcal{A}$, there is some $\mu \in \mathbb{R}$ with $|\mu| \geq N\mu_0$ such that $\beta_{a,b}$ is a normal probability distribution with variance $N\varsigma^2$ and mean $N\mu \in \mathbb{R}$. For example, if N is large and β believes that the scores $\{s_a^n;\ n \in [1 \ldots N] \text{ and } a \in \mathcal{A}\}$ are independent random variables, then the Central Limit Theorem suggests this is a good approximation. Then the culture $\boldsymbol{\mathcal{B}}$ is regular for F. (This is a consequence of Proposition 2 below.) ◇

[5] Note that the designer does *not* suppose that these random variables are independent; indeed, we do not assign to the designer any particular probabilistic beliefs about the profile of utility functions across the entire population.

[6] We suppose that N is large enough so that the voter's own preferences make up only a tiny part of the profile.

[7] A voting rule is *neutral* if it treats all alternatives the same.

Example 1 can be generalized as follows. Let $\Gamma:\mathbb{R}\longrightarrow\mathbb{R}_+$ be any unimodal probability density function with its mode at zero. Let $(\sigma_N)_{N=1}^\infty$ be a sequence such that

$$(1.1)\quad \lim_{N\to\infty}\sigma_N \;=\; \infty \qquad \text{but} \qquad (1.2)\quad \lim_{N\to\infty}\frac{\sigma_N}{N} \;=\; 0. \tag{1}$$

For all $N\in\mathbb{N}$, and any $\mu\in\mathbb{R}$, let γ_μ^N be the probability density function defined by:

$$\gamma_\mu^N(x) \;:=\; \frac{1}{\sigma_N}\,\Gamma\left(\frac{x-N\,\mu}{\sigma_N}\right), \qquad \text{for all } x\in\mathbb{R}. \tag{2}$$

Thus, the mode of γ_μ^N is shifted to $N\,\mu$, while the horizontal scale of γ_μ^N is stretched by a factor of σ_N, relative to Γ.

Proposition 2 *Fix $\mu_0>0$, let $(\sigma_N)_{N=1}^\infty$ be a sequence satisfying* (1)*, and let $(\epsilon_N)_{N=1}^\infty$ be a sequence such that*

$$\lim_{N\to\infty} N\epsilon_N \;=\; 0. \tag{3}$$

Let F be a scoring rule, and suppose that, for all $N\in\mathbb{N}$, every belief $\beta\in\mathcal{B}_N$ and all distinct $a,b\in\mathcal{A}$, there is some $\mu_{a,b}\in\mathbb{R}$ with $|\mu_{a,b}|\geq\mu_0$ such that $\left\|\beta_{a,b}-\gamma^N_{\mu_{a,b}}\right\|_\infty\leq\epsilon_N$. Then the culture $\boldsymbol{\mathcal{B}}$ is regular for F.

Example 1 is obtained as a special case of Proposition 2, by letting Γ be the standard normal distribution, and setting $\sigma_N:=\sqrt{N}\,\varsigma$ for all $N\in\mathbb{N}$. This implies that γ_μ^N is the normal distribution with mean $N\,\mu$ and variance $\sigma_N^2=N\,\varsigma^2$. However, Γ could be any unimodal distribution in Proposition 2, even one with infinite variance such as a Cauchy distribution. The coefficient σ_N essentially plays the role of the "standard deviation" of a random vote distribution for a population of size N. Proposition 2 makes the plausible assumption that σ_N grows sub-linearly as $N\to\infty$. Indeed, if a voter believes that the other voters were independent random variables, then she would expect that $\sigma_N=\mathcal{O}(\sqrt{N})$. We now propose another sufficient condition for a culture to be regular. Let $\Gamma:\mathbb{R}\longrightarrow\mathbb{R}_+$ be any probability density function with $\|\Gamma\|_\infty<\infty$. Given any sequence $(\sigma_N)_{N=1}^\infty$ and any $\mu\in\mathbb{R}$, we define a sequence $(\gamma_\mu^N)_{N=1}^\infty$ of probability density functions as in equation (2).

Proposition 3 *Let $(\epsilon_N)_{N=1}^\infty$ be a sequence satisfying condition* (3), and let $(\sigma_N)_{N=1}^\infty$ be a sequence such that

$$\lim_{N\to\infty}\frac{\sigma_N}{N} \;=\; \infty. \tag{4}$$

Let F be a scoring rule, and suppose that, for all $N\in\mathbb{N}$, every $\beta\in\mathcal{B}_N$, and all distinct $a,b\in\mathcal{A}$, there is some $\mu_{a,b}\in\mathbb{R}$ such that $\left\|\beta_{a,b}-\gamma^N_{\mu_{a,b}}\right\|_\infty\leq\epsilon_N$. Then the culture $\boldsymbol{\mathcal{B}}$ is regular for F.

To understand the difference between Propositions 2 and 3, note that Proposition 2 required each voter to believe that there is a clear asymmetry in each two-way race: $\mu_{a,b}$ must be bounded away from zero for every pair of distinct alternatives $a,b\in\mathcal{A}$. (Indeed, if $\mu_0=0$ in Example 1, then $\boldsymbol{\mathcal{B}}$ might *not* be regular for F. For instance, suppose that, for all $N\in\mathbb{N}$, there is some $\beta\in\mathcal{B}_N$ with $\beta_{a,b}=0$; then $\tau(\mathcal{B}_N,F_N)$ will decay to zero no faster than $\frac{1}{\sqrt{N}}$ as $N\to\infty$.) Proposition 3 relaxes this assumption; a voter could regard all two-way races as perfectly symmetric (i.e. $\mu_{a,b}$ could be zero for all $a,b\in\mathcal{A}$) without jeopardizing regularity. Proposition 3 also relaxes the unimodality assumption, but it assumes that σ_N grows *super*-linearly as $N\to\infty$ (as in condition (4)), rather than *sub*-linearly (as in condition (1.2)). This could occur, for example, if a voter believed that the other voters were highly correlated due to "information cascades" or "herding behaviour".

Anonymous cultures. A belief $\beta\in\Delta(\mathcal{V}^N)$ is *anonymous* if, for any permutation $\sigma:[1\ldots N]\longrightarrow[1\ldots N]$, and any profile $\mathbf{v}\in\mathcal{V}^N$, we have $\beta(\mathbf{v})=\beta(\sigma_*(\mathbf{v}))$, where we define the profile $\sigma_*(\mathbf{v}):=(v'_n)_{n=1}^N$ by setting $v'_n:=v_{\sigma(n)}$ for all $n\in[1\ldots N]$. In effect, this means that β does not identify any specific voting behaviour with any specific voter; it only provides aggregate information about the *number* of voters who are likely to deploy a particular voting behaviour. This is realistic in a large population, where a voter cannot be expected to have individual-specific beliefs about every other voter.

In our results, we will not need to assume that voters have anonymous beliefs. However, we *could have* made this assumption without loss of generality, in a sense we now explain. Let $\mathcal{S}_N$ be the set of all permutations of $[1 \ldots N]$; this set contains $N!$ elements. For any belief $\beta \in \Delta(\mathcal{V}^N)$, its *anonymization* is the belief $\beta^* \in \Delta(\mathcal{V}^N)$ defined as follows:

$$\beta^*(\mathbf{v}) \quad := \quad \frac{1}{N!} \sum_{\sigma \in \mathcal{S}_N} \beta[\sigma_*(\mathbf{v})], \qquad \text{for all } \mathbf{v} \in \mathcal{V}^N.$$

It is easy to verify two things: first, β^* is an anonymous belief, and second, $\tau(\beta^*, F_N) = \tau(\beta, F_N)$ (because F_N is also invariant under all permutations in $\mathcal{S}_N$). Now, given any culture $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)_{N=1}^{\infty}$, we define its *anonymization* to be the culture $\boldsymbol{\mathcal{B}}^* = (\mathcal{B}_N^*)_{N=1}^{\infty}$, where, for all $N \in \mathbb{N}$, we set $\mathcal{B}_N^* := \{\beta^*; \;\; \beta \in \mathcal{B}_N\}$. It follows that $\tau(\mathcal{B}_N^*, F_N) = \tau(\mathcal{B}_N, F_N)$ for all $N \in \mathbb{N}$ (because the rule F is anonymous, by definition). Thus, the culture $\boldsymbol{\mathcal{B}}^*$ is regular for F if and only if $\boldsymbol{\mathcal{B}}$ is regular for F. Thus, in all the results which follow, we could have assumed without loss of generality that we were dealing with anonymous cultures. But we will not require this assumption.

5 Truth-revealing ordinal voting rules

Consider an election held under an ordinal voting rule F. As usual, we assume that each voter is endowed with a vNM utility function in $\mathcal{U}$, which determines a preference order over $\mathcal{A}$ in the obvious way. Each voter knows her own utility function, but not those of other voters. The mechanism designer does not know any of their utility functions. The designer's problem is that voters might not disclose their true ordinal preferences when participating in F. We will now design a stochastic ordinal voting rule $\widetilde{F}$ which is asymptotically equal to F, but which is also *asymptotically ordinally truth-revealing*. Roughly speaking, this means that, according any beliefs that a designer could entertain about the utility functions of the voters, it is highly likely that any voter in a sufficiently large population will find optimal to reveal her true preferences over $\mathcal{A}$, regardless of her beliefs about the other voters.

To be precise, let $\rho \in \Delta(\mathcal{U})$ be a probability measure, describing the beliefs of the designer. This designer does not know the true vNM utility function of the voters, so she regards them as ρ-random variables (not necessarily independent). For simplicity, we will say they are *ρ-random voters*. Let $\widetilde{F}$ be a stochastic ordinal voting rule, and let $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)_{N=1}^{\infty}$ be a culture. For any $N \in \mathbb{N}$ and any $\beta \in \mathcal{B}_N$, let $\mathrm{Tr}(\beta, \rho, \widetilde{F}_N)$ be the ρ-probability that a ρ-random voter will find it optimal, in sense of maximizing expected utility, to reveal her true preference order over $\mathcal{A}$ in the voting rule $\widetilde{F}$. Finally, let $\mathrm{Tr}(\mathcal{B}_N, \rho, \widetilde{F}_N) := \inf_{\beta \in \mathcal{B}_N} \mathrm{Tr}(\beta, \rho, \widetilde{F}_N)$.

Definition. For any culture $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)_{N=1}^{\infty}$, the rule $\widetilde{F}$ is *asymptotically ordinally truth-revealing* for $\boldsymbol{\mathcal{B}}$ if $\lim_{N \to \infty} \mathrm{Tr}(\mathcal{B}_N, \rho, \widetilde{F}_N) = 1$ for all $\rho \in \Delta(\mathcal{U})$.

Note that the probability $\mathrm{Tr}(\beta, \rho, \widetilde{F}_N)$ describes the beliefs of a mechanism designer (not a voter), since it is computed using the *designer's* probability distribution ρ. If the rule $\widetilde{F}$ is asymptotically ordinally truth-revealing, then any designer will believe that each voter in a large enough population will, with very high probability, find it optimal to reveal her true ordinal preferences, regardless of her beliefs about the other voters. Thus, with very high probability, most of the voters in a large population will vote honestly. A small number of voters might vote dishonestly (either because they are irrational or because this is actually their optimal strategy), but this small number is unlikely to be enough to change the outcome of the vote. Our first main result says that, for any regular culture, *any* ordinal voting rule can be asymptotically approximated by a stochastic ordinal voting rule which is asymptotically ordinally truth-revealing.

Theorem 4 *Let F be any ordinal voting rule. Let $\boldsymbol{\mathcal{B}}$ be any regular culture for F. Then there is a stochastic ordinal voting rule $\widetilde{F}$ which is asymptotically equal to F, and which is asymptotically ordinally truth-revealing with respect to $\boldsymbol{\mathcal{B}}$.*

The rule $\widetilde{F}$ in Theorem 4 works roughly as follows. With a very high probability, $\widetilde{F}_N$ yields exactly the same outcome as F_N. However, with a tiny probability q_N, the rule $\widetilde{F}_N$ instead selects a random voter n and two random alternatives a and b, and makes n the "dictator" in the choice between a and b. If voter n stated that she prefers a over b, then a is chosen; otherwise, b is chosen. Obviously, such a "random dictatorship" will likely produce a socially suboptimal outcome. But since q_N is tiny (and becomes smaller as N gets large), the probability of such a suboptimal outcome occuring is very small; with very high probability (i.e. $1 - q_N$), the rule $\widetilde{F}_N$ will agree with F_N. Nevertheless, the tiny possibility that she *might* be the random dictator is enough to incentivize voter n to express her true ordinal preferences. The reason is that her optimal voting strategy is determined only by the cases where her vote could make a difference: namely, the case where the profile is nearly tied, and the case where she is the random dictator. If N is large, then the probability that n is chosen as a random dictator, although tiny, is still *much larger* than the probability of a nearly-tied profile according to her beliefs (as specified by the culture $\boldsymbol{\mathcal{B}}$). Thus, n's optimal strategy is driven by the "random dictatorship" case (where it is best for her to be honest), rather than the "nearly tied" case (where it might be optimal to be dishonest).

6 Truth-revealing approval voting rules

Now that we have shown how to elicit honesty ordinal voting rules, we will apply a similar technique to approval voting. But first we need some notation. Let $\mathcal{V}_0 := \{v : \mathcal{A} \longrightarrow \{0,1\};\ v(a) = 0 \text{ and } v(b) = 1 \text{ for some } a, b \in \mathcal{A}\}$ be the set of all "binary utility functions". A *binary voting rule* is a sequence $F = (F_N)_{N=1}^{\infty}$, where $F_N : \mathcal{V}_0^N \longrightarrow \mathcal{A}$ for all $N \in \mathbb{N}$. Likewise, a *stochastic binary voting rule* is a sequence $\widetilde{F} = (\widetilde{F}_N)_{N=1}^{\infty}$, where $\widetilde{F}_N : \mathcal{V}_0^N \longrightarrow \widetilde{\mathcal{A}}$, for all $N \in \mathbb{N}$. The most well-known binary voting rule is the *approval voting* rule, which simply selects the alternative a having the highest approval score $\sum_{n=1}^{N} v_n(a)$, with some tie-breaking rule in the event of a tie (Brams and Fishburn, 1983). There is no universally agreed definition of "honesty" in approval voting, beyond the minimal criterion that $v(a) \geq v(b)$ whenever the voter prefers a over b. We will adopt the following criterion. Given a cardinal utility function $u \in \mathcal{U}$, a binary signal $v \in \mathcal{V}_0$ is *truthful* for u if it endorses only alternatives whose utilities are above average, according to u. Formally:

$$\text{For all } a \in \mathcal{A}, \quad \Big(u(a) > \overline{u}\Big) \implies \Big(v(a) = 1\Big) \qquad (5)$$
$$\text{while} \quad \Big(u(a) < \overline{u}\Big) \implies \Big(v(a) = 0\Big), \qquad \text{where} \quad \overline{u} \quad := \quad \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} u(a).$$

Note that, if $u(a) = \overline{u}$, then both $v(a) = 0$ and $v(a) = 1$ are considered truthful.

The definition of sincerity under approval voting is far from obvious, because there is no natural way of transforming each possible utility function for a voter into a unique signal in $\mathcal{V}$, as with most other voting rules such as plurality or Borda. Merill and Nagel (1987) discuss several notions of sincerity in approval voting (see Núñez (2014) for a strategic analysis of these concepts in a Poisson game). Among these definitions, the least restrictive one is *no-skipping sincerity*, which requires that if a voter approves of some alternative x, then he also approves of all the alternatives that he prefers to x. All the definitions suggested by Merill and Nagel (1987) respect no-skipping sincerity, and so does ours. In particular, the one we have chosen (approving of all the alternatives better than the average utility) is called *Pure Sincerity* by Merill and Nagel. This choice is made only for simplicity and illustration purposes. By a slight modification of the design, we can also construct asymptotically truth-revealing voting rules which implement of Expansive and Restrictive Sincerity.[8]

Let ρ be the probability distribution over $\mathcal{U}$ describing the designer's beliefs about the utility functions of the voters. Our goal is to construct a stochastic binary voting rule such that it is highly probable (according to ρ) that any voter will find optimal to vote truthfully in the sense of definition (5), regardless of her beliefs about the other voters.

Let $\widetilde{F}$ be a stochastic binary voting rule, and let $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)_{N=1}^{\infty}$ be a culture. Let $\rho \in \Delta(\mathcal{U})$ be a probability distribution (representing the possible beliefs of a mechanism designer). For any $N \in \mathbb{N}$, and any $\beta \in \mathcal{B}_N$, let $\mathrm{Tr}(\beta, \rho, \widetilde{F}_N)$ be the ρ-probability that a ρ-random voter with beliefs β will find it optimal (in the sense of maximizing expected utility) to vote truthfully in the voting rule $\widetilde{F}$, given that her vNM utility function is a ρ-random variable. Let $\mathrm{Tr}(\mathcal{B}_N, \rho, \widetilde{F}_N) := \inf_{\beta \in \mathcal{B}_N} \mathrm{Tr}(\beta, \rho, \widetilde{F}_N)$.

Definition. For any culture $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)_{N=1}^{\infty}$, the rule $\widetilde{F}$ is *asymptotically binarily truth-revealing* for $\boldsymbol{\mathcal{B}}$ if $\lim_{N \to \infty} \mathrm{Tr}(\mathcal{B}_N, \rho, \widetilde{F}_N) = 1$ for all $\rho \in \Delta(\mathcal{U})$.

[8] In Merill and Nagel's terminology, *Expansive* (resp. *Restrictive*) *Sincerity* requires a voter to approve of a strict superset (resp. subset) of the Pure Sincerity set, while still satisfying the no-skipping condition.

Our next result says: for any regular culture, we can asymptotically approximate approval voting by a stochastic binary voting rule which is asymptotically binarily truth-revealing.

Theorem 5 *Let $\boldsymbol{\mathcal{B}}$ be any regular culture for approval voting. There is a stochastic binary voting rule $\widetilde{F}$ which is asymptotically equal to approval voting, and which is asymptotically binarily truth-revealing for $\boldsymbol{\mathcal{B}}$.*

7 Truth-revealing cardinal voting rules

Let F be a cardinal voting rule. We want to design a stochastic cardinal voting rule which asymptotically approximates F, and such that for any voter in a large population, it is highly probable that this voter will find it optimal to reveal her true utility function, regardless of her beliefs about the other voters. First, we need some notation. Let $\widetilde{F}$ be a stochastic cardinal voting rule, and let $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)_{N=1}^{\infty}$ be a culture. Let $\rho \in \Delta(\mathcal{U})$ be a probability distribution (representing the beliefs of a mechanism designer). For any $N \in \mathbb{N}$, any belief $\beta \in \mathcal{B}_N$, and any $\epsilon > 0$ let $\mathrm{Tr}_\epsilon(\beta, \rho, \widetilde{F}_N)$ be the ρ-probability that a voter with beliefs β and a ρ-random utility function $\widetilde{u}$ will find it optimal (in the sense that it maximizes her expected utility) to declare a vNM function u' such that $\|\widetilde{u} - u'\|_\infty < \epsilon$. Let $\mathrm{Tr}_\epsilon(\mathcal{B}_N, \rho, \widetilde{F}_N) := \inf_{\beta \in \mathcal{B}_N} \mathrm{Tr}_\epsilon(\beta, \rho, \widetilde{F}_N)$.

Definition. For any culture $\boldsymbol{\mathcal{B}}$ and $\rho \in \Delta(\mathcal{U})$, the rule $\widetilde{F}$ is *asymptotically cardinally truth-revealing* with respect to $\boldsymbol{\mathcal{B}}$ and ρ if $\lim_{N\to\infty} \mathrm{Tr}_\epsilon(\mathcal{B}_N, \rho, \widetilde{F}_N) = 1$ for all $\epsilon > 0$.

If $\widetilde{F}$ is asymptotically cardinally truth-revealing, then the designer believes that any voter in a large enough population will, with very high probability, find it optimal to reveal something very close to her true vNM utility function, regardless of her beliefs about the other voters. Thus, with very high probability, most of the voters in a large population will vote honestly (modulo an ϵ-sized error, which can be made arbitrarily small).

We will also need to introduce a weakened notion of "asymptotic equality". Let $F = (F_N)_{N=1}^{\infty}$ and $G = (G_N)_{N=1}^{\infty}$ be two cardinal voting rules. For any $\epsilon > 0$ and $N \in \mathbb{N}$, we will say that G_N is *ϵ-similar* to F_N if, for any utility profile $\mathbf{u} \in \mathcal{U}^N$, there is some profile $\mathbf{u}' \in \mathcal{U}^N$ with $\|\mathbf{u}' - \mathbf{u}\|_\infty < \epsilon$ such that $G_N(\mathbf{u}) = F_N(\mathbf{u}')$. In other words, applying the rule G_N to the profile $\mathbf{u}$ is equivalent to applying F_N to some profile $\mathbf{u}'$ which is an "ϵ-perturbation" of $\mathbf{u}$. Heuristically, this means that F_N and G_N will agree "most of the time" (i.e. unless the profile $\mathbf{u}$ is ϵ-close to a boundary between two outcomes). We will say that G is *asymptotically similar* to F if, for any $\epsilon > 0$, there is some $N_\epsilon \in \mathbb{N}$ such that G_N is ϵ-similar to F_N for all $N \geq N_\epsilon$. This means that, for large populations, F and G will almost always produce the same outcome. (Indeed, in practice, the utility functions of the voters can only be recorded with finite precision anyways; thus, in a large enough population, it would be practically impossible to distinguish between F and G.) Finally, let $\widetilde{F} = (\widetilde{F}_N)_{N=1}^{\infty}$ be a stochastic cardinal voting rule. We say that $\widetilde{F}$ is *asymptotically similar to* F if there is some (deterministic) cardinal voting rule $G = (G)_{N=1}^{\infty}$ such that $\widetilde{F}$ is asymptotically equal to G, and G is asymptotically similar to F.

We will also need to slightly strengthen the regularity requirement on the culture. Let $\epsilon > 0$, and let $\widetilde{\mathbf{v}}, \widetilde{\mathbf{u}}$ be two random utility profiles in $\mathcal{U}^N$; we will say that $\widetilde{\mathbf{v}}$ is an *ϵ-distortion* of $\widetilde{\mathbf{u}}$ if $\|\widetilde{\mathbf{v}} - \widetilde{\mathbf{u}}\|_\infty < \epsilon$ almost surely. Given two probabilistic beliefs $\beta, \beta' \in \Delta(\mathcal{U}^N)$, we say that β and β' are *ϵ-close* if β is the probability distribution of a random utility profile $\widetilde{\mathbf{u}}$, and β' is the probability distribution of another random utility profile $\widetilde{\mathbf{v}}$ which is an ϵ-distortion of $\widetilde{\mathbf{u}}$. Finally, let $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)_{N=1}^{\infty}$ be a culture. For any $\epsilon > 0$ and any $N \in \mathbb{N}$, we define

$$\mathcal{B}_N^{(\epsilon)} \quad := \quad \left\{\beta' \in \Delta(\mathcal{U}^N)\ ;\ \beta' \text{ is } \epsilon\text{-close to some } \beta \in \mathcal{B}_N\right\}. \tag{6}$$

In other words, $\mathcal{B}_N^{(\epsilon)}$ is the set of all probabilistic beliefs which a voter could have about the voting behaviour of the other voters, if her belief about the *true* profile of utility functions was some $\beta \in \mathcal{B}_N$, but she also believed that the profile of utility functions *actually* reported would be some ϵ-distortion of the true profile of utility functions. Finally, define $\boldsymbol{\mathcal{B}}^{(\epsilon)} := (\mathcal{B}_N^{(\epsilon)})_{N=1}^{\infty}$. We will say that the culture $\boldsymbol{\mathcal{B}}$ is *robustly regular* there is some $\epsilon > 0$ such that the culture $\boldsymbol{\mathcal{B}}^{(\epsilon)}$ is regular for the voting rule F. This means that, if any voter in a large population expects that the profile of utility functions actually reported by the other voters will be an ϵ-distortion of the true profile of utility functions, then she will estimate the probability of a near-tie to be extremely small. Note that any robustly regular culture is regular (because $\mathcal{B}_N \subseteq \mathcal{B}_N^{(\epsilon)}$ for all $N \in \mathbb{N}$ and $\epsilon > 0$). Our next result says that for any beliefs ρ and any robustly regular culture, *any* cardinal voting rule can be approximated by an asymptotically similar stochastic cardinal voting rule which is asymptotically cardinally truth-revealing.

Theorem 6 *Let F be any cardinal voting rule, let ρ be any probability distribution on $\mathcal{U}$, and let $\boldsymbol{\mathcal{B}}$ be any robustly regular culture for F. Then there is a stochastic cardinal voting rule $\widetilde{F}$ which is asymptotically similar to F, and which is asymptotically cardinally truth-revealing with respect to $\boldsymbol{\mathcal{B}}$ and ρ.*

Remark: We could extend Theorem 6 to the case where ρ is replaced by some finite (or even countable) set of priors.

The rule in Theorem 6 is somewhat more complicated than the rule in Theorems 4 and 5. For any $\epsilon > 0$, we first construct a stochastic cardinal voting rule $\widetilde{F}^\epsilon$ which is "asymptotically ϵ-similar" to F and "asymptotically ϵ-truth revealing", in the sense that it incentivizes each voter to reveal something within ϵ of her true utility function. To construct $\widetilde{F}^\epsilon$, we approximate the set $\Delta(\mathcal{A})$ of all lotteries over $\mathcal{A}$ with a finite set $\mathcal{A}_\epsilon$, consisting of all lotteries involving exactly two alternatives and probabilities which are integer multiples of ϵ.[9] A voter's vNM utility function on $\mathcal{A}$ is determined, up to an ϵ-sized error, by her ordinal preferences over $\mathcal{A}_\epsilon$. We can then elicit the voters' ordinal preferences over $\mathcal{A}_\epsilon$ by constructing an asymptotically *ordinally* truth-revealing voting rule $\widetilde{F}^\epsilon = (\widetilde{F}^\epsilon_N)_{N=1}^\infty$, using Theorem 4. Finally, let $\{\epsilon_m\}_{m=1}^\infty$ be a sequence tending to zero; we define the stochastic rule $\widetilde{F} = (\widetilde{F}_N)_{N=1}^\infty$ by constructing a suitable diagonal sequence of the two-dimensional array $(\widetilde{F}^{\epsilon_m}_N)_{N,m=1}^\infty$.

8 Asymptotic Bayesian Nash implementation

Theorems 4, 5, and 6 assumed that the voters' beliefs were drawn from a "culture" $\boldsymbol{\mathcal{B}}$ but they did not endogenize these beliefs as part of an equilibrium; thus, they were not standard implementation results. However, we will now show how special cases of these theorems yield truth-revealing Bayesian Nash equilibrium implementations in large populations. To do so, we must first build a Bayesian voting game in our framework.

Types. We will assume that each voter has a *type*, which is known only to her, and which determines both her utility function and her beliefs about the other voters. A voter of type t will simply be called a *t-voter*. For all $N \in \mathbb{N}$, and all $m \in [1 \ldots N]$, let $\mathcal{T}_m^N$ be the set of possible types for the mth voter in a population of size N. These sets could be finite or infinite. For simplicity, we will assume that the type-sets of different voters are disjoint —i.e. $\mathcal{T}_n^N \cap \mathcal{T}_m^N = \emptyset$ for any distinct $n, m \in [1 \ldots N]$. (This simplifies notation but does not alter our results.) For all $m \in [1 \ldots N]$, and all $t \in \mathcal{T}_m^N$, let $u_t \in \mathcal{U}$ be the vNM utility function of a t-voter.

Let $\boldsymbol{\mathcal{T}}^N := \prod_{m=1}^N \mathcal{T}_m^N$. An element $\mathbf{t} = (t_n)_{n=1}^N \in \boldsymbol{\mathcal{T}}^N$ will be called a *type profile* for a population of size N. Given any type profile $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$, we define the corresponding *utility profile* $\mathbf{u_t} := (u_{t_n})_{n=1}^N \in \mathcal{U}^N$.

Beliefs. Let $N \in \mathbb{N}$. For any voter $m \in [1 \ldots N]$, the set

$$\boldsymbol{\mathcal{T}}_{-m}^N := \prod_{\substack{n=1 \\ n \neq m}}^N \mathcal{T}_n^N$$

represents the set of possible type-profiles of all the *other* voters. For any $t \in \mathcal{T}_m^N$, let $\pi_t \in \Delta(\boldsymbol{\mathcal{T}}_{-m}^N)$, be a probability distribution, representing the beliefs of a t-voter about the possible types of all the other voters. The data $\mathbf{s}^N := \left[(t, u_t, \pi_t);\ m \in [1 \ldots N] \text{ and } t \in \mathcal{T}_m^N\right]$ will be called a *community* of size N. The sequence $\mathbf{S} := (\mathbf{s}^N)_{N=1}^\infty$ will be called a *society*.[10]

[9] For example, if $\epsilon = \frac{1}{5}$, then these would be lotteries over some pair $\{a, b\}$ of alternatives, where $p(a) \in \{0, \frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}, 1\}$, and $p(b) = 1 - p(a)$.

[10] In most Bayesian game models, all players share a common prior probability $\pi \in \Delta(\boldsymbol{\mathcal{T}}^N)$, and the type-$t$ belief π_t is obtained by Bayesian updating of π conditional on t. But we do not need to assume this. Also, since we are interested only in anonymous voting rules on very large populations, it would be reasonable to suppose that the community $\mathbf{s}^N$ is invariant under all permutations of $[1 \ldots N]$, so that all voters are, in effect, *ex ante* indistinguishible. But we do not need to assume this, either.

Bayesian Game. Let $\widetilde{F} = (\widetilde{F}_N)_{N=1}^\infty$ be a stochastic voting rule using some set $\mathcal{V}$ of signals. Thus, for all $N \in \mathbb{N}$, we have a function $\widetilde{F}_N : \mathcal{V}^N \longrightarrow \widetilde{\mathcal{A}}$. The ordered pair $G_N := (\widetilde{F}_N, \mathbf{s}^N)$ defines an N-player *Bayesian game*: each voter's strategy space is $\mathcal{V}$, the outcome is determined by applying $\widetilde{F}_N$ to obtain a random alternative in $\widetilde{\mathcal{A}}$, and the possible types, beliefs, and utility functions the voters are determined by $\mathbf{s}^N$. For any $m \in [1 \ldots N]$, a (pure) *voting strategy* for voter m in G_N is a function $V_m : \mathcal{T}_m^N \longrightarrow \mathcal{V}$. A *strategy profile* in G_N is an N-tuple $\mathbf{V} = (V_n)_{n=1}^N$. Given any strategy profile $\mathbf{V}$ and type profile $\mathbf{t} = (t_n)_{n=1}^N \in \boldsymbol{\mathcal{T}}^N$, we obtain a vote profile $\mathbf{V}(\mathbf{t}) := (v_n)_{n=1}^N$ by setting $v_n := V_n(t_n)$ for all $n \in [1 \ldots N]$.

Given a strategy profile $\mathbf{V}$, each voter-type can compute the probability that any particular profile of votes will occur. For any voter $m \in [1 \ldots N]$ and type $t_m \in \mathcal{T}_m^N$, let $p_{t_m} \in \Delta(\mathcal{V}^N)$ describe the beliefs of a t_m-voter about the vote profile that will occur given $\mathbf{V}$. Formally, for any possible vote profile $\mathbf{v} \in \mathcal{V}^N$, if $v_m = \mathbf{V}_m(t_m)$, then we define

$$p_{t_m}(\mathbf{v} \mid \mathbf{V}) \quad := \quad \pi_{t_m}\Big\{\mathbf{t}_{-m} \in \boldsymbol{\mathcal{T}}_{-m}^N \ ; \ \mathbf{V}_n(t_n) = v_n, \text{ for all } n \in [1 \ldots N] \setminus \{m\}\Big\}, \tag{7}$$

whereas $p_{t_m}(\mathbf{v} \mid \mathbf{V}) := 0$ if $v_m \neq \mathbf{V}_m(t_m)$. In other words, $p_{t_m}(\cdot \mid \mathbf{V})$ represents the beliefs of a t_m-voter that the vote profile equals $\mathbf{v}$ given her beliefs π_{t_m} over the types of the rest of the voters and the strategy profile $\mathbf{V}$. For short, we write simply $p_{t_m}(\mathbf{v})$.

Equilibrium. A strategy profile $\mathbf{V}$ is a (pure strategy) *Bayesian Nash Equilibrium* for the game G_N if, for all $m \in [1 \ldots N]$ and all types $t_m \in \mathcal{T}_m^N$, the vote $V_m(t_m)$ is type t_m's best response, in the sense that it maximizes the expected value of u_{t_m} given the beliefs p_{t_m} defined by (7). In other words: each voter-type's strategy is optimal given her beliefs, while at the same time, her beliefs correctly account for the strategies of the other voter-types. For all $N \in \mathbb{N}$, let $\mathbf{V}^N$ be a strategy profile for $\widetilde{F}_N$. The sequence $\mathbb{V} := (\mathbf{V}^N)_{N=1}^\infty$ is an *eventual Bayesian Nash equilibrium* for the sequence $(G_N)_{N=1}^\infty$ if there is some $N_0 \in \mathbb{N}$ such that, for all $N \geq N_0$, the strategy profile $\mathbf{V}^N$ is a Bayesian Nash equilibrium of the game G_N.

Asymptotic implementation. Let $F = (F_N)_{N=1}^\infty$ be a voting rule, and let $\widetilde{F} = (\widetilde{F}_N)_{N=1}^\infty$ be a stochastic voting rule. We will now define what it means for $\widetilde{F}$ to implement F in Bayesian Nash equilibrium, for sufficiently large populations. The definitions for ordinal voting rules and cardinal voting rules are slightly different, so we treat them separately.

First, suppose F is a cardinal voting rule. Let $N \in \mathbb{N}$, and fix a community $\mathbf{s}$ of size N. For any strategy profile $\mathbf{V}^N$ for $\widetilde{F}_N$, and any type profile $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$, let $P_{\mathbf{t}}(\widetilde{F}_N, F_N, \mathbf{s}, \mathbf{V}^N)$ be the probability that $\widetilde{F}_N[\mathbf{V}^N(\mathbf{t})] = F_N(\mathbf{u_t})$.

Now suppose F is an ordinal voting rule. Thus, $\mathcal{V}$ is the set of all possible preference orders over $\mathcal{A}$. Let $N \in \mathbb{N}$, and fix a community $\mathbf{s}$ of size N. For any type profile $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$, let $\mathbf{v}_{\mathbf{t}}^* \in \mathcal{V}^N$ be the profile of preference orders defined by the utility profile $\mathbf{u_t}$. For any strategy profile $\mathbf{V}^N$ for $\widetilde{F}_N$, we now define $P_{\mathbf{t}}(\widetilde{F}_N, F_N, \mathbf{s}, \mathbf{V}^N)$ to be the probability that $\widetilde{F}_N[\mathbf{V}^N(\mathbf{t})] = F_N(\mathbf{v}_{\mathbf{t}}^*)$.

Finally, when F is either ordinal or cardinal, and $\mathbf{V}^N$ is a strategy profile for $\widetilde{F}_N$, we define

$$P(\widetilde{F}_N, F_N, \mathbf{s}, \mathbf{V}^N) \quad := \quad \inf_{\mathbf{t} \in \boldsymbol{\mathcal{T}}^N} P_{\mathbf{t}}(\widetilde{F}_N, F_N, \mathbf{s}, \mathbf{V}^N). \tag{8}$$

Definition. Let $\mathbf{S} = (\mathbf{s}^N)_{N=1}^\infty$ be a society. The rule $\widetilde{F}$ *asymptotically implements F in Bayesian Nash equilibrium* for the society $\mathbf{S}$ if there is an eventual Bayesian Nash equilibrium $\mathbb{V} = (\mathbf{V}^N)_{N=1}^\infty$ for the pair $(\widetilde{F}, \mathbf{S})$ such that $\lim_{N \to \infty} P(\widetilde{F}_N, F_N, \mathbf{s}^N, \mathbf{V}^N) = 1$.

First, we will give a sufficient condition for the asymptotic implementation of ordinal voting rules; then we will turn to cardinal rules. For any voter $m \in [1 \ldots N]$ and any possible type $t_m \in \mathcal{T}_m^N$, let $v_{t_m}^* \in \mathcal{V}$ denote the preference order induced by the utility function u_{t_m}. Then define the probability distribution $\beta_{t_m} \in \Delta(\mathcal{V}^N)$, as follows: for any other preference profile $\mathbf{v} = (v_n)_{n=1}^N \in \mathcal{V}^N$

$$\beta_{t_m}(\mathbf{v}) \ := \ \begin{cases} \pi_{t_m}\Big\{\mathbf{t} \in \boldsymbol{\mathcal{T}}_{-m}^N \ ; \ v_{t_n}^* = v_n \text{ for all } n \in [1 \ldots N] \setminus \{m\}\Big\} & \text{if } \ v_{t_m}^* = v_m; \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

In other words, β_{t_m} is the probabilistic beliefs of a type-t_m voter about the true preference profile of the whole population (including herself). For all $N \in \mathbb{N}$, let

$$\mathcal{B}_N \quad := \quad \left\{\beta_t \, ; \; t \in \mathcal{T}_m^N \text{ for some } m \in [1 \ldots N]\right\} \quad \subseteq \quad \Delta(\mathcal{V}^N). \tag{10}$$

Then define $\boldsymbol{\mathcal{B}}_{\mathbf{S}} := (\mathcal{B}_N)_{N=1}^{\infty}$. In other words, $\boldsymbol{\mathcal{B}}_{\mathbf{S}}$ is the (ordinal) culture determined by all possible beliefs which could be held by any voter of any type in the society $\mathbf{S}$, about the true preference profile of the other voters.

Meanwhile, let $\mathcal{U}_{\mathbf{S}} := \{u_t; \;\; t \in \mathcal{T}_m^N$ for some $N \in \mathbb{N}$ and some $m \in [1 \ldots N]\}$. In other words, $\mathcal{U}_{\mathbf{S}}$ is the set of all possible vNM utility functions for any voter of any type, in any size of population. For any utility function $u \in \mathcal{U}_{\mathbf{S}}$, define $\gamma(u) := \min\{|u(a) - u(b)|; \;\; a, b \in \mathcal{A}$ and $u(a) \neq u(b)\}$. We say that the society $\mathbf{S}$ is F-*regular* if the culture $\boldsymbol{\mathcal{B}}_{\mathbf{S}}$ is regular for F, and there exists some $\epsilon > 0$ such that, for any $u \in \mathcal{U}_{\mathbf{S}}$, we have $\gamma(u_t) > \epsilon$. For example, if the set of voter types is finite (a common assumption in the literature) and all voter types have strict preferences, then the condition on γ is automatically true.

Theorem 7 *Let* $\mathbf{S}$ *be a society. Let* F *be an ordinal voting rule, and let* $\widetilde{F}$ *be the stochastic voting rule from Theorem 4. If* $\mathbf{S}$ *is* F*-regular, then* $\widetilde{F}$ *asymptotically implements* F *in Bayesian Nash equilibrium for the society* $\mathbf{S}$*.*

We will now give a sufficient condition for the asymptotic implementation of a cardinal voting rule. Fix a society $\mathbf{S}$. For any voter $m \in [1 \ldots N]$ and any possible type $t_m \in \mathcal{T}_m^N$, define the probability distribution $\beta_{t_m} \in \Delta(\mathcal{U}^N)$, as follows: for any utility profile $\mathbf{w} \in \mathcal{U}^N$

$$\beta_{t_m}(\mathbf{w}) \;:=\; \begin{cases} \pi_{t_m}\left\{\mathbf{t} \in \boldsymbol{\mathcal{T}}_{-m}^N \, ; \; u_{t_n} = w_n \text{ for all } n \in [1 \ldots N] \setminus \{m\}\right\} & \text{if} \quad w_m = u_{t_m}; \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

In other words, β_{t_m} is the probabilistic beliefs of a type-t_m voter about the true utility profile of the whole population (including herself). For all $N \in \mathbb{N}$, let

$$\mathcal{B}_N \quad := \quad \left\{\beta_t \, ; \; t \in \mathcal{T}_m^N \text{ for some } m \in [1 \ldots N]\right\} \quad \subseteq \quad \Delta(\mathcal{U}^N). \tag{12}$$

Then define $\boldsymbol{\mathcal{B}}_{\mathbf{S}} := (\mathcal{B}_N)_{N=1}^{\infty}$. In other words, $\boldsymbol{\mathcal{B}}_{\mathbf{S}}$ is the culture determined by all possible beliefs which could be held by any voter of any type in the society $\mathbf{S}$, about the true utility profile of the other voters. We will say that the society $\mathbf{S}$ is F-*regular* if:

(R1) $\mathcal{U}_{\mathbf{S}}$ is finite;

(R2) The culture $\boldsymbol{\mathcal{B}}_{\mathbf{S}}$ is robustly regular for the voting rule F; and

(R3) There is some $\epsilon > 0$ and $N_0 \in \mathbb{N}$ such that, for any $N \geq N_0$ and any type profile $\mathbf{t} \in \mathcal{T}^N$, we have $F_N(\mathbf{v}) = F_N(\mathbf{u_t})$ for all $\mathbf{v} \in \mathcal{U}^N$ with $\|\mathbf{v} - \mathbf{u_t}\|_{\infty} < \epsilon$.

The meanings of (R1) and (R2) are clear. Condition (R3) says that, in a large population, the outcome of the voting rule will be robust under ϵ-distortions, for any type profile.

Theorem 8 *For any cardinal voting rule* F*, and* F*-regular society* $\mathbf{S}$*, there is a stochastic voting rule that asymptotically implements* F *in Bayesian Nash equilibrium for* $\mathbf{S}$*.*

Through similar techniques, we can state and prove an asymptotic Bayesian Nash implementation result for approval voting, using the definition of "truthful" given in Section 6. We leave the details to the reader.

Appendix: Proofs

Proof of Proposition 2. We are interested in the asymptotic behaviour of $N \cdot \tau(\mathcal{B}_N, F_N)$ as $N \to \infty$. So we can assume without loss of generality that $N \geq \frac{1}{\mu_0}$. Let $\beta \in \mathcal{B}_N$; thus, β is a probability distribution over $\mathcal{V}^N$, the set of N-voter profiles. Let $\widetilde{\mathbf{s}} := (\widetilde{s}_a^n; \;\; n \in [1 \ldots N]$ and $a \in \mathcal{A})$ be a β-random profile of N voters. For all $a \in \mathcal{A}$, let $\widetilde{S}_a := \sum_{n=1}^{N} \widetilde{s}_a^n$ be the total score of alternative a in this profile; this is a random number. Fix some alternatives $a, b \in \mathcal{A}$. Without loss of generality, suppose $\mu_{a,b} > 0$ (the other case is analogous). Thus, $N\,\mu_{a,b} \geq N\,\mu_0 \geq 1$ (because $N \geq \frac{1}{\mu_0}$ by assumption). Let $\tau_{a,b}(\beta, F_N)$ be the β-probability that $\widetilde{\mathbf{s}}$ is a nearly-tied profile having a and b as its two top candidates. Then

$$\begin{aligned}
\tau_{a,b}(\beta,F_N) \quad &\leq \quad \text{Prob}\left[|\widetilde{S}_a - \widetilde{S}_b| \leq 1\right] \quad = \quad \int_{-1}^{1} \beta_{a,b}(x)\ \mathrm{d}x \qquad \text{(A1)}\\
&\leq \quad 2\left\|\beta_{a,b} - \gamma^N_{\mu_{a,b}}\right\|_\infty \; + \; \int_{-1}^{1} \gamma^N_{\mu_{a,b}}(x)\ \mathrm{d}x \\
&\leq \quad 2\epsilon_N \; + \; \int_{-1}^{1} \gamma^N_{\mu_{a,b}}(x)\ \mathrm{d}x \quad = \quad 2\epsilon_N \; + \; \frac{1}{\sigma_N}\int_{-1}^{1} \Gamma\left(\frac{x - N\,\mu_{a,b}}{\sigma_N}\right)\ \mathrm{d}x \\
&\leq \quad 2\epsilon_N \; + \; \frac{2}{\sigma_N} \sup_{x\in[-1,1]} \Gamma\left(\frac{x - N\,\mu_{a,b}}{\sigma_N}\right) \quad \underset{(*)}{\leq} \quad 2\epsilon_N \; + \; \frac{2}{\sigma_N}\Gamma\left(\frac{1 - N\,\mu_0}{\sigma_N}\right).
\end{aligned}$$

Here, $(*)$ is because Γ is unimodal with its mode at zero, and hence, nondecreasing on $(-\infty, 0]$, while $0 \geq 1 - N\,\mu_0 \geq x - N\,\mu_{a,b}$ for all $x \in [-1, 1]$, because $\mu_0 \leq \mu_{a,b}$ and $N \geq \frac{1}{\mu_0}$. Summing inequality (A1) over all pairs $\{a, b\} \subseteq \mathcal{A}$, we obtain

$$\tau(\beta, F_N) \quad \leq \quad \sum_{\substack{a,b\in\mathcal{A}\\ a\neq b}} \tau_{a,b}(\beta, F_N) \quad \leq \quad \frac{A(A-1)}{2}\left[2\epsilon_N + \frac{2}{\sigma_N}\Gamma\left(\frac{1-N\,\mu_0}{\sigma_N}\right)\right]. \qquad \text{(A2)}$$

Let $K := A(A - 1)$. Taking the supremum of inequality (A2) over all $\beta \in \mathcal{B}_N$, we obtain:

$$\tau(\mathcal{B}_N, F_N) \quad \leq \quad K\left(\epsilon_N + \frac{1}{\sigma_N}\Gamma\left(\frac{1-N\,\mu_0}{\sigma_N}\right)\right).$$

$$\begin{aligned}
\text{Thus,}\quad \lim_{N\to\infty} N\cdot\tau(\mathcal{B}_N, F_N) \quad &\leq \quad K \lim_{N\to\infty} N\epsilon_N \; + \; K \lim_{N\to\infty} \frac{N}{\sigma_N}\Gamma\left(\frac{1-N\,\mu_0}{\sigma_N}\right) \\
&\underset{\text{(a)}}{=} \quad 0 \; + \; K \lim_{N\to\infty}\left(\frac{N}{1-N\,\mu_0}\right)\left(\frac{1-N\,\mu_0}{\sigma_N}\right)\Gamma\left(\frac{1-N\,\mu_0}{\sigma_N}\right) \\
&= \quad K\left(\lim_{N\to\infty}\frac{N}{1-N\,\mu_0}\right)\cdot\left(\lim_{N\to\infty}\frac{1-N\,\mu_0}{\sigma_N}\,\Gamma\left(\frac{1-N\,\mu_0}{\sigma_N}\right)\right) \\
&\underset{\text{(b)}}{=} \quad \frac{-K}{\mu_0}\cdot\lim_{x\to-\infty} x\cdot\Gamma(x) \quad \underset{\text{(c)}}{=} \quad 0,
\end{aligned}$$

as desired. Here, (a) is by hypothesis (3). Meanwhile, (b) is the change of variables $x := \frac{1-N\,\mu_0}{\sigma_N}$, using condition (1.2) to obtain $\lim_{N\to\infty} \frac{1-N\,\mu_0}{\sigma_N} = -\infty$. Finally, (c) is because Γ is integrable, so that $\lim_{x\to-\infty} x\cdot\Gamma(x) = 0$, concluding the proof. □

Proof of Proposition 3. Let $G := \|\Gamma\|_\infty$; then $G < \infty$ by hypothesis. Let $\beta \in \mathcal{B}_N$, and let $a, b \in \mathcal{A}$. We have

$$\tau_{a,b}(\beta, F_N) \quad \underset{(*)}{\leq} \quad 2\epsilon_N \; + \; \frac{2}{\sigma_N}\sup_{x\in[-1,1]}\Gamma\left(\frac{x-N\,\mu_{a,b}}{\sigma_N}\right) \quad \leq \quad 2\epsilon_N \; + \; \frac{2G}{\sigma_N}.$$

where the proof of $(*)$ is exactly the same as the first few steps in inequality (A1). Thus, as in inequality (A2), we obtain

$$\tau(\beta, F_N) \quad \leq \quad \sum_{\substack{a,b\in\mathcal{A}\\ a\neq b}} \tau_{a,b}(\beta, F_N) \quad \leq \quad \frac{A(A-1)}{2}\left(2\epsilon_N + \frac{2G}{\sigma_N}\right). \qquad \text{(A3)}$$

Let $K := A(A - 1)$. Taking the supremum of inequality (A3) over all $\beta \in \mathcal{B}_N$, we obtain:

$$\tau(\mathcal{B}_N, F_N) \quad \leq \quad K\left(\epsilon_N + \frac{G}{\sigma_N}\right).$$

$$\text{Thus,}\quad \lim_{N\to\infty} N\cdot\tau(\mathcal{B}_N, F_N) \quad \leq \quad K\lim_{N\to\infty} N\epsilon_N \; + \; K\lim_{N\to\infty}\frac{NG}{\sigma_N} \quad \underset{(*)}{=} \quad 0 \; + \; 0 \quad = \quad 0,$$

as desired. Here, $(*)$ is by hypotheses (3) and (4). □

Proof of Theorem 4. For any utility function $u \in \mathcal{U}$, let $\gamma(u) := \min\{|u(a) - u(b)|;\ \ a, b \in \mathcal{A}$ and $u(a) \neq u(b)\}$. In other words, $\gamma(u)$ is the minimum utility *gap* between two nonindifferent alternatives.

For any preference order $p \in \mathcal{P}$, let $g(p) \in \Delta(\mathcal{A})$ be the probability distribution describing the outcome of the following random procedure:

1. Choose two distinct alternatives $a, b \in \mathcal{A}$ uniformly at random.
2. If $a \succ_p b$, then select a. Otherwise select b.

This defines a function $g : \mathcal{P} \longrightarrow \Delta(\mathcal{A})$. Our first claim says that g is "truth-revealing" in the following sense: if a single voter is told that the outcome will be decided by the preference order she feeds into g, then honesty is the unique policy which maximizes her expected utility.

Claim 1: *Let $u \in \mathcal{U}$ be any utility function, with ordinal preferences $p \in \mathcal{P}$. Then for any other preference order $p' \in \mathcal{P}$, we have $\mathbb{E}u[g(p)] - \mathbb{E}u[g(p')] \geq \frac{2}{A(A-1)}\gamma(u)$.*

Proof. Let $\mathcal{L}(p')$ be the set of all pairs $\{a, b\}$ for which p' disagrees with p. Let $\{a, b\}$ be the pair randomly chosen in Step 1 of the procedure defining g. If $\{a, b\} \notin \mathcal{L}(p')$, then p and p' yield the same outcome in Step 2. (Unless both p and p' are indifferent between a and b, in which case they might yield different outcomes, but with the same utility.) But if $\{a, b\} \in \mathcal{L}(p')$, then p and p' yield opposite outcomes in Step 2, and the outcome-utility of p' is at least $\gamma(u)$ less than the outcome-utility of p. There are $\frac{A(A-1)}{2}$ pairs, so the probability of any particular pair being chosen is $\frac{2}{A(A-1)}$. Thus, $\mathbb{E}u[g(p)] - \mathbb{E}u[g(p')] = \frac{2}{A(A-1)} \sum_{\{a,b\} \in \mathcal{L}(p')} |u(a) - u(b)| \geq \frac{2}{A(A-1)}\gamma(u)$. ◇ Claim 1

We now define a stochastic voting rule $\widetilde{G}$ through the following random procedure. For any $N \in \mathbb{N}$, and any ordinal preference profile $\mathbf{p} = (p_n)_{n=1}^N \in \mathcal{P}^N$:

1. Let $\widetilde{n} \in [1 \ldots N]$ be a uniformly distributed random voter.
2. Randomly select an alternative according to the distribution $g(p_{\widetilde{n}})$.

From Claim 1, it is easy to see that $\widetilde{G}$ is truth-revealing. Let $(q_N)_{N=1}^\infty$ be a sequence of real numbers in the interval $[0, 1]$. Consider the stochastic voting rule $\widetilde{F}$ defined as follows. For any $N \in \mathbb{N}$, and any ordinal profile $\mathbf{p} \in \mathcal{P}^N$:

- With probability $1 - q_N$, set $\widetilde{F}(\mathbf{p}) := F(\mathbf{p})$.
- With probability q_N, let $\widetilde{F}(\mathbf{p}) := \widetilde{G}(\mathbf{p})$.

Claim 2: *If $\lim_{N\to\infty} q_N = 0$, then $\widetilde{F}$ is asymptotically equal to F.*

Proof. Clearly, $P_N(F, \widetilde{F}) \geq 1 - q_N$. The claim follows. ◇ Claim 2

Claim 3: *If $\lim_{N\to\infty} \dfrac{N \cdot \tau(\mathcal{B}_N, F_N)}{q_N} = 0$, then $\widetilde{F}$ is asymptotically ordinally truth-revealing for $\boldsymbol{\mathcal{B}}$.*

Proof. Let $u \in \mathcal{U}$ be a utility function, and let $p \in \mathcal{P}$ be the corresponding preference order. Let $\beta \in \mathcal{B}_N$. For any preference order $p' \in \mathcal{P}$ and any voting rule F, let $\mathbb{E}u(F, p', \beta)$ be the expected utility of declaring the preference p' in the voting rule F, given beliefs β. Then for any $p' \in \mathcal{P}_N \setminus \{p\}$, Claim 1 says that

$$\begin{aligned} \mathbb{E}u[g(p)] - \mathbb{E}u[g(p')] &\geq \frac{2}{A(A-1)}\gamma(u), \\ \text{and thus,}\quad \mathbb{E}u(\widetilde{G}, p, \beta) - \mathbb{E}u(\widetilde{G}, p', \beta) &\underset{(*)}{=} \frac{1}{N}\Big(\mathbb{E}u[g(p)] - \mathbb{E}u[g(p')]\Big) \\ &= \frac{2}{N\,A\,(A-1)}\gamma(u), \end{aligned} \tag{A4}$$

where $(*)$ is because each voter has a $1/N$ probability of getting picked in Step 1 of the procedure defining $\widetilde{G}$. Meanwhile, given any profile $\mathbf{p}$ of the other $N - 1$ voters in the population, even if $(p, \mathbf{p})$

is a nearly-tied profile, we have $u[F(p', \mathbf{p})] - u[F(p, \mathbf{p})] \leq 1$, by the definition of $\mathcal{U}$. (This is any voter's maximum possible benefit from voting strategically.) Thus,

$$\begin{aligned} \mathbb{E}u(F, p', \beta) - \mathbb{E}u(F, p, \beta) &\leq \text{Prob}_\beta \Big(\text{the profile is nearly tied}\Big) \\ &= \tau(\beta, F) \quad \leq \quad \tau(\mathcal{B}, F). \end{aligned} \tag{A5}$$

Thus,

$$\begin{aligned} &\mathbb{E}u(\widetilde{F}, p, \beta) - \mathbb{E}u(\widetilde{F}, p', \beta) \\ &= \quad q_N \cdot \Big[\mathbb{E}u(\widetilde{G}, p, \beta) - \mathbb{E}u(\widetilde{G}, p', \beta)\Big] + (1 - q_N) \cdot [\mathbb{E}u(F, p, \beta) - \mathbb{E}u(F, p', \beta)] \\ &\underset{(*)}{\geq} \quad \frac{q_N}{N\,A\,(A-1)}\gamma(u) \; - \; (1 - q_N) \cdot \tau(\mathcal{B}_N, F_N), \end{aligned}$$

where $(*)$ is by inequalities (A4) and (A5). A rational voter will vote honestly if and only if this expression is nonnegative. But

$$\left(\frac{q_N}{N\,A\,(A-1)}\gamma(u) \; - \; (1 - q_N)\tau(\mathcal{B}_N, F_N) \; \geq \; 0\right) \quad \Longleftrightarrow \quad \Big(\gamma(u) \geq \epsilon_N\Big),$$

$$\text{where} \quad \epsilon_N \; := \; A\,(A-1)\,(1-q_N)\,\frac{N\,\tau(\mathcal{B}_N, F_N)}{q_N}.$$

Thus, if $\gamma(u) \geq \epsilon_N$, then the voter will vote honestly. Note that the hypothesis of Claim 3 implies that $\lim_{N\to\infty} \epsilon_N = 0$.

Now, let ρ be any probability distribution on $\mathcal{U}$. If $\widetilde{u}$ is a random utility function drawn from the distribution ρ, then

$$\lim_{\epsilon \searrow 0} \text{Prob}_\rho[\gamma(\widetilde{u}) > \epsilon] \; = \; 1, \tag{A6}$$

because $\gamma(\widetilde{u})$ is almost-surely nonzero, by definition. Thus,

$$\text{Tr}(\beta, \rho, \widetilde{F}_N) \quad \geq \quad \text{Prob}\,[\gamma(\widetilde{u}) \geq \epsilon_N] \quad \xrightarrow[N\to\infty]{(*)} \quad 1,$$

as desired. Here, $(*)$ is by equation (A6), because $\lim_{N\to\infty} \epsilon_N = 0$. ◇ `Claim 3`

Now, if $\mathcal{B}$ is regular, then $\lim_{N\to\infty} N \cdot \tau(\mathcal{B}_N, F_N) \; = \; 0$. Then it is always possible to find a sequence $\{q_N\}_{N=1}^\infty$ which simultaneously satisfies the condition of Claims 2 and 3. For example, define $q_N := \min\{1, \sqrt{N \cdot \tau(\mathcal{B}_N, F_N)}\}$. Then clearly $\lim_{N\to\infty} q_N = 0$, because $\lim_{N\to\infty} N \cdot \tau(\mathcal{B}_N, F_N) \; = \; 0$; thus, Claim 2 says that $\widetilde{F}$ is asymptotically equal to F. But also,

$$\lim_{N\to\infty} \frac{N \cdot \tau(\mathcal{B}_N, F_N)}{q_N} \quad = \quad \lim_{N\to\infty} \sqrt{N \cdot \tau(\mathcal{B}_N, F_N)} \; = \; 0.$$

Thus, Claim 3 says that $\widetilde{F}$ is asymptotically ordinally truth-revealing for $\mathcal{B}$. □

Proof of Theorem 5. For any $u \in \mathcal{U}$, let $\overline{u} := \frac{1}{|\mathcal{A}|}\sum_{a\in\mathcal{A}} u(a)$, and then define

$$\begin{aligned} \mathcal{S}_u &:= \{a \in \mathcal{A}\,;\; u(a) > \overline{u}\}, \\ \mathcal{I}_u &:= \{a \in \mathcal{A}\,;\; u(a) < \overline{u}\}, \\ \text{and} \quad \mathcal{C}_u &:= \{a \in \mathcal{A}\,;\; u(a) = \overline{u}\}. \end{aligned}$$

Mnemonically, $\mathcal{S}_u$ and $\mathcal{I}_u$ are the "superior" and "inferior" alternatives, according to u (in the sense that they are above average and below average, respectively), while $\mathcal{C}_u$ is the "cutoff" set (the exactly average alternatives). The set $\mathcal{C}_u$ might be empty, but the sets $\mathcal{S}_u$ and $\mathcal{I}_u$ must be nonempty, because u is non-constant (by definition of $\mathcal{U}$). Now define

$$\gamma(u) \quad := \quad \min_{s\in\mathcal{S}_u} u(s) \; - \; \overline{u} \quad \text{and} \quad \beta(u) \quad := \quad \overline{u} \; - \; \max_{i\in\mathcal{I}_u} u(i).$$

Both of these values are strictly positive. Thus, if we define $\alpha(u) := \min\{\beta(u), \gamma(u)\}$, then $\alpha(u) > 0$, for any $u \in \mathcal{U}$. Thus if $\widetilde{u}$ is a random utility function drawn from the distribution ρ, then

$$\lim_{\epsilon \searrow 0} \text{Prob}_\rho[\alpha(\widetilde{u}) > \epsilon] \;=\; 1. \tag{A7}$$

For any binary signal $v \in \mathcal{V}_0$, let $\widetilde{h}(v) \in \widetilde{\mathcal{A}}$ be the random alternative obtained from the following random procedure:

1. Pick an element $a \in \mathcal{A}$ uniformly at random.
2. If $v(a) = 1$, then select a.
3. Otherwise, let $c \in \mathcal{A}$ be random element drawn uniformly from $\mathcal{A}$. Select c.

This defines a function $\widetilde{h} : \mathcal{V}_0 \longrightarrow \widetilde{\mathcal{A}}$. Our first claim says that $\widetilde{h}$ is "truth-revealing" in the following sense: if a single voter is told that the outcome will be decided by the binary utility function she feeds into $\widetilde{h}$, then the unique policy which maximizes her expected utility is to be truthful in the sense of definition (5).

Claim 1: *Let $u \in \mathcal{U}$ be any utility function. Let $v, v' \in \mathcal{V}_0$ denote, respectively, a truthful and a non-truthful function. Then $\mathbb{E}u[\widetilde{h}(v)] - \mathbb{E}u[\widetilde{h}(v')] \geq \frac{\alpha(u)}{A}$.*

Proof. Define

$$\mathcal{S}_u^* \;=\; \{s \in \mathcal{S}_u \;;\; v'(s) = 0\} \qquad \text{and} \qquad \mathcal{I}_u^* \;=\; \{i \in \mathcal{I}_u \;;\; v'(i) = 1\}.$$

Since v' is non-truthful, it violates (5); thus, at least one of $\mathcal{S}_u^*$ or $\mathcal{I}_u^*$ is nonempty. Let $\widetilde{a}$ denote the random alternative chosen in Step 1 of the mechanism. Suppose $\widetilde{a} = s$ for some $s \in \mathcal{S}_u^*$. Then Step 2 would yield $\widetilde{h}(v) = s$. Meanwhile, Step 3 would yield $\widetilde{h}(v') = c$, where c is chosen uniformly at random. Thus, $\mathbb{E}\left[u[\widetilde{h}(v)] \;\middle|\; \widetilde{a} = s\right] = u(s)$, whereas $\mathbb{E}\left[u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} = s\right] = \overline{u}$. Thus, $\mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} = s\right] = u(s) - \overline{u} \geq \gamma(u)$. This holds for all $s \in \mathcal{S}_u^*$. Thus,

$$\mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} \in \mathcal{S}_u^*\right] \quad \geq \quad \gamma(u). \tag{A8}$$

Next, suppose $\widetilde{a} = i$ for some $i \in \mathcal{I}_u^*$. Then Step 3 would yield $\widetilde{h}(v) = c$, where c is chosen uniformly at random. Meanwhile, Step 2 would yield $\widetilde{h}(v') = i$. Thus, $\mathbb{E}\left[u[\widetilde{h}(v)] \;\middle|\; \widetilde{a} = i\right] = \overline{u}$, whereas $\mathbb{E}\left[u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} = i\right] = u(i)$. Thus, $\mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} = i\right] = \overline{u} - u(i) \geq \beta(u)$. This holds for all $i \in \mathcal{I}_u^*$. Thus,

$$\mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} \in \mathcal{I}_u^*\right] \quad \geq \quad \beta(u). \tag{A9}$$

We must now deal with the case when $\widetilde{a}$ is not in $\mathcal{S}_u^*$ or $\mathcal{I}_u^*$. There are three subcases. If $\widetilde{a} = s$ for some $s \in \mathcal{S}_u \setminus \mathcal{S}_u^*$, then Step 2 yields both $\mathbb{E}\left[u[\widetilde{h}(v)] \;\middle|\; \widetilde{a} = s\right] = u(s)$ and $\mathbb{E}\left[u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} = s\right] = u(s)$. If $\widetilde{a} = i$ for some $i \in \mathcal{I}_u \setminus \mathcal{I}_u^*$, then Step 3 yields both $\mathbb{E}\left[u[\widetilde{h}(v)] \;\middle|\; \widetilde{a} = i\right] = \overline{u}$ and $\mathbb{E}\left[u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} = i\right] = \overline{u}$. Finally, if $\widetilde{a} = c$ for some $c \in \mathcal{C}_u$, then we will have $\mathbb{E}\left[u[\widetilde{h}(v)] \;\middle|\; \widetilde{a} = c\right] = \overline{u}$ by applying either Step 2 (if $v(c) = 1$) or Step 3 (if $v(c) = 0$). Likewise, we have $\mathbb{E}\left[u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} = c\right] = \overline{u}$ by applying either Step 2 (if $v'(c) = 1$) or Step 3 (if $v'(c) = 0$). Let $\mathcal{A}_0 := \mathcal{C} \sqcup (\mathcal{S}_u \setminus \mathcal{S}_u^*) \sqcup (\mathcal{I}_u \setminus \mathcal{I}_u^*)$. Then combining the three subcases, we see that

$$\mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} \in \mathcal{A}_0\right] \quad = \quad 0. \tag{A10}$$

Now, clearly $\mathcal{A} = \mathcal{A}_0 \sqcup \mathcal{S}_u^* \sqcup \mathcal{I}_u^*$. Thus,

$$\begin{aligned}
\mathbb{E}[u[\widetilde{h}(v)] - u[\widetilde{h}(v')]] \quad &= \quad \mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} \in \mathcal{A}_0\right] \cdot \text{Prob}\left[\widetilde{a} \in \mathcal{A}_0\right] \\
&\quad + \; \mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} \in \mathcal{S}_u^*\right] \cdot \text{Prob}\left[\widetilde{a} \in \mathcal{S}_u^*\right] \\
&\quad + \; \mathbb{E}\left[u[\widetilde{h}(v)] - u[\widetilde{h}(v')] \;\middle|\; \widetilde{a} \in \mathcal{I}_u^*\right] \cdot \text{Prob}\left[\widetilde{a} \in \mathcal{I}_u^*\right] \\
&\underset{(*)}{\geq} \quad 0 \cdot \text{Prob}[\widetilde{a} \in \mathcal{A}_0] + \frac{\gamma(u) \cdot |\mathcal{S}_u^*|}{A} + \frac{\beta(u) \cdot |\mathcal{I}_u^*|}{A} \quad \underset{(\dagger)}{\geq} \quad \frac{\alpha(u)}{A},
\end{aligned}$$

as desired. Here, (∗) is by combining equation (A10) with inequalities (A8) and (A9), while (†) is because $\widetilde{a}$ is uniformly distributed, and either $\mathcal{S}_u^*$ or $\mathcal{I}_u^*$ is nonempty, because v' violates (5), ◇ Claim 1

We now define a stochastic voting rule $\widetilde{H}$ through the following random procedure. For any $N \in \mathbb{N}$, and any binary utility profile $\mathbf{v} = (v_n)_{n=1}^N \in \mathcal{U}_0^N$:

1. Let $\widetilde{n} \in [1 \ldots N]$ be a uniformly distributed random voter.
2. Select the random alternative $\widetilde{h}(v_{\widetilde{n}})$.

From Claim 1, it is easy to see that $\widetilde{H}$ is binarily truth-revealing. Let $(q_N)_{N=1}^\infty$ be a sequence of real numbers in the interval $[0,1]$. Let F be the approval voting rule. Consider the stochastic voting rule $\widetilde{F}$ defined as follows. For any $N \in \mathbb{N}$, and any ordinal profile $\mathbf{v} \in \mathcal{V}^N$:

- With probability $1 - q_N$, set $\widetilde{F}(\mathbf{v}) := F(\mathbf{v})$.
- With probability q_N, let $\widetilde{F}(\mathbf{v}) := \widetilde{H}(\mathbf{v})$.

The rest of the proof proceeds as in the proof of Theorem 4. First, we prove first the next two claims (the proofs are omitted for the sake for brevity):

Claim 2: *If* $\lim_{N\to\infty} q_N = 0$, *then* $\widetilde{F}$ *is asymptotically equal to* F.

Claim 3: *If* $\lim_{N\to\infty} \dfrac{N \cdot \tau(\mathcal{B}_N, F_N)}{q_N} = 0$, *then* $\widetilde{F}$ *is asymptotically binarily truth-revealing for* $\boldsymbol{\mathcal{B}}$.

We can now conclude the proof. Take any regular $\boldsymbol{\mathcal{B}}$. Then $\lim_{N\to\infty} N \cdot \tau(\mathcal{B}_N, F_N) = 0$. So it is always possible to find a sequence $\{q_N\}_{N=1}^\infty$ which simultaneously satisfies the condition of Claims 2 and 3, which proves the result. □

The proof of Theorem 6 depends on a preliminary result. For any $\epsilon > 0$, we say that a stochastic cardinal voting rule $\widetilde{F} = (\widetilde{F}_N)_{N=1}^\infty$ is *asymptotically ϵ-truth-revealing* for $\boldsymbol{\mathcal{B}}$ if

$$\lim_{N\to\infty} \mathrm{Tr}_\epsilon(\mathcal{B}_N, \rho, \widetilde{F}_N) \quad = \quad 1, \qquad \text{for all } \rho \in \Delta(\mathcal{U}). \tag{A11}$$

Thus, $\widetilde{F}$ is asymptotically cardinally truth-revealing for $\boldsymbol{\mathcal{B}}$ if it is asymptotically ϵ-truth-revealing for all $\epsilon > 0$. Meanwhile, let $F = (F_N)_{N=1}^\infty$ be a (deterministic) cardinal voting rule. We will say that $\widetilde{F}$ is *asymptotically ϵ-similar* to F if $\widetilde{F}$ is asymptotically equal to some cardinal voting rule $H = (H_N)_{N=1}^\infty$, and for all $N \in \mathbb{N}$, H_N is ϵ-similar to F_N.

Proposition A1 *Let F be any cardinal voting rule. Let $\boldsymbol{\mathcal{B}}$ be any robustly regular culture for F. Then for any sufficiently small $\epsilon > 0$, there is a stochastic cardinal voting rule $\widetilde{F}^\epsilon$ which is asymptotically ϵ-similar to F, and which is asymptotically ϵ-truth-revealing for $\boldsymbol{\mathcal{B}}$.*

Proof. Since $\boldsymbol{\mathcal{B}}$ is robustly regular for F, there is some $\epsilon_0 > 0$ such that the "ϵ_0-distorted" culture $\boldsymbol{\mathcal{B}}^{(\epsilon_0)}$ is regular for F. Thus, for any $\epsilon \in (0, \epsilon)$, the culture $\boldsymbol{\mathcal{B}}^{(\epsilon)}$ is also regular for F, because $\mathcal{B}_N^{(\epsilon)} \subseteq \mathcal{B}_N^{(\epsilon_0)}$ for all $N \in \mathbb{N}$.

Let $\Delta(\mathcal{A})$ be the set of probability distributions over $\mathcal{A}$. For all $a \in \mathcal{A}$, let $\delta_a \in \Delta(\mathcal{A})$ be the point mass at a. For any $a, b \in \mathcal{A}$ and $r \in [0,1]$, let $\delta_{a,b}^r$ be the lottery which gives probability r to a and probability $1 - r$ to b. For any $\epsilon > 0$, we define

$$\begin{aligned} \mathcal{R}_\epsilon &:= \{m\epsilon \,;\, m \in \mathbb{N} \text{ such that } 0 \le m\epsilon \le 1\} \\ \text{and } \mathcal{A}_\epsilon &:= \left\{\delta_{a,b}^r \,;\, a, b \in \mathcal{A} \text{ and } r \in \mathcal{R}_\epsilon\right\}. \end{aligned} \tag{A12}$$

In other words, $\mathcal{R}_\epsilon$ is an ϵ-spaced mesh of points in the interval $[0,1]$, and $\mathcal{A}_\epsilon$ is the set of all two-outcome lotteries in $\Delta(\mathcal{A})$ with probabilities drawn from $\mathcal{R}_\epsilon$. Each alternative $a \in \mathcal{A}$ can be identified with the point mass δ_a, which is an element of $\mathcal{A}_\epsilon$; thus, we can treat $\mathcal{A}$ as a subset of $\mathcal{A}_\epsilon$. Note that $\mathcal{A}_\epsilon$ is finite. Let $\mathcal{P}_\epsilon$ be the set of all preference orders over $\mathcal{A}_\epsilon$. Any $u \in \mathcal{U}$ determines a preference order $p_\epsilon^u \in \mathcal{P}_\epsilon$.

Claim 1: *Let $u, u' \in \mathcal{U}$. If $\|u - u'\|_\infty > \epsilon$, then $p_\epsilon^u \neq p_\epsilon^{u'}$.*

Proof. If u and u' induce different preference orders on $\mathcal{A}$, then we are done (because $\mathcal{A}$ is a subset of $\mathcal{A}_\epsilon$). So, assume without loss of generality that they induce the same preference order on $\mathcal{A}$. Let l and o be, respectively, the u-maximal and u-minimal elements of $\mathcal{A}$. Since u' induces the same preference order on $\mathcal{A}$, these elements are also u'-maximal and u'-minimal. Thus, $u(l) = 1 = u'(l)$ and $u(o) = 0 = u'(o)$. Thus, for any $r \in [0,1]$, we have $\mathbb{E}u[\delta^r_{l,o}] = \mathbb{E}u'[\delta^r_{l,o}] = r$.

Now, suppose $\|u - u'\|_\infty > \epsilon$. Then there is some $a \in \mathcal{A}$ such that $|u(a) - u'(a)| > \epsilon$. Without loss of generality, assume $u(a) > u'(a)$; then there is some $r \in \mathcal{R}_\epsilon$ such that $u(a) > r > u'(a)$. In other words, $\mathbb{E}u(\delta_a) > \mathbb{E}u[\delta^r_{l,o}]$ while $\mathbb{E}u'[\delta^r_{l,o}] > \mathbb{E}u'(\delta_a)$. Thus, $\delta_a \succ_{p^u_\epsilon} \delta^r_{l,o}$, whereas $\delta_a \prec_{p^{u'}_\epsilon} \delta^r_{l,o}$. Thus, $p^u_\epsilon \neq p^{u'}_\epsilon$. ◇ Claim 1

For each preference order $q \in \mathcal{P}_\epsilon$, let us fix some $u_q \in \mathcal{U}$ such that $p^{u_q}_\epsilon = q$. Finally, for all $u \in \mathcal{U}$, let $u^\epsilon := u^q$ where $q := p^u_\epsilon$.

Claim 2: $\|u^\epsilon - u\|_\infty \leq \epsilon$.

Proof. (by contradiction) Suppose $\|u^\epsilon - u\|_\infty > \epsilon$. Then Claim 1 says that $p_{u^\epsilon} \neq p_u$. This contradicts the definition of u^ϵ. ◇ Claim 2

For any $N \in \mathbb{N}$, define the function $G^\epsilon_N : \mathcal{P}^N_\epsilon \longrightarrow \mathcal{A} \subset \mathcal{A}_\epsilon$ by

$$G^\epsilon_N(p_1, \ldots, p_N) \quad := \quad F_N(u_{p_1}, \ldots, u_{p_N}) \qquad \text{for all } (p_1, \ldots, p_N) \in \mathcal{P}^N_\epsilon. \tag{A13}$$

Now let $G^\epsilon = (G^\epsilon_N)^\infty_{N=1}$; this is an ordinal voting rule on $\mathcal{A}_\epsilon$, which we can regard as an "ϵ-approximation" of F. (Note: although G^ϵ is formally an ordinal voting rule on $\mathcal{A}_\epsilon$, its output is always an element of $\mathcal{A}$.)

For any probability distribution $\beta \in \Delta(\mathcal{U}^N)$, we can define a probability distribution $\beta^\epsilon \in \Delta(\mathcal{P}^N_\epsilon)$ as follows: for any preference profile $\mathbf{q} := (q_1, \ldots, q_N) \in \mathcal{P}^N_\epsilon$, we set $\beta^\epsilon(\mathbf{q}) := \beta\{\mathbf{u} \in \mathcal{U}^N;\ p^{u_n}_\epsilon = q_n$, for all $n \in [1 \ldots N]\}$. Thus, if $\widetilde{\mathbf{u}} = (\widetilde{u}_1, \ldots, \widetilde{u}_N)$ is a β-random utility profile, then $(p^{\widetilde{u}_1}_\epsilon, \ldots, p^{\widetilde{u}_N}_\epsilon)$ is a random preference profile with distribution β^ϵ. If $\mathcal{B}_N \subseteq \Delta(\mathcal{U}^N)$ is a set of such distributions, then we define $\mathcal{B}^\epsilon_N := \{\beta^\epsilon;\ \beta \in \mathcal{B}_N\}$, which is a subset of $\Delta(\mathcal{P}^N_\epsilon)$. Finally, given the culture $\boldsymbol{\mathcal{B}} = (\mathcal{B}_N)^\infty_{N=1}$ (concerning $\mathcal{U}$-profiles), we define the culture $\boldsymbol{\mathcal{B}}^\epsilon := (\mathcal{B}^\epsilon_N)^\infty_{N=1}$ (concerning $\mathcal{P}_\epsilon$-profiles).

Claim 3: *$\boldsymbol{\mathcal{B}}^\epsilon$ is a regular culture for G^ϵ.*

Proof. Let $N \in \mathbb{N}$, let $\beta \in \mathcal{B}_N$, and let $\widetilde{\mathbf{u}} = (\widetilde{u}_1, \ldots, \widetilde{u}_N)$ be a β-random utility profile. Let $\widetilde{\mathbf{u}}^\epsilon := (\widetilde{u}^\epsilon_1, \ldots, \widetilde{u}^\epsilon_N)$; then Claim 2 implies that $\widetilde{\mathbf{u}}^\epsilon$ is an ϵ-distortion of $\widetilde{\mathbf{u}}$. Thus, if β' denotes the probability distribution of $\widetilde{\mathbf{u}}^\epsilon$, then β' is ϵ-close to β; thus, $\beta' \in \mathcal{B}^{(\epsilon)}_N$, by defining formula (6). Now, for all $N \in \mathbb{N}$, define $\mathcal{B}'_N := \{\beta';\ \beta \in \mathcal{B}_N\}$. Then define the culture $\boldsymbol{\mathcal{B}}' := (\mathcal{B}'_N)^\infty_{N=1}$. By hypothesis, the culture $\boldsymbol{\mathcal{B}}^{(\epsilon)}$ is regular for F. Thus, the culture $\boldsymbol{\mathcal{B}}'$ is also regular for F, because we have $\mathcal{B}'_N \subseteq \mathcal{B}^{(\epsilon)}_N$ for all $N \in \mathbb{N}$.

Now, let $\mathbf{p} = (p_1, \ldots, p_N) \in \mathcal{P}^N_\epsilon$ be a preference profile. From the defining formula (A13) for the rule G^ϵ, it is clear that

$$\Big((p_1, \ldots, p_N) \text{ is almost-tied for } G^\epsilon_N\Big) \iff \Big((u_{p_1}, \ldots, u_{p_N}) \text{ is almost-tied for } F_N\Big). \tag{A14}$$

Suppose $\beta \in \mathcal{B}_N$, so that $\beta^\epsilon \in \mathcal{B}^\epsilon_N$ and $\beta' \in \mathcal{B}'_N$. If $\widetilde{\mathbf{p}} = (\widetilde{p}_1, \ldots, \widetilde{p}_N)$ is a β^ϵ-random preference profile, then $(u_{\widetilde{p}_1}, \ldots, u_{\widetilde{p}_N})$ is a β'-random utility profile. Thus, from statement (A14), we deduce that

$$\mathrm{Prob}_{\beta^\epsilon}\Big((\widetilde{p}_1, \ldots, \widetilde{p}_N) \text{ is almost-tied for } G^\epsilon_N\Big) = \mathrm{Prob}_{\beta'}\Big((u_{\widetilde{p}_1}, \ldots, u_{\widetilde{p}_N}) \text{ is almost-tied for } F_N\Big).$$

In other words, $\tau(\beta^\epsilon, G^\epsilon_N) = \tau(\beta', F_N)$. Taking the supremum over all $\beta \in \mathcal{B}_N$ (hence, over all $\beta^\epsilon \in \mathcal{B}^\epsilon_N$ and $\beta' \in \mathcal{B}'_N$), we deduce that $\tau(\mathcal{B}^\epsilon_N, G^\epsilon_N) = \tau(\mathcal{B}'_N, F_N)$. However, as we have already argued, the culture $\boldsymbol{\mathcal{B}}'$ is regular for F. Thus, we conclude that $\lim_{N\to\infty} N \cdot \tau(\mathcal{B}^\epsilon_N, G^\epsilon_N) = \lim_{N\to\infty} N \cdot \tau(\mathcal{B}'_N, F_N) = 0$, as desired. ◇ Claim 3

Theorem 4 and Claim 3 jointly imply that there is a stochastic ordinal voting rule $\widetilde{G}^\epsilon = (\widetilde{G}^\epsilon_N)^\infty_{N=1}$ on $\mathcal{A}_\epsilon$ which is asymptotically equal to G^ϵ and asymptotically ordinally truth-revealing for the culture $\boldsymbol{\mathcal{B}}^\epsilon$. For any $N \in \mathbb{N}$, define the function $\widetilde{F}^\epsilon_N : \mathcal{U}^N \longrightarrow \widetilde{\mathcal{A}}$ as follows: for any utility profile $\mathbf{u} = (u_1, \ldots, u_N) \in \mathcal{U}^N$, let $\mathbf{p} := (p^{u_1}_\epsilon, \ldots, p^{u_N}_\epsilon) \in \mathcal{P}^N_\epsilon$ be the corresponding profile of ordinal preferences over $\mathcal{A}_\epsilon$, and let $\widetilde{\alpha} :=$

$\widetilde{G}^\epsilon_N(\mathbf{p})$; thus, $\widetilde{\alpha}$ is an random element of $\mathcal{A}_\epsilon$ —i.e. a random probability distribution over $\mathcal{A}$. Finally, let $\widetilde{F}^\epsilon_N(\mathbf{u})$ be an $\widetilde{\alpha}$-random alternative. This yields a stochastic cardinal voting rule $\widetilde{F}^\epsilon = (\widetilde{F}^\epsilon_N)_{N=1}^\infty$ ranging over $\mathcal{A}$. We will now show that $\widetilde{F}^\epsilon$ is both asymptotically ϵ-truth-revealing and asymptotically ϵ-similar to F.

First, a bit of notation. For any $\delta > 0$, and any random variables $\tilde{x}$ and $\tilde{y}$, we will write "$\tilde{x} \overset{\delta}{\approx} \tilde{y}$" to mean "$\text{Prob}[\tilde{x} = \tilde{y}] > 1 - \delta$".

Claim 4: *$\widetilde{F}^\epsilon$ is asymptotically ϵ-similar to F.*

Proof. For all $N \in \mathbb{N}$, and any utility profile $(u_1, \ldots, u_N) \in \mathcal{U}^N$, define $H_N(u_1, \ldots, u_N) := F_N(u_1^\epsilon, \ldots, u_N^\epsilon)$; this yields a function $H_N : \mathcal{U}^N \longrightarrow \mathcal{A}$, and whence, a voting rule $H = (H_N)_{N=1}^\infty$ on $\mathcal{A}$. Observe that H is ϵ-similar to F by construction.

We claim that $\widetilde{F}^\epsilon$ is asymptotically equal to H. To see this, recall that $\widetilde{G}^\epsilon$ is asymptotically equal to G^ϵ. Thus, for any $\delta > 0$, there is some $N_\delta \in \mathbb{N}$ such that, for all $N \geq N_\delta$, we have $P_N(\widetilde{G}^\epsilon, G^\epsilon) > 1 - \delta$. In other words, for any $(p_1, \ldots, p_N) \in \mathcal{P}_\epsilon^N$, we have $\widetilde{G}^\epsilon_N(p_1, \ldots, p_N) \overset{\delta}{\approx} G^\epsilon_N(p_1, \ldots, p_N)$.

Now let $\mathbf{u} = (u_1, \ldots, u_N) \in \mathcal{U}^N$; then $\widetilde{F}^\epsilon_N(\mathbf{u})$ is an $\widetilde{\alpha}$-random alternative, where $\widetilde{\alpha} := \widetilde{G}^\epsilon_N(p_\epsilon^{u_1}, \ldots, p_\epsilon^{u_N})$. But if $N \geq N_\delta$, then we have

$$\begin{aligned} \widetilde{\alpha} \quad = \quad \widetilde{G}^\epsilon_N(p_\epsilon^{u_1}, \ldots, p_\epsilon^{u_N}) \quad &\overset{\delta}{\approx} \quad G^\epsilon_N(p_\epsilon^{u_1}, \ldots, p_\epsilon^{u_N}) \underset{\text{(a)}}{=} F_N(u_{p_\epsilon^{u_1}}, \ldots, u_{p_\epsilon^{u_N}}) \\ &\underset{\text{(b)}}{=} F_N(u_1^\epsilon, \ldots, u_N^\epsilon) \underset{\text{(c)}}{=} H_N(u_1, \ldots, u_N), \end{aligned}$$

where (a) is by the defining formula (A13) for the rule G^ϵ_N, (b) is by the definition of $u_1^\epsilon, \ldots, u_N^\epsilon$, and (c) is by the definition of H_N. In other words, with probability greater than $1 - \delta$, $\widetilde{\alpha}$ is the point mass at $H_N(\mathbf{u})$, in which case an $\widetilde{\alpha}$-random variable must simply take the value $H_N(\mathbf{u})$. Thus, since $\widetilde{F}^\epsilon_N(\mathbf{u})$ is an $\widetilde{\alpha}$-random variable, we have $\widetilde{F}^\epsilon_N(\mathbf{u}) \overset{\delta}{\approx} H_N(\mathbf{u})$.

This holds for all $\mathbf{u} \in \mathcal{U}^N$; thus $P_N(\widetilde{F}^\epsilon, H) > 1 - \delta$. This holds for all $N \geq N_\delta$, and we can construct such an N_δ for any $\delta > 0$. Thus, $\widetilde{F}^\epsilon$ is asymptotically equal to H, as desired. ◇ Claim 4

Claim 5: *$\widetilde{F}^\epsilon$ is asymptotically ϵ-truth-revealing for the culture $\boldsymbol{\mathcal{B}}$.*

Proof. Fix $\rho \in \Delta(\mathcal{U})$. We must verify the limit (A11). Let $N \in \mathbb{N}$, let $\beta \in \mathcal{B}_N$, and consider a voter with utility function $u \in \mathcal{U}$ and beliefs β. She participates in $\widetilde{F}^\epsilon_N$ by reporting a utility function —say, u' —in $\mathcal{U}$. Let $p_\epsilon^{u'}$ be the preference order that u' induces on $\mathcal{A}_\epsilon$. By definition of $\widetilde{F}^\epsilon_N$, the β-expected value of u generated by reporting the utility function u' in the rule $\widetilde{F}^\epsilon_N$ is equal to the β^ϵ-expected value of u generated by reporting the preference order $p_\epsilon^{u'}$ in the rule $\widetilde{G}^\epsilon_N$. Thus, if reporting u' is this voter's best response for the rule $\widetilde{F}^\epsilon_N$ (i.e. it maximizes the β-expected value of u), then reporting $p_\epsilon^{u'}$ is her best response for the rule $\widetilde{G}^\epsilon_N$ (i.e. it maximizes the β^ϵ-expected value of u) .

Now suppose that u is a ρ-random variable, and recall that $\text{Tr}_\epsilon(\beta, \rho, \widetilde{F}^\epsilon_N)$ is the ρ-probability that $\|u - u'\|_\infty < \epsilon$, while $\text{Tr}(\beta^\epsilon, \rho, \widetilde{G}^\epsilon_N)$ is the ρ-probability that $p_\epsilon^{u'} = p_\epsilon^u$. If $p_\epsilon^{u'} = p_\epsilon^u$, then the contrapositive of Claim 1 yields $\|u - u'\|_\infty \leq \epsilon$. Thus,

$$\text{Tr}_\epsilon(\beta, \rho, \widetilde{F}^\epsilon_N) \quad \geq \quad \text{Tr}(\beta^\epsilon, \rho, \widetilde{G}^\epsilon_N). \tag{A15}$$

This argument applies to any $\beta \in \mathcal{B}_N$. Taking the infimum over all $\beta \in \mathcal{B}_N$, we get:

$$\begin{aligned} 1 \quad \geq \quad \text{Tr}_\epsilon(\mathcal{B}_N, \rho, \widetilde{F}^\epsilon_N) \underset{\text{(a)}}{=} \inf_{\beta \in \mathcal{B}_N} \text{Tr}_\epsilon(\beta, \rho, \widetilde{F}^\epsilon_N) \underset{\text{(b)}}{\geq} \inf_{\beta \in \mathcal{B}_N} \text{Tr}(\beta^\epsilon, \rho, \widetilde{G}^\epsilon_N) \\ \underset{\text{(c)}}{=} \inf_{\beta' \in \mathcal{B}_N^\epsilon} \text{Tr}(\beta', \rho, \widetilde{G}^\epsilon_N) \underset{\text{(d)}}{=} \text{Tr}(\mathcal{B}_N^\epsilon, \rho, \widetilde{G}^\epsilon_N) \xrightarrow[N\to\infty]{\text{(e)}} 1, \end{aligned}$$

which implies limit (A11). Here, (a) is the definition of $\text{Tr}_\epsilon()$ (from Section 7), (b) is by inequality (A15), (c) is by the definition of $\mathcal{B}_N^\epsilon$, and (d) is the definition of $\text{Tr}()$ (from Section 5). Finally, (e) is because $\widetilde{G}^\epsilon$ is asymptotically ordinally truth-revealing for the culture $\boldsymbol{\mathcal{B}}^\epsilon$, by hypothesis.

This argument holds for any $\rho \in \Delta(\mathcal{U})$; thus, $\widetilde{F}^\epsilon$ is asymptotically ϵ-truth-revealing for the culture $\boldsymbol{\mathcal{B}}$. ◇ Claim 5

Claims 4 and 5 together prove the theorem. □

Proof of Theorem 6. Let $(\epsilon_m)_{m=1}^{\infty}$ be a decreasing sequence converging to zero. For every $m \in \mathbb{N}$, Proposition A1 says there is a stochastic cardinal voting rule $\widetilde{F}^{\epsilon_m} = (\widetilde{F}_N^{\epsilon_m})_{N=1}^{\infty}$ which is asymptotically ϵ_m-similar to F and asymptotically ϵ_m-truth-revealing for $\boldsymbol{\mathcal{B}}$. Thus, there is a cardinal voting rule $H^m = (H_N^m)_{N=1}^{\infty}$, such that for all $N \in \mathbb{N}$, H_N is ϵ_m-similar to F_N, and such that

$$\lim_{N\to\infty} P_N(H^m, \widetilde{F}^{\epsilon_m}) \quad = \quad 1 \text{ and } \lim_{N\to\infty} \mathrm{Tr}_{\epsilon_m}(\mathcal{B}_N, \rho', \widetilde{F}_N^{\epsilon_m}) \quad = \quad 1, \quad \text{for all } \rho' \in \Delta(\mathcal{U}).$$

Thus, for any fixed $m \in \mathbb{N}$ and $\rho \in \Delta(\mathcal{U})$, there exists some $N_m \in \mathbb{N}$ such that

$$P_N(H^m, \widetilde{F}^{\epsilon_m}) \quad \geq \quad 1 - \epsilon_m, \qquad \text{for all } N \geq N_m, \tag{A16}$$

$$\text{and} \quad \mathrm{Tr}_{\epsilon_m}(\mathcal{B}_N, \rho, \widetilde{F}_N^{\epsilon_m}) \quad \geq \quad 1 - \epsilon_m, \qquad \text{for all } N \geq N_m. \tag{A17}$$

Without loss of generality, we can assume that $N_1 \leq N_2 \leq N_3 \leq \cdots$. (Otherwise, just replace N_m by $\max\{N_1, N_2, \ldots, N_m\}$ for each $m \in \mathbb{N}$.) Now define the stochastic cardinal rule $\widetilde{F}$ as follows: for any $N \in \mathbb{N}$, $\widetilde{F}_N := \widetilde{F}_N^{\epsilon_{m(N)}}$, where $m(N)$ is the largest $m \in \mathbb{N}$ such that $N_m \leq N$. In particular, for any $k \in \mathbb{N}$, if $N \geq N_k$, then $m(N) \geq k$. Meanwhile, define the (deterministic) cardinal rule G as follows: for any $N \in \mathbb{N}$, $G_N := H_N^{m(N)}$, where $(H_N^m)_{N=1}^{\infty}$ and $m(N)$ are defined as above.

Claim 1: *$\widetilde{F}$ is asymptotically equal to G.*

Proof. Let $\epsilon > 0$. Find $k \in \mathbb{N}$ such that $\epsilon_k \leq \epsilon$. Then for any $N \geq N_k$, we have

$$P_N(G, \widetilde{F}) \underset{\text{(a)}}{=} P_N(H^{m(N)}, \widetilde{F}^{\epsilon_{m(N)}}) \underset{\text{(b)}}{\geq} 1 - \epsilon_{m(N)} \underset{\text{(c)}}{\geq} 1 - \epsilon_k \geq 1 - \epsilon,$$

as desired. Here (a) is by the definitions of $\widetilde{F}$ and G, (b) is by inequality (A16), because $N \geq N_{m(N)}$ by the definition of $m(N)$. Finally, (c) is because $N \geq N_k$, so $m(N) \geq k$, so that $\epsilon_{m(N)} \leq \epsilon_k$. ◇ Claim 1

Claim 2: *$\widetilde{F}$ is asymptotically similar to F.*

Proof. Because of Claim 1, it suffices to show that the rule G is asymptotically similar to F. So, let $\epsilon > 0$. We must find some N_ϵ such that G_N is ϵ-similar to F_N for all $N \geq N_\epsilon$. Since $\lim_{m\to\infty} \epsilon_m = 0$, there exists some m_0 such that $0 < \epsilon_m < \epsilon$ for all $m > m_0$. Now let $N_\epsilon := N_{m_0}$. Then for any $N \geq N_\epsilon$, we have $m(N) \geq m_0$. Thus, $G_N = H_N^m$ for some $m \geq m_0$. But by hypothesis, H_N^m is ϵ_m-similar to F_N, and $\epsilon_m < \epsilon$. Thus, G_N is ϵ-similar to F_N, as desired. ◇ Claim 2

Claim 3: *$\widetilde{F}$ is asymptotically cardinally truth-revealing for $\boldsymbol{\mathcal{B}}$ and ρ.*

Proof. Let $\epsilon, \eta > 0$. Find k such that $\epsilon_k \leq \min\{\epsilon, \eta\}$. Then for any $N \geq N_k$, we have

$$\begin{aligned}\mathrm{Tr}_\epsilon(\mathcal{B}_N, \rho, \widetilde{F}_N) &\underset{\text{(a)}}{=} \mathrm{Tr}_\epsilon(\mathcal{B}_N, \rho, \widetilde{F}_N^{\epsilon_{m(N)}}) \underset{\text{(b)}}{\geq} \mathrm{Tr}_{\epsilon_{m(N)}}(\mathcal{B}_N, \rho, \widetilde{F}_N^{\epsilon_{m(N)}}) \\ &\underset{\text{(c)}}{\geq} 1 - \epsilon_{m(N)} \underset{\text{(d)}}{\geq} 1 - \eta.\end{aligned} \tag{A18}$$

Here (a) is the definition of $\widetilde{F}$, while (b) is because $N \geq N_k$, so $m(N) \geq k$, so that $\epsilon_{m(N)} \leq \epsilon_k \leq \epsilon$, and thus, ϵ-truth revelation is at least as probable as $\epsilon_{m(N)}$-truth revelation. Next, (c) by inequality (A17), because $N \geq N_{m(N)}$ by the definition of $m(N)$. Finally, (d) is because $m(N) \geq k$, so that $\epsilon_{m(N)} \leq \epsilon_k \leq \eta$.

Now fix $\epsilon > 0$. For any $\eta > 0$, the argument above shows that there is some N_η such that inequality (A18) holds for all $N > N_\eta$. Thus, $\lim_{N\to\infty} \mathrm{Tr}_\epsilon(\mathcal{B}_N, \rho, \widetilde{F}_N) = 1$. Repeat this argument for all $\epsilon > 0$ to conclude that $\widetilde{F}$ is asymptotically cardinally truth-revealing. ◇ Claim 3

Claims 2 and 3 together prove the theorem. □

Proof of Theorem 7. Let F be an ordinal voting rule, and let $\widetilde{F}$ be the stochastic voting rule defined in the proof of Theorem 4. Let $\mathbf{S} = (\mathbf{s}_N)_{N=1}^{\infty}$ be an F-regular society. For any $N \in \mathbb{N}$, we define the N-player Bayesian game $G_N := (\widetilde{F}_N, \mathbf{s}_N)$, and we define $\mathcal{B}_N$ as in equation (10). For any voter type t, let v_t^* be the preference order induced by the utility function u_t.

Claim 1: *There exists an $N_0 \in \mathbb{N}$ such that, for any $N \geq N_0$, any voter $m \in [1 \ldots N]$, any type $t_m \in \mathcal{T}_m^N$, and any strategy profile $\mathbf{V}$, if p_{t_m} is the probabilistic belief which voter-type t obtains from $\mathbf{V}$ via formula (7), and $p_{t_m} \in \mathcal{B}_N$, then her best response to p_{t_m} in the game G_N is to set $V_m(t_m) = v_{t_m}^*$.*

Proof. As argued in the proof of Theorem 4, a voter with utility function u will vote honestly in the rule $\widetilde{F}_N$ if only if $\gamma(u) \geq \epsilon_N$, where $\epsilon_N := A\,(A-1)\,(1-q_N)\,\frac{N\,\tau(\mathcal{B}_N, F_N)}{q_N}$. Now, $\gamma(u) > \epsilon$ for all $u \in \mathcal{U}$, because $\mathbf{S}$ is F-regular. Meanwhile, $\lim_{N\to\infty} \epsilon_N = 0$, because the culture $\boldsymbol{\mathcal{B}}_{\mathbf{S}}$ is regular for F. Thus, there is some N_0 such that, if $N > N_0$, then $\epsilon_N < \epsilon$, and thus, $\epsilon_N < \gamma(u)$ for all $u \in \mathcal{U}_{\mathbf{S}}$. In particular, for any $m \in [1 \ldots N]$ and $t_m \in \mathcal{T}_m^N$, we have $\epsilon_N < \gamma(u_{t_m})$. Thus, if $p_{t_m} \in \mathcal{B}_N$, then voter-type t_m's best response in the game G_N is to set $V_m(t_m) = v_{t_m}^*$. ◇ `Claim 1`

Suppose that $N \geq N_0$. For all $m \in [1 \ldots N]$ let $V_m : \mathcal{T}_m^N \to \mathcal{V}$ be the voting strategy such that $V_m(t_m) = v_{t_m}^*$ for all $t_m \in \mathcal{T}_m^N$ (the best response defined by Claim 1). Let $\mathbf{V}^N := (V_n)_{n=1}^N$ denote the resulting strategy profile with N players. For all $m \in [1 \ldots N]$ and all $t_m \in \mathcal{T}_m^N$, define $p_{t_m} \in \Delta(\mathcal{V}^N)$ by applying formula (7) to $\mathbf{V}^N$. Then clearly, $p_{t_m} = \beta_{t_m}$, where β_{t_m} is as defined in formula (9). Thus, $p_{t_m} \in \mathcal{B}_N$, by defining formula (10). Thus, for all $m \in [1 \ldots N]$ and $t \in \mathcal{T}_m^N$, Claim 1 says that $V_m(t_m)$ is voter-type t_m's best response, given her beliefs induced by $\mathbf{V}^N$. Therefore, $\mathbf{V}^N$ defines a Bayesian Nash equilibrium for G_N. Meanwhile, for all $N \in [1 \ldots N_0)$, let $\mathbf{V}^N$ be an arbitrary strategy profile. Let $\mathbb{V} := (\mathbf{V}^N)_{N=1}^{\infty}$; then $\mathbb{V}$ is an eventual Bayesian Nash equilibrium for the sequence $(G_N)_{N=1}^{\infty}$.

For any $N \geq N_0$, and any type profile $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$, if $\mathbf{v}_{\mathbf{t}}^*$ is the ordinal preference profile defined by $\mathbf{t}$, then $P_{\mathbf{t}}(\widetilde{F}_N, F_N, \mathbf{s}, \mathbf{V}^N) = \text{Prob}[\widetilde{F}^N(\mathbf{v}_{\mathbf{t}}^*) = F^N(\mathbf{v}_{\mathbf{t}}^*)]$, by the definition of V^N. Thus, formula (8) implies that $P(\widetilde{F}_N, F_N, \mathbf{s}, \mathbf{V}^N) \geq P_N(F, \widetilde{F})$, where $P_N(F, \widetilde{F}) := \inf_{\mathbf{v}\in\mathcal{V}^N} \text{Prob}\left[\widetilde{F}_N(\mathbf{v}) = F_N(\mathbf{v})\right]$. But $\lim_{N\to\infty} P_N(F, \widetilde{F}) = 1$ because $\widetilde{F}$ is asymptotically equal to F, by the construction from Theorem 4. Thus, $\widetilde{F}$ asymptotically implements F, as desired. □

Proof of Theorem 8. The strategy is the same as the proof of Theorem 7, but the details are more complicated, so we will break the proof into three steps.

Step 1. (Definition of the stochastic voting rule $\widetilde{G}^{\epsilon}$) Let $\mathbf{S} = (\mathbf{s}_N)_{N=1}^{\infty}$ be an F-regular society, and define $\boldsymbol{\mathcal{B}}_{\mathbf{S}} = (\mathcal{B}_N)_{N=1}^{\infty}$ as in equation (12). By hypothesis (R2), the culture $\boldsymbol{\mathcal{B}}_{\mathbf{S}}$ is robustly regular for F; thus, there is some $\epsilon_1 > 0$ such that the "ϵ_1-distorted" culture $\boldsymbol{\mathcal{B}}_{\mathbf{S}}^{(\epsilon_1)}$ is regular for F. Meawhile, let $\epsilon_2 > 0$ be the constant satisfying condition (R3). Let $\epsilon := \min\{\epsilon_1, \epsilon_2\}$; then ϵ satisfies both (R2) and(R3).

Define $\mathcal{R}_{\epsilon}$ and $\mathcal{A}_{\epsilon}$ as in formula (A12) in the proof of Proposition A1, and let $\mathcal{P}_{\epsilon}$ be the set of all ordinal preferences over $\mathcal{A}_{\epsilon}$. Thus, $\mathcal{P}_{\epsilon}$ is finite because $\mathcal{A}_{\epsilon}$ is finite. Any utility function $u \in \mathcal{U}$ determines a preference order p_{ϵ}^u in $\mathcal{P}_{\epsilon}$. As in the proof of Proposition A1, we can obtain a stochastic ordinal voting rule $\widetilde{G}^{\epsilon} = (\widetilde{G}_N^{\epsilon})_{N=1}^{\infty}$ which is asymptotically ordinally truth-revealing of the voters' ordinal preferences on $\mathcal{A}_{\epsilon}$ in the culture $\boldsymbol{\mathcal{B}}_{\mathbf{S}}^{(\epsilon)}$. For any $N \in \mathbb{N}$, we define the N-player Bayesian game $G_N := (\widetilde{G}_N^{\epsilon}, \mathbf{s}_N)$.

Step 2. (Construction of an eventual Bayesian Nash equilibrium) For any $p \in \mathcal{P}_{\epsilon}$, fix some $u_p \in \mathcal{U}$ which represents p. Let $\mathcal{U}^{\epsilon} := \{u_p;\ p \in \mathcal{P}_{\epsilon}\}$; this is a finite subset of $\mathcal{U}$. There is an obvious bijection between $\mathcal{P}_{\epsilon}$ and $\mathcal{U}^{\epsilon}$. Thus, without loss of generality, we can assume that the voting rule $\widetilde{G}^{\epsilon}$ uses the signal-set $\mathcal{V} = \mathcal{U}^{\epsilon}$. For any $u \in \mathcal{U}$, let u^{ϵ} be the (unique) element of $\mathcal{U}^{\epsilon}$ which represents the preference order p_{ϵ}^u. Claim 2 in the proof of Proposition A1 implies that

$$\|u^{\epsilon} - u\|_{\infty} \quad < \quad \epsilon, \qquad \text{for all } u \in \mathcal{U}. \tag{A19}$$

Claim 1: *There exists an $N_0 \in \mathbb{N}$ such that, for any $N \geq N_0$, any voter $m \in [1 \ldots N]$, any type $t \in \mathcal{T}_m^N$, and any strategy profile $\mathbf{V}$, if p_t is the probabilistic belief which voter-type t obtains from $\mathbf{V}$ via formula (7), and $p_t \in \mathcal{B}_N^{(\epsilon)}$, then her best response to p_t in the voting rule $\widetilde{G}_N^{\epsilon}$ is to set $v_t = u_t^{\epsilon}$.*

Proof. For any $u \in \mathcal{U}_{\mathbf{S}}$, let ρ_u be the probability measure on $\mathcal{U}$ which assigns probability 1 to u. Then $\lim_{N\to\infty} \mathrm{Tr}(\mathcal{B}_N^{(\epsilon)}, \rho_u, \widetilde{G}_N^\epsilon) = 1$, because $\widetilde{G}^\epsilon$ is asymptotically ordinally truth-revealing with respect to $\boldsymbol{\mathcal{B}}_{\mathbf{S}}^{(\epsilon)}$. Thus, there is some $N_u \in \mathbb{N}$ such that, for any $N \geq N_u$, we have $\mathrm{Tr}(\mathcal{B}_N^{(\epsilon)}, \rho_u, \widetilde{G}_N^{\epsilon_N}) > 0$. However, for any $\beta \in \mathcal{B}_N^{(\epsilon)}$, either $\mathrm{Tr}(\beta, \rho_u, \widetilde{G}_N^\epsilon) = 0$ or $\mathrm{Tr}(\beta, \rho_u, \widetilde{G}_N^\epsilon) = 1$ (because either a voter with utility function u and beliefs β finds it optimal to reveal her true preferences over $\mathcal{A}_\epsilon$, or she does not). Taking the infimum over all $\beta \in \mathcal{B}_N^{(\epsilon)}$, we deduce that either $\mathrm{Tr}(\mathcal{B}_N^{(\epsilon)}, \rho_u, \widetilde{G}_N^\epsilon) = 0$ or $\mathrm{Tr}(\mathcal{B}_N^{(\epsilon)}, \rho_u, \widetilde{G}_N^\epsilon) = 1$. Thus, if $N \geq N_u$, then $\mathrm{Tr}(\mathcal{B}_N^{(\epsilon)}, \rho_u, \widetilde{G}_N^\epsilon) = 1$. Now define $N_0 := \max_{u\in\mathcal{U}_{\mathbf{S}}} N_u$. Then N_0 is finite, because $\mathcal{U}_{\mathbf{S}}$ is finite by hypothesis (R1), because $\mathbf{S}$ is regular. Thus, for any $N \geq N_0$ and any $u \in \mathcal{U}_{\mathbf{S}}$, we have $\mathrm{Tr}(\mathcal{B}_N^{(\epsilon)}, \rho_u, \widetilde{G}_N^\epsilon) = 1$. In particular, for any $m \in [1 \ldots N]$, and any $t \in \mathcal{T}_m^N$, this implies that $\mathrm{Tr}(\beta, \rho_{u_t}, \widetilde{G}_N^\epsilon) = 1$ for any $\beta \in \mathcal{B}_N^{(\epsilon)}$. In other words, given any beliefs β in $\mathcal{B}_N^{(\epsilon)}$ about the behaviour of the other voters, the unique best response of voter-type t in the rule $\widetilde{G}_N^\epsilon$ will be $v_t = u_t^\epsilon$. In particular, this holds for $\beta = p_t$. ◇ `Claim 1`

Suppose $N \geq N_0$. For all $m \in [1 \ldots N]$, we can define a voting strategy $V_m : \mathcal{T}_m^N \longrightarrow \mathcal{U}$ by setting $V_m(t) := u_t^\epsilon$ for all $t \in \mathcal{T}_m^N$ (the best response from Claim 1). This yields a strategy profile $\mathbf{V}^N := (V_n)_{n=1}^N$.

Claim 2: *For all $m \in [1 \ldots N]$ and all $t_m \in \mathcal{T}_m^N$, define $p_{t_m} \in \Delta(\mathcal{U}^N)$ by applying formula (7) to $\mathbf{V}^N$. Then $p_{t_m} \in \mathcal{B}_N^{(\epsilon)}$.*

Proof. For any type profile $\mathbf{t} = (t_n)_{n=1}^N \in \boldsymbol{\mathcal{T}}^N$, we have $\mathbf{V}^N(\mathbf{t}) = (u_{t_n}^\epsilon)_{n=1}^N$, while $\mathbf{u_t} = (u_{t_n})_{n=1}^N$. Inequality (A19) implies that

$$\left\|\mathbf{V}^N(\mathbf{t}) - \mathbf{u_t}\right\|_\infty \quad < \quad \epsilon. \tag{A20}$$

Now suppose that $\widetilde{\mathbf{t}} = (\widetilde{t}_n)_{n=1}^N$ is a random type profile (distributed, e.g. according to the beliefs of some voter-type). Then $\mathbf{V}^N(\widetilde{\mathbf{t}}) = (u_{\widetilde{t}_n}^\epsilon)_{n=1}^N$ and $\mathbf{u}_{\widetilde{\mathbf{t}}} = (u_{\widetilde{t}_n})_{n=1}^N$ are two random utility profiles in $\mathcal{U}^N$, and inequality (A20) tells us that $\mathbf{V}^N(\widetilde{\mathbf{t}})$ is an ϵ-distortion of $\mathbf{u}_{\widetilde{\mathbf{t}}}$. In particular, fix $m \in [1 \ldots N]$ and $t_m \in \mathcal{T}_m^N$, and suppose that $\widetilde{\mathbf{t}}_{-m}$ is distributed according to the beliefs π_{t_m} of voter-type t_m, while $\mathrm{Prob}[\widetilde{t}_m = t_m] = 1$. Then in formula (7), p_{t_m} describes t_m's probabilistic beliefs about the random utility profile $\mathbf{V}(\widetilde{\mathbf{t}})$.[11] Meanwhile, let β_{t_m} be voter-type t_m's beliefs about the *true* utility profile $\mathbf{u}_{\widetilde{\mathbf{t}}}$, as defined by formula (11). As already noted, $\mathbf{V}^N(\widetilde{\mathbf{t}})$ is an ϵ-distortion of $\mathbf{u}_{\widetilde{\mathbf{t}}}$. Thus, p_{t_m} is ϵ-close to β_{t_m}. But $\beta_{t_m} \in \mathcal{B}_N$, by defining formula (12). Thus, we conclude that $p_{t_m} \in \mathcal{B}_N^{(\epsilon)}$, by defining formula (6). ◇ `Claim 2`

For any $N \geq N_0$, in the Bayesian game determined by the rule $\widetilde{G}_N^\epsilon$ and the community $\mathbf{s}_N$, Claims 1 and 2 together imply that the strategy profile $\mathbf{V}^N$ is the unique Bayesian Nash equilibrium which is consistent with beliefs in $\mathcal{B}_N^{(\epsilon)}$. Meanwhile, for all $N \in [1 \ldots N_0)$, let $\mathbf{V}^N$ be an arbitrary strategy profile. Let $\mathbb{V} := (\mathbf{V}^N)_{N=1}^\infty$; then $\mathbb{V}$ is an eventual Bayesian Nash equilibrium for $\widetilde{G}^\epsilon$ and $\mathbf{S}$.

Step 3. (This equilibrium asymptotically implements F) Let $G^\epsilon = (G_N^\epsilon)_{N=1}^\infty$ be the ordinal voting rule on $\mathcal{A}_\epsilon$ defined by formula (A13) the proof of Proposition A1. As with $\widetilde{G}^\epsilon$, we can assume that G^ϵ uses the signal-set $\mathcal{V} = \mathcal{U}^\epsilon$. But with this assumption, the function G_N^ϵ is just the restriction of the function F_N to utility profiles in $(\mathcal{U}^\epsilon)^N$. (In other words, $G_N^\epsilon(u_1, \ldots, u_N) = F_N(u_1, \ldots, u_N)$ for all $u_1, \ldots, u_N \in \mathcal{U}^\epsilon$.) By the construction in the proof of Proposition A1, $\widetilde{G}^\epsilon$ is asymptotically equal to G^ϵ; thus, for any $\delta > 0$, there is some $N_\delta \in \mathbb{N}$ such that $\mathrm{Prob}\left[\widetilde{G}_N^\epsilon(\mathbf{u}) = G_N^\epsilon(\mathbf{u})\right] > 1 - \delta$, for all $N \geq N_\delta$ and all $\mathbf{u} \in (\mathcal{U}^\epsilon)^N$. However, as we have just noted, $G_N^\epsilon(\mathbf{u}) = F_N(\mathbf{u})$ for all $\mathbf{u} \in (\mathcal{U}^\epsilon)^N$. Thus, we obtain:

$$\mathrm{Prob}\left[\widetilde{G}_N^\epsilon(\mathbf{u}) = F_N(\mathbf{u})\right] \quad > \quad 1 - \delta, \qquad \text{for all } N \geq N_\delta \text{ and all } \mathbf{u} \in (\mathcal{U}^\epsilon)^N. \tag{A21}$$

Now, for any type profile $\mathbf{t} = (t_n)_{n=1}^N$ in $\boldsymbol{\mathcal{T}}^N$, recall that $\mathbf{u_t} := (u_{t_n})_{n=1}^N \in \mathcal{U}^N$ is the corresponding profile of vNM utility functions. Define $\mathbf{u}_\mathbf{t}^\epsilon := (u_{t_n}^\epsilon)_{n=1}^N \in (\mathcal{U}^\epsilon)^N$. Then for any $N \geq N_\delta$ and all $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$,

[11] To be precise, for each utility profile $\mathbf{v} = (v_n)_{n=1}^N \in \mathcal{V}^N$, if $v_m = u_{t_m}^\epsilon$, then $p_{t_m}(\mathbf{v}) = \pi_{t_m}\{\mathbf{t}_{-m} \in \boldsymbol{\mathcal{T}}_{-m}^N;\ u_{t_n}^\epsilon = v_n$ for all $n \in [1 \ldots N] \setminus \{m\}\}$, whereas if $v_m \neq u_{t_m}^\epsilon$, then $p_{t_m}(\mathbf{v}) := 0$.

statement (A21) implies that Prob $\left[\widetilde{G}^\epsilon(\mathbf{u}^\epsilon_\mathbf{t}) = F(\mathbf{u}^\epsilon_\mathbf{t})\right] > 1-\delta$. However, for any type profile $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$, condition (R3) says that $F(\mathbf{u}^\epsilon_\mathbf{t}) = F(\mathbf{u}_\mathbf{t})$, because $\|\mathbf{u}^\epsilon_\mathbf{t} - \mathbf{u}_\mathbf{t}\|_\infty < \epsilon$ (because $\|u^\epsilon_t - u_t\|_\infty < \epsilon$ for all $t \in \mathcal{T}$). Thus, we obtain Prob $\left[\widetilde{G}^\epsilon(\mathbf{u}^\epsilon_\mathbf{t}) = F(\mathbf{u}_\mathbf{t})\right] > 1-\delta$. But $\mathbf{V}^N(\mathbf{t}) = \mathbf{u}^\epsilon_\mathbf{t}$ for all $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$; thus, we obtain Prob $\left[\widetilde{G}^\epsilon(\mathbf{V}^N(\mathbf{t})) = F(\mathbf{u}_\mathbf{t})\right] > 1-\delta$, for all $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$. Taking the infimum over all $\mathbf{t} \in \boldsymbol{\mathcal{T}}^N$ as in formula (8), we get $P(\widetilde{G}^\epsilon_N, F_N, \mathbf{s}^N, \mathbf{V}^N) \geq 1-\delta$. This holds for all $N \geq N_\delta$. We can find such an N_δ for any $\delta > 0$. Thus, we conclude that $\lim_{N\to\infty} P(\widetilde{G}^\epsilon_N, F_N, \mathbf{s}^N, \mathbf{V}^N) = 1$, as desired. □

References

Abreu, D., Matsushima, H., 1992. Virtual implementation in iteratively undominated strategies: complete information. Econometrica 60 (5), 993–1008.

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

Artemov, G., Kunimoto, T., Serrano, R., 2013. Robust virtual implementation: Toward a reinterpretation of the Wilson doctrine. Journal of Economic Theory 148 (2), 424 – 447.

Azevedo, E. M., Budish, E., August 2015. Strategy-proofness in the large. (preprint).

Birrell, E., Pass, R., 2011. Approximately strategy-proof voting. In: IJCAI'11.

Brams, S. J., Fishburn, P. C., 1983. Approval voting. Birkhäuser Boston, Mass.

Carroll, G., May 2013. A quantitative approach to incentives: Application to voting rules. (preprint).

Chierichetti, F., Kleinberg, J., 2014. Voting with limited information and many alternatives. SIAM Journal on Computing 43 (5), 1615–1653.

Dhillon, A., Mertens, J.-F., 1999. Relative utilitarianism. Econometrica 67, 471–498.

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

Ehlers, L., Peters, H., Storcken, T., 2004. Threshold strategy-proofness: on manipulability in large voting problems. Games and Economic Behavior 49 (1), 103–116.

Gibbard, A., 1973. Manipulation of voting schemes: a general result. Econometrica 41, 587–602.

Gibbard, A., 1977. Manipulation of schemes that mix voting with chance. Econometrica 45, 665–681.

Jackson, M. O., 2001. A crash course in implementation theory. Social choice and welfare 18 (4), 655–708.

Kalai, E., Smorodinsky, M., 1975. Other solutions to Nash's bargaining problem. Econometrica 43, 513–518.

Laslier, J., Weibull, J. W., 2013. An incentive-compatible Condorcet jury theorem. The Scandinavian Journal of Economics 115 (1), 84–108.

Laslier, J.-F., Sanver, M. R. (Eds.), 2010. Handbook on approval voting. Studies in Choice and Welfare. Springer, Heidelberg.

Lehtinen, A., Kuorikoski, J., 2007. Unrealistic assumptions in rational choice theory. Philosophy of the Social Sciences 37 (2), 115–138.

Leung, S., Lui, E., Pass, R., 2015. Voting with coarse beliefs. In: Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science. ACM, pp. 61–61.

Matsushima, H., 1988. A new approach to the implementation problem. Journal of Economic Theory 45 (1), 128–144.

McLennan, A., 2011. Manipulation in elections with uncertain preferences. Journal of Mathematical Economics 47, 370–375.

Merill, S., Nagel, J., 1987. The Effect of Approval Balloting on Strategic Voting Under Alternative Decision Rules. American Political Science Review 81, 509–524.

Moulin, H., 1980. On strategy-proofness and single peakedness. Public Choice 35 (4), 437–455.

Núñez, M., Laslier, J.-F., 2014. Preference intensity representation: strategic overstating in large elections. Social Choice and Welfare 42 (2), 313–340.

Núñez, M., 2014. The strategic sincerity of approval voting. Economic Theory 56 (1), 157–189.

Procaccia, A., 2010. Can approximation circumvent gibbard-satterthwaite? In: AAAI.

Renault, R., Trannoy, A., 2007. The Bayesian average voting game with a large population. Économie publique 17 (2).

Renault, R., Trannoy, A., 2011. Assessing the extent of strategic manipulation: the average vote example. SERIEs: Journal of the Spanish Economic Association 2 (4), 497–513.

Satterthwaite, M. A., 1975. Strategy-proofness and arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. Journal of Economic Theory 10, 187–217.

Serrano, R., Vohra, R., 2005. A characterization of virtual Bayesian implementation. Games and Economic Behavior. 50 (2), 312–331.

Tsetlin, I., Regenwetter, M., Grofman, B., 2003. Impartial culture maximizes the probability of majority cycles. Social Choice and Welfare 21 (3), 387–398.