Information aggregation under non-strategic voting*

Ignacio Esponda (NYU Stern)

Demian Pouzo (UC Berkeley)

August 14, 2010

Abstract

A presumed benefit of group decision-making is to select the best alternative by aggregating privately dispersed information. In reality, people often learn to make decisions based on previous experience. When the consequences of unchosen past alternatives (i.e., counterfactuals) are not observed, learning takes place from a biased sample. We investigate the extent to which information aggregation is precluded in such a learning environment. We apply the notion of a behavioral equilibrium (Esponda, 2008) to a benchmark voting game in order to formalize the assumption that players fail to account for selection bias. We then characterize equilibrium in games with a large number of players, provide necessary and sufficient conditions for information to be aggregated (and, therefore, for biases to be inconsequential in large games), and characterize optimal voting rules. Our results provide a more-nuanced view of the benefits of using group decision-making for the purpose of information aggregation.

*We thank Heski Bar-Isaac, Ernesto Dal Bó, Erik Eyster, Navin Kartik, Alessandro Lizzeri, Santiago Oliveros, Wolfgang Pesendorfer, Andy Skrzypacz, and several seminar participants for helpful comments. Esponda: 44 W Fourth Street KMC 7-76, New York, NY 10012, iesponda@stern.nyu.edu; Pouzo: 508-1 Evans Hall #3880, Berkeley, CA 94704, dpouzo@econ.berkeley.edu.

1 Introduction

One rationale for elections is that better outcomes are chosen by aggregating information that is dispersed in the population. We study settings where members of a group, such as a committee, have a particular objective (to elect or hire the best candidate, to choose the best treatment for a patient) and obtain private information (campaign advertising, personal interviews, physical exams) about how best to achieve their objective. We deviate from the literature by assuming that people learn to make decisions from past experience. In this context, counterfactuals are usually not observed and, consequently, learning suffers from a selection problem. For example, when deciding between two political parties, voters will consider the past performance of each party. While voters can judge the elected party's performance in office, they do not observe whether the losing party would have performed better or worse had it been elected.[1] Our objective is to evaluate the extent to which group decision-making aggregates information in a learning environment with unobserved counterfactuals.

The setup is a standard voting environment (e.g., Feddersen and Pesendorfer, 1997) with a non-standard behavioral assumption. Voters simultaneously decide which of two alternatives to support. The best alternative depends on the state of the world, and votes are cast after observing private signals that are correlated with the state. The outcome of the election is determined by a particular voting rule (e.g., majority voting).

The seminal result in the literature, due to Feddersen and Pesendorfer (1997), is that information is aggregated when voters play Nash equilibrium under *any* non-unanimous voting rule.[2] The assumption of Nash equilibrium is sometimes questioned due to the sophistication required from players: They must not only realize that they should vote as if their vote were pivotal, but they must also make correct inferences from the information that their vote is pivotal. In another paper (Esponda and Pouzo, 2010) we show that, provided that players understand the pivotal idea, there exists a learning foundation for Nash equilibrium that does not require players to have correct beliefs about the primitives of the game or the strategies being followed

[1] In the case of doctors deciding whether or not to perform surgery, the consequences of treating an untreated patient and of not treating a treated patient are both unobservable.

[2] With multidimensional state variables, Feddersen and Pesendorfer (1997) show that information may fail to aggregate.

by other players.

Nevertheless, it is natural to ask what would happen if players did not apply the pivotal idea in such learning environments. We assume that voters naively take information at face value, thus failing to account for the possibility that the sample from which they learn is biased. We model this behavior using the notion of a behavioral equilibrium (Esponda, 2008), which builds on the idea of a self-confirming equilibrium (Battigalli (1987), Fudenberg and Levine (1993), Dekel, Fudenberg, and Levine (2004)). In another paper (Esponda and Pouzo, 2010) we also provide a learning foundation for the notion of behavioral equilibrium in the context of voting games and show that not understanding the pivotal idea is analogous to not accounting for sample selection when learning. We also argue that even sophisticated players may decide to ignore issues of sample selection. While we believe the Nash assumption–which implicitly assumes that players can perfectly account for selection problems–to be sensible in several settings, we view our alternative behavioral assumption as a modeling device for understanding how serious the selection problem can be when counterfactuals are not observed.

We develop an approach that allows us to characterize (behavioral) equilibria of the voting game with a large number of players. We are able to find necessary and sufficient conditions for equilibrium by slightly perturbing players' payoffs and by keeping track of the average strategy that each type of player follows. The key insight leading to our characterization result is that, in the perturbed game, the probability that a player is pivotal (i.e., decides the election) goes to zero as the number of players increases.

We then apply the characterization result to investigate the extent to which information aggregation (i.e., efficiency) obtains in equilibrium with sufficiently many players. On the one hand, a source of bias disappears as players become negligible because their biased decisions have a negligible impact on their own learning. On the other hand, the aggregate biased decisions of all other players do have an impact on each player's learning. We show that information may or may not be aggregated and provide necessary and sufficient conditions on the primitives (including the voting rule) under which information is aggregated as the number of players goes to infinity. We also characterize the voting rules that maximize social welfare and, in particular, provide a new rationale for optimality of majority voting in symmetric settings where players have sufficiently accurate signals.

The results that information may or may not be aggregated and that institutional details (e.g., voting rules) matter are in stark contrast to the results obtained by Feddersen and Pesendorfer (1997) under the Nash equilibrium assumption. The difference in results highlights the importance of making further progress in understanding which behavioral assumptions are appropriate in different contexts.

Our results also provide guidance to a planner who must determine whether to promote decentralized learning in committees, as opposed to, for example, promoting coordinated learning through randomized trials. We show that the welfare loss from sample-selection issues is less of a concern when the two alternatives result in similar payoffs when adopted in the states of the world in which they are best; surprisingly, the costs from choosing the wrong alternative play a relatively minor role.

For example, suppose that voters choose between two political parties, A and B. Party A is actually best if the underlying (unobservable) state of the economy is strong, while party B is best if the economy is weak. Voters get imperfect signals correlated with the state of the economy. In this case, equilibrium will be inefficient even in large elections. Roughly, the intuition is that, if equilibrium were efficient, so that party A were elected in a strong economy and party B in a weak economy, then voters would always observe party A doing better than party B (since it is easier to govern in a strong economy). Hence, all voters would prefer to vote for party A, thus contradicting the hypothesis that the right party is chosen in its corresponding state of the world. In equilibrium, party A will have to be occasionally elected into office in a weak economy; this mistake will then reduce party A's popularity and provide incentives for voters to choose both parties in equilibrium.

This paper relates to several strands of literature: voting (Austen-Smith and Banks (1996), Feddersen and Pesendorfer (1997, 1998)); information aggregation in both auctions (Milgrom (1979, 1981b), Pesendorfer and Swinkels (2000), Perry and Reny (2006)) and elections (Feddersen and Pesendorfer (1997, 1998), McLennan (1998)); and equilibrium concepts with boundedly rational players (Jehiel (2005), Eyster and Rabin (2005), Jehiel and Samet (2007), Jehiel and Koessler (2008), Esponda (2008)). We deviate from the voting literature by proposing an alternative behavioral assumption that provides a complementary view of the information-aggregation question.[3] We also provide a full characterization of (not necessarily type-symmetric)

[3]Two other behavioral assumptions have been applied to voting games. Eyster and Rabin (2005) apply their notion of (partially) cursed equilibrium to voting games, thus capturing a convex combi-

equilibria that have the property that both alternatives are chosen with non-vanishing probability, and, therefore, the potential to aggregate information.

In Section 2, we present an example that illustrates the motivation for our behavioral assumption, the relationship to other assumptions in the literature, and the intuition for some of our results. In Section 3, we present the voting stage game and the notion of equilibrium. In Section 4, we present the setup for games with many players, and in Section 5, we characterize equilibrium as the number of players goes to infinity. In Section 6, we apply these results to provide necessary and sufficient conditions for information aggregation and to characterize optimal voting rules. We briefly conclude in Section 7.

2 Motivation and examples

A group of n players chooses between alternatives A and B. The payoffs are summarized in Figure 1(a): A is best (because $2 > 1$) in state ω_A and B is best in state ω_B. Before casting their vote, players observe private signals $s \in \{a, b\}$ that are independently drawn, conditional on the state; in particular,

$$\Pr(a \mid \omega_A) = \Pr(b \mid \omega_B) = q > 1/2.$$

Hence, signal a is more favorable about ω_A than signal b and vice versa for state ω_B. After observing their signals, players simultaneously cast their vote for one of the two alternatives. The group adopts A if and only if the proportion of votes in favor of A is higher than some threshold ρ. We later generalize this setup by allowing for heterogeneity in preferences and information structure among players.

The literature has focused on two different behavioral assumptions. In the first case, players know the primitives of the game and vote for the best alternative given their information. In our example, players would vote for A after observing signal a and for B after observing b. A well known result, dating back to Condorcet (1785), states that, if signals are sufficiently precise (i.e., $q > 1/2$), then such *sincere* voting under majority rule ($\rho = 1/2$) selects the best alternative with probability that goes

nation of the sincere and Nash approaches discussed in Section 2. Costinot and Kartik (2007) study voters who are level-k thinkers (Stahl and Wilson (1995), Nagel (1995)) and show that, under homogeneous preferences, the optimal voting rule is the same regardless of whether players are sincere, Nash, level-k thinkers, or mixtures among all of these.

	ω^A	ω^B
A	2	0
B	1	1

(a)

	$\omega_1{}^A$	$\omega_2{}^A$	ω^B
A	4	2	0
B	1	1	3

(b)

Figure 1: States, outcomes, and payoffs

to 1 as the group size increases–i.e., information is aggregated.[4] But how do players learn to associate signals a and b with states ω_A and ω_B, respectively?

Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1997,1998) emphasize a different concern: Sincere voting does not necessarily constitute a Nash equilibrium of the voting game. In a Nash equilibrium, voters may sometimes vote against the best alternative, given their private information alone. The reason is that a vote is relevant only if it changes the outcome of the election, so voters should choose the alternative that is optimal, conditional on the information that they can infer from the hypothetical event that they are pivotal. Figure 2(a) illustrates this argument for voting rule $\rho > q$. If all players were voting sincerely, then a player's vote would be pivotal with vanishing probability. However, conditional on the event that a vote is pivotal, it is much more likely that the state is ω_A rather than ω_B. Therefore, a player should ignore her private information and vote as if the state were ω_A, thus contradicting that sincere voting constitutes a Nash equilibrium.

Despite sincere voting not necessarily being a Nash equilibrium, Feddersen and Pesendorfer (1997) show that information is aggregated under *any* non-unanimous voting rule when voters play a Nash equilibrium.[5] Nash behavior, however, relies on players being sophisticated enough to realize that there is information to be inferred from other players' votes, and that they should, therefore, condition their choice on

[4]By the law of large numbers, the proportion of players that observe signal a and, therefore, vote for A, concentrates around q, conditional on state ω_A, and around $1-q$, conditional on ω_B (see Figure 2(a)). Consequently, if $1-q < \rho = 1/2 < q$, then the best alternative is chosen with probability that goes to 1 in each state.

[5]To see some intuition, suppose that both alternatives were chosen with positive probability but that information were not aggregated in equilibrium, as depicted in Figure 2(b): In state ω_A, A is correctly chosen with probability that goes to 1; however, in state ω_B, both A and B are chosen with non-negligible probability. Now, the information that a vote is pivotal suggests that the state is ω_B. But then, no one would want to vote for A, contradicting the assumption that this case can arise in equilibrium.

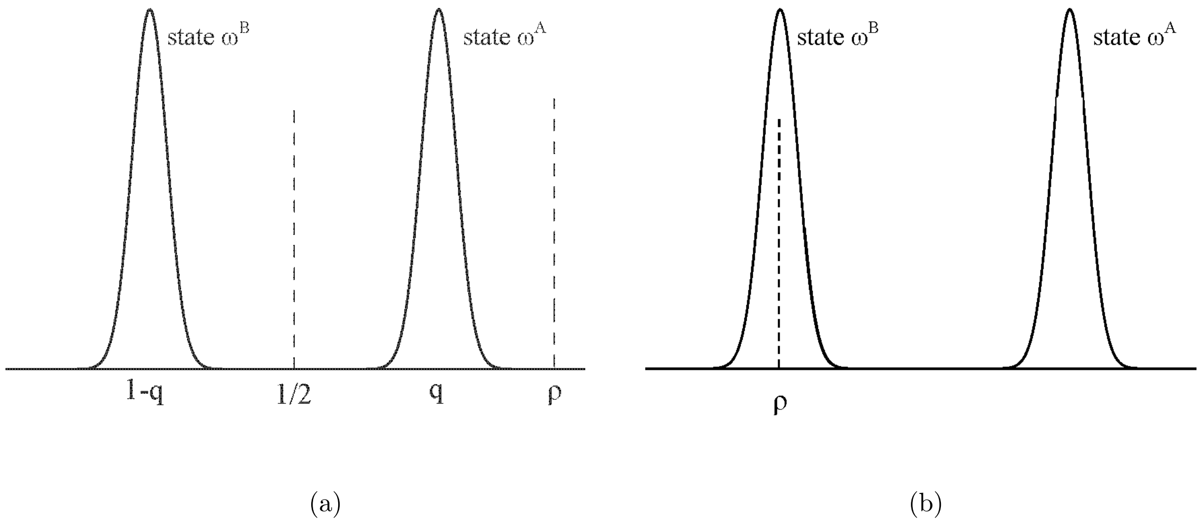

Figure 2: Comparison of sincere, Nash, and non-strategic voting

the hypothetical event that their vote is pivotal. The behavioral assumption that we postulate in this paper, in contrast, assumes that players do not account for the information content of others' actions.

The notion of a non-strategic (or naive) equilibrium is motivated by a learning environment in which players play the same stage game every period and learn to make decisions by observing the outcome of previous choices; it may be viewed as the analog of sincere voting, but when the primitives of the game must be learned (see Esponda and Pouzo (2010) for a learning foundation).

We illustrate the learning rule that provides a foundation for our solution concept by using Figure 3, which depicts a particular history of past outcomes observed by a player after playing the game for 8 periods. The private history includes her signals, her vote, the outcome of the election, and her payoff in each period. Suppose that in period 9, the player observes signal a. We postulate the following behavior. First, the player forms beliefs about the expected benefit of outcome A. These beliefs are given by the average observed payoff obtained from A when the observed signal was a, which in this case is $(0+2+2)/3 = 4/3$. Second, the player votes for the alternative that she believes has the highest expected payoffs: in this case $4/3 > 1$ and, therefore, she votes for A.

The learning rule does not take into account two sources of sample selection. The first source is exogenous: Estimates are likely to be biased upwards if alternatives

time	signal	vote	election outcome	observed payoff	
1	**a**	A	A	**0**	
2	**a**	B	A	**2**	
3	**a**	A	B	1	→ counterfactual not observed
4	**b**	B	A	0	
5	**b**	B	A	0	
6	**a**	A	A	**2**	
7	**a**	A	B	1	→ counterfactual not observed
8	**b**	B	A	2	

Figure 3: History of signals, outcomes, and payoffs.

tend to be chosen when they are most likely to be successful–which is to be expected if players use their private information to make decisions. In Figure 3, counterfactual payoffs for A are not observed in periods 3 and 7, but the fact that A was not chosen makes it likely that counterfactual payoffs would have been lower, on average, than observed payoffs for A. The second source is endogenous: A player's vote affects the sample that she will observe. For example, suppose that the player was pivotal in period 1. Then, had she voted for B instead of A, B would have been the outcome, and no payoff would have been observed for A in period 1. If all other votes were unchanged, then in period 9, the player would have even stronger beliefs of $(2+2)/2 = 2$ in favor of A. In both the exogenous and endogenous cases, the underlying source of the bias is that other players use their private information to make decisions. Failing to account for selection in a learning environment is, then, analogous to failing to account for the informational content of other players' actions.

Consider, again, the example in Figure 1(a) with voting rule $\rho > q$. We argued that sincere voting was not a Nash equilibrium; a related argument shows that sincere voting cannot be a non-strategic equilibrium either. If it were a non-strategic equilibrium, then A would be chosen with probability that goes to zero as the number of players increases. However, beliefs about the benefits from choosing A would come from those instances where A is chosen–an event that is much more likely to happen when the state is ω_A rather than ω_B. Therefore, players would mostly observe a payoff of 2 from alternative A, thus concluding that A is the best choice and contradicting that sincere voting is an equilibrium. This example highlights that what Nash and non-strategic behavior have in common is that beliefs are endogenously restricted by the strategies being followed by all players.

Nash and non-strategic behavior, however, could be fundamentally different. We argued that the situation depicted in Figure 2(b) cannot be a Nash equilibrium and that information must always be aggregated. However, in our example, information cannot be aggregated in a non-strategic equilibrium. Suppose that information were close to being aggregated in a non-strategic equilibrium: Then, players would almost always observe that alternative A has a payoff of 2, thus contradicting the assumption that any of them would ever vote for B. In fact, a non-strategic equilibrium will have the features of the situation in Figure 2(b). In order to induce players to vote for B, the committee must make enough mistakes so that a payoff of 0 is sometimes observed for A, thus counterbalancing the high payoff of 2 that is observed when A is chosen in the right state. Hence, in this example, mistakes are an inevitable feature of equilibrium outcomes.

While the example illustrates lack of information aggregation in a non-strategic equilibrium, there are cases where non-strategic behavior yields information aggregation. An example is given by the payoff structure in Figure 1(b), where there are now 3 states of the world.[6] If information were aggregated, then players would observe payoffs of 4 and 2 for alternative A and a payoff of 3 for B. Now suppose that there are two signals, and that the weighted average of 4 and 2 is higher than 3 conditional on one of the signals and lower than 3 conditional on the other. Unlike the example in Figure 1(a), players now have incentives to make both choices in equilibrium, and information will be aggregated provided that the voting rule is chosen appropriately. The rest of the paper formalizes and generalizes the arguments in this section and provides additional insights into the nature of the information-aggregation problem.

[6]Naturally, the assumption is that players do not know the structure of payoffs in Figure 1(b); otherwise, they could infer a counterfactual payoff of 1 after observing a payoff of 4 for alternative A, contradicting our motivating assumption that counterfactuals are not observed.

3 Voting framework

3.1 Voting stage game

A group of n players must choose between two alternatives, A and B. A state $\omega \in \Omega$ is first drawn according to a probability distribution G and, conditional on the state, players observe independently drawn signals s_i. In addition, players observe an idiosyncratic payoff shock, v_i, drawn independently from the state of the world and independently across players. Players then simultaneously submit a vote for either alternative, $x_i \in X = \{A, B\}$. Votes are aggregated according to a threshold voting rule: The committee chooses alternative A if and only if $k > 0$ or more players vote for A; otherwise, it chooses alternative B.

We model heterogeneity (in preferences and information) by assuming that each player i is of a particular type $\theta_i \in \Theta$, where Θ is a finite set of types.[7] Player i's utility is

$$u_{\theta_i}(o(x), \omega) + 1\{o(x) = B\} v_i,$$

where ω is the state, $v_i \in \mathbb{R}$ is a privately-observed payoff perturbation drawn from a probability distribution F_{θ_i}, and $o(x) \in \{A, B\}$ is the alternative chosen by the committee, given votes $x = (x_1, ..., x_n)$. Each signal s_i is drawn independently from a finite, totally ordered set S_{θ_i} and, conditional on the state ω, with probability $q_{\theta_i}(s_i \mid \omega) > 0$.

We make the following additional assumptions on the primitives, for all types $\theta \in \Theta$:

A1. $\Omega = [-1, 1]$ and G is an absolutely continuous probability distribution over Ω with density g.

A2. (i) $u_\theta(A, \cdot) : \Omega \to \mathbb{R}$ is nondecreasing and $u_\theta(B, \cdot) : \Omega \to \mathbb{R}$ is nonincreasing, and one of them holds strictly; (ii) $\sup_{o=\{A,B\}, \omega\in\Omega} |u_\theta(o, \omega)| < K < \infty$.

A3. There exists $z > 0$ such that for $\omega' > \omega$, and $s'_\theta > s_\theta$:

$$\frac{q_\theta(s'_\theta|\omega')}{q_\theta(s'_\theta|\omega)} - \frac{q_\theta(s_\theta|\omega')}{q_\theta(s_\theta|\omega)} = z(\omega' - \omega).$$

[7] Introducing the θ notation is redundant with a finite number of players, but it plays an important role when we take the number of players to infinity while maintaining a constant share of each type in the population.

A4. F_θ is absolutely continuous and satisfies $F_\theta(-2K) > 0$ and $F_\theta(2K) < 1$; its density f_θ satisfies $\inf_{x\in[-2K,2K]} f_\theta(x) > 0$.

A5. (i) $\inf_\Omega g(\omega) > 0$; (ii) there exists $d > 0$ such that $q_\theta(s_\theta|\omega) > d$ for all $s_\theta \in S_\theta$ and $\omega \in \Omega$; (iii) $q_\theta(s_\theta \mid \cdot)$ is continuous for all $s_\theta \in S_\theta$.

A6. $u_\theta(A, \cdot)$ and $u_\theta(B, \cdot)$ are both continuously differentiable.

A7. $u_\theta(A, 1) - Eu_\theta(B, \omega \mid s_\theta^L) > 0$ and $Eu_\theta(A, \omega \mid s_\theta^H) - u_\theta(B, -1) < 0$, where s_θ^L and s_θ^H denote the lowest and highest signals in S_θ.

Assumptions A1-A3 provide an ordering between states, information, and players' preferences, as well as a uniform bound on the utility function.[8],[9] In particular, A3 requires that the strict MLRP (monotone likelihood ration property) holds. We actually need a slight strengthening of strict MLRP: There must be a uniform bound on the rate at which the likelihood ratio changes.

Payoff perturbations (A4) play several important roles. Two of these roles are discussed by Esponda and Pouzo (2010): First, perturbations guarantee that both alternatives are chosen in equilibrium with positive probability, thus eliminating weakly dominated strategies (and, consequently, trivial equilibria where everyone votes for the same alternative) and providing the necessary experimentation to pin down beliefs in equilibrium. Second, the perturbations are important for providing a learning foundation for mixed-strategy (as opposed to correlated) equilibrium. In this paper, the payoff perturbations play two additional roles. First, payoff perturbations guarantee that the variance of the probability of voting for an alternative stays bounded away from zero. We can then apply a version of the central limit theorem to show that the probability that players are pivotal (i.e., that their vote decides the election) goes to zero. This result is crucial for characterizing equilibrium in games with many players. Second, we later consider sequences of equilibria where the perturbation vanishes. Here, perturbations can be seen as playing the standard Harsanyi (1973) role of purifying mixed strategies in games without perturbations.

Assumption A5 requires densities to be uniformly bounded (in particular, "strong

[8]It is important that $u_\theta(A, \cdot)$ and $u_\theta(B, \cdot)$ are separately increasing; it does not suffice that their difference is increasing. Bhattacharya (2008) shows that information need not be aggregated when the assumption of monotone preferences in Feddersen and Pesendorfer (1997) is relaxed.

[9]The interval state space allows us to conveniently characterize equilibrium in terms of a cutoff state above which A is chosen and below which B is chosen with probability going to 1 as the number of players increases. With a finite state space (as in the example in Section 2), we would have the same characterization, but we would also have to indicate the probability that A is chosen at the cutoff state.

signals" (Milgrom, 1979) are ruled out) and, for simplicity in the statement of results, continuity of $q_\theta(s_\theta \mid \cdot)$. Assumption A6 is for convenience but can be relaxed. Assumption A7 guarantees that, as the perturbation vanishes, there exist non-extreme equilibria where the probability of choosing each alternative is bounded away from zero for all n.

We integrate out the payoff perturbations and denote the resulting mixed strategy of player i by $\alpha_i \in \mathcal{A}_{\theta_i}$, where

$$\mathcal{A}_{\theta_i} = \{\alpha_i : F_{\theta_i}(-2K) \leq \alpha_i(s_i) \leq F_{\theta_i}(2K)\ \forall s_i \in S_{\theta_i}\}$$

is the set of player i's strategies, and $\alpha_i(s_i)$ is the probability that player i votes for A after observing signal s_i. The restriction to the set of strategies in $\mathcal{A}_{\theta_i}$ is a consequence of assuming that players know the bound K on utility (Esponda and Pouzo (2010) provide the details of this construction).

Each strategy profile $\alpha = (\alpha_1, ..., \alpha_n) \in \mathcal{A} \equiv \mathcal{A}_{\theta_1} \times ... \times \mathcal{A}_{\theta_n}$ induces a distribution over outcomes $P^n(\alpha) \in \Delta(Z^n)$, where $Z^n \equiv X^n \times S^n \times \Omega$, $S^n \equiv \prod_{i=1}^n S_{\theta_i}$, and

$$P(\alpha)(X', S', \Omega') = \sum_{x \in X'} \sum_{s \in S'} \int_{\Omega'} \left(\prod_{i=1}^n \alpha_i(s_i)^{1\{x_i = A\}} \left(1 - \alpha_i(s_i)\right)^{1\{x_i = B\}} q_{\theta_i}(s_i \mid \omega) G(d\omega) \right),$$

for any $X' \subset X^n$, $S' \subset S^n$, and $\Omega' \in \mathcal{B}$, where $\mathcal{B}$ is the Borel sigma algebra over $[-1, 1]$. Whenever an expectation E_P has a subscript P, this means that the probabilities are taken with respect to the distribution P.

3.2 Definition of equilibrium

A behavioral (naive) equilibrium (Esponda, 2008) combines the idea of a self-confirming equilibrium (Battigalli (1987), Fudenberg and Levine (1993), Dekel, Fudenberg, and Levine (2004)) with an information-processing bias. In the current voting context, Esponda and Pouzo (2010) provide a learning foundation for this equilibrium concept.

To gain some intuition for the solution concept, suppose that player i repeatedly faces a sequence of stage games where players use strategies α every period. Then, under the assumption that the payoff to alternative A is observed only whenever A is chosen, player i will come to observe that, conditional on observing signal s_i, alternative A yields in expectation $E_{P^n(\alpha)}\left[u_{\theta_i}(A, \omega) \mid o = A, s_i\right]$. A similar expression holds for alternative B.

A non-strategic (or naive) player who observes v_i and s_i believes that expected utility is maximized by voting for A whenever $\Delta_i(P^n(\alpha), s_i) - v_i > 0$ and voting for B otherwise,[10] where

$$\Delta_i(P^n, s_i) \equiv E_{P^n(\alpha)}\left[u_{\theta_i}(A,\omega) \mid o = A, s_i\right] - E_{P^n(\alpha)}\left[u_{\theta_i}(B,\omega) \mid o = B, s_i\right]$$

is well-defined because of the payoff perturbations.

Definition 1. A strategy profile $\alpha \in \mathcal{A}$ is a *(non-strategic or naive) equilibrium* of the stage game if for every player $i = 1, ..., n$ and for every $s_i \in S_i$,

$$\alpha_i(s_i) = F_i\left(\Delta_i(P^n(\alpha), s_i)\right).$$

We refer to $P^n(\alpha) \in \Delta(Z)$ as an equilibrium distribution.

In equilibrium, players best respond to beliefs that are endogenously determined by both their own strategy and those of other players and that are consistent with observed equilibrium outcomes. Non-strategic players, however, do not account for the correlation between others' votes and the state of the world (conditional on their own private information). Following a standard application of Brower's fixed-point theorem, Esponda and Pouzo (2010) show than equilibrium always exists.

4 Voting framework: a large number of players

We analyze games in which the number of players goes to infinity by studying sequences of voting games. We build such sequences by independently drawing infinite sequences of types $\xi = (\theta_1, \theta_2, ..., \theta_n, ...) \in \Xi$ according to the probability distribution $\phi \in \Delta(\Theta)$; we denote the distribution over Ξ by Φ. We interpret each sequence of types as describing an infinite number of n-player games by letting the first n elements of ξ represent the types of the n players.

[10] Implicitly, we assume that players (correctly) believe that they can be pivotal with strictly positive probability.

Let $\boldsymbol{\alpha}$ denote a *strategy mapping* from sequences of types Ξ to sequences of strategy profiles–i.e., for all $\xi \in \Xi$, let $\boldsymbol{\alpha}(\xi) = (\alpha^1(\xi), ..., \alpha^n(\xi), ...)$, where

$$\alpha^n(\xi) = (\alpha_1^n(\xi), ..., \alpha_n^n(\xi)) \in \prod_{i=1}^{n} \mathcal{A}_{\theta_i}$$

is the strategy profile that is played in the n-player game with types $\theta_1, ..., \theta_n$. Let $P^n(\boldsymbol{\alpha}(\xi))$ be the probability distribution over $X^n \times S^n \times \Omega$ induced by the strategy profile $\alpha^n(\xi)$ in the n-player game.

We define three properties of strategy mappings. The first property requires that, for large enough n, players play strategies that constitute an ε equilibrium. Our notion of equilibrium will require this property to hold for all $\varepsilon > 0$. This condition is slightly weaker than requiring that strategies constitute an equilibrium, but it allows us to obtain a full characterization of equilibrium. In particular, our result that an equilibrium is a fixed point of a particular correspondence remains true under the stronger requirement that strategies constitute an equilibrium. But the converse result, that any fixed point is also an equilibrium, relies on the notion of ε equilibrium.

Definition 2. A strategy mapping $\boldsymbol{\alpha}$ is an *ε-equilibrium mapping* if there exists n_ε such that for all $n \geq n_\varepsilon$, $i = 1, ..., n$, and $s_i \in S_i$,

$$|\alpha_i^n(\xi)(s_i) - F_i\left(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i)\right)| \leq \varepsilon \tag{1}$$

for all $\xi \in \Xi$.

The second property requires that the probabilities of choosing A and B remain bounded away from zero as the number of players increases. We will provide a full characterization of equilibria that have such a property and, therefore, the potential to aggregate information.[11]

Definition 3. A strategy mapping $\boldsymbol{\alpha}$ is $\Xi'-$*asymptotically interior* if

$$\liminf_{n\to\infty} P^n(\boldsymbol{\alpha}(\xi))\,(o = A) > 0 \quad \text{and} \quad \limsup_{n\to\infty} P^n(\boldsymbol{\alpha}(\xi))\,(o = A) < 1 \tag{2}$$

[11] Assumption A7 guarantees existence of such interior equilibria when the perturbation is small enough; for general perturbation structures, footnote 13 provides a sufficient condition.

a.s.$-\Xi'$.

The final property specifies that, as the number of players increases, the probability that the committee chooses A goes to 1 for states above a cutoff and goes to zero for states below it. We will show that equilibrium can be characterized by this convenient property.

Definition 4. A strategy mapping $\boldsymbol{\alpha}$ is Ξ'-*asymptotically* c-*cutoff* if there exists $c \in (-1, 1)$ such that

$$\lim_{n\to\infty} P^n\left(\boldsymbol{\alpha}(\xi)\right)(o = A \mid \omega) = \begin{cases} 1 & \text{for } \omega > c \\ 0 & \text{for } \omega < c \end{cases}$$

a.s.$-\Xi'$.

In addition to characterizing the equilibrium c-cutoff, our objective is to characterize the entire profile of equilibrium strategies. A complete characterization of equilibrium strategies is cumbersome due to the nature of the equilibrium object: As the number of players increases, the dimension of α^n also increases. We overcome this inconvenience by characterizing the limit, as the number of players increases, of the *average* strategy chosen by each type of player. But, unlike most of the related literature, we do *not* a priori restrict players of the same type to following the same strategy.

For a given strategy mapping $\boldsymbol{\alpha}$ and a sequence of types $\xi \in \Xi$, let $\sigma^n(\xi;\boldsymbol{\alpha}) = (\sigma^n_\theta(\xi;\boldsymbol{\alpha}))_{\theta\in\Theta} \in \mathcal{A}^* \equiv \prod_{\theta\in\Theta}\mathcal{A}_\theta$ denote the average strategy played by each type in the n-player game with types $(\theta_1, ..., \theta_n)$ and strategy profile $\alpha^n(\xi)$. Formally,

$$\sigma^n_\theta(\xi;\boldsymbol{\alpha})(s_i) = \frac{\sum_{i=1}^n 1\left\{\theta_i(\xi) = \theta\right\}\alpha^n_i(\xi)(s_i)}{\sum_{i=1}^n 1\left\{\theta_i(\xi) = \theta\right\}} \in \mathcal{A}_\theta \tag{3}$$

whenever $\sum_{i=1}^n 1\left\{\theta_i(\xi) = \theta\right\} > 0$, and arbitrary otherwise. We call any element $\sigma \in \mathcal{A}^*$ an *average strategy profile* and say that σ is *increasing* if for each type $\theta \in \Theta$, $s'_\theta > s_\theta$ implies $\sigma_\theta(s'_\theta) > \sigma_\theta(s_\theta)$.

Definition 5. An average strategy profile $\sigma \in \mathcal{A}^*$ is a *limit ε-equilibrium* if there exists $\boldsymbol{\alpha}$ and Ξ' with $\Phi(\Xi') > 0$ such that:

1. $\boldsymbol{\alpha}$ is an ε-equilibrium mapping
2. $\boldsymbol{\alpha}$ is $\Xi'-$asymptotically interior
3. $\lim_{n\to\infty} \|\sigma^n(\xi;\boldsymbol{\alpha}) - \sigma\| = 0$ for all $\xi \in \Xi'$

If, in addition, $\boldsymbol{\alpha}$ is Ξ'-asymptotically c-cutoff, then σ is a *c-cutoff limit ε-equilibrium.*

Definition 6. An average strategy profile $\sigma \in \mathcal{A}^*$ is a *[c-cutoff] limit equilibrium* if it is a [c-cutoff] limit ε-equilibrium for all $\varepsilon > 0$.

5 Characterization of equilibrium in large games

We first characterize limit equilibrium for a fixed payoff perturbation structure, and we then characterize it for vanishing perturbations.

5.1 Limit equilibrium

The intuition behind the characterization results can be grasped by thinking about a voting game with a continuum of players.[12] For a given average strategy profile $\sigma \in \mathcal{A}^*$, we interpret

$$\kappa(\sigma \mid \omega) = \sum_{\theta\in\Theta} \phi(\theta) \sum_{s_\theta \in S_\theta} q_\theta(s_\theta \mid \omega)\, \sigma_\theta(s_\theta) \tag{4}$$

as the proportion of players that vote for A conditional on ω.

For any $c \in (-1, 1)$, let

$$v_\theta(s_\theta; c) \equiv E(u_\theta(A,\omega) \mid \omega \geq c, s_\theta) - E(u_\theta(B,\omega) \mid \omega \leq c, s_\theta)$$

denote the expected difference in observed utility of type θ from alternatives A and B, conditional on signal s_θ and conditional on observing the payoff of A whenever $\omega \geq c$ and the payoff of B whenever $\omega \leq c$.

[12] Of course, if the game were really one of a continuum of players, then each player would be pivotal with probability zero and anything would constitute an equilibrium.

We will show that if σ is a limit equilibrium, there exists $c^*(\sigma) \in (-1, 1)$ such that

$$\kappa\left(\sigma \mid \omega\right) \left\{ \begin{array}{c} > \\ < \end{array} \right\} \rho \text{ if } \omega \left\{ \begin{array}{c} > \\ < \end{array} \right\} c^*(\sigma),$$

so that alternative A is chosen in states $\omega > c^*(\sigma)$ and alternative B is chosen in states $\omega < c^*(\sigma)$. We can, therefore, interpret $v_\theta\left(s; c^*(\sigma)\right)$ as type θ's belief about the difference in expected payoff from A and B in a limit equilibrium σ. A limit equilibrium average strategy of type θ will, therefore, satisfy

$$\sigma_\theta(s_\theta) = F_\theta\left(v_\theta\left(s_\theta; c^*(\sigma)\right)\right). \tag{5}$$

Finally, $c^*(\sigma)$ will be the solution to $\kappa(\sigma \mid c^*(\sigma)) = \rho$, implying that $c^*(\sigma) = c^*$ for any limit equilibrium σ, where c^* is the unique solution to[13]

$$\sum_{\theta \in \Theta} \phi(\theta) \sum_{s_\theta \in S_\theta} q_\theta\left(s_\theta \mid c^*\right) F_\theta\left(v_\theta\left(s_\theta; c^*\right)\right) = \rho. \tag{6}$$

Therefore, equilibrium strategies and outcomes can be characterized from knowledge of the primitives by using (5) and (6) above. The remainder of the section shows that the above claims, inspired by thinking about a game with a continuum of players, are formally correct in the limit as the number of players goes to infinity.

Lemma 1. *There exists $\overline{\varepsilon} > 0$ and $\gamma(\varepsilon)$ with $\lim_{\varepsilon \to 0} \gamma(\varepsilon) = 0$ such that for all $\varepsilon < \overline{\varepsilon}$: if σ is a limit ε-equilibrium, then (i) σ is increasing and is a $c^*(\sigma)$-cutoff limit ε-equilibrium, where*

$$\kappa(\sigma \mid c^*(\sigma)) = \rho, \tag{7}$$

and (ii) for all $\theta \in \Theta$ and $s_\theta \in S_\theta$,

$$\left|\sigma_\theta(s_\theta) - F_\theta\left(v_\theta\left(s_\theta; c^*(\sigma)\right)\right)\right| \leq \gamma(\varepsilon).$$

[13] As shown formally in the proofs, these statements are correct as long as both alternatives are chosen in the limit with strictly positive probability; a sufficient condition is that the LHS of equation (6) is lower than ρ for $c^* = -1$ and higher than ρ for $c^* = 1$ (note that A7 guarantees that this condition holds for small enough perturbations). Further, the solution to (6) is unique because the LHS is increasing in c^* (Claim 2.1 in the Appendix).

We now provide a discussion and proof of Lemma 1, relegating some of the details to the Appendix. The proof relies on the following Lemma.

Lemma 1.1. *Suppose that there exists* $\boldsymbol{\alpha}$ *and* Ξ' *with* $\Phi(\Xi') > 0$ *such that* $\boldsymbol{\alpha}$ *is* $\Xi'-$*asymptotically interior and for all* $\xi \in \Xi'$

$$\lim_{n\to\infty} ||\sigma^n(\xi;\boldsymbol{\alpha}) - \sigma|| = 0,$$

where σ *is increasing. Then,* $\boldsymbol{\alpha}$ *is* Ξ'*-asymptotically* $c^*(\sigma)$*-cutoff, where* $\kappa(\sigma \mid c^*(\sigma)) = \rho$*, and for all* i, s_i*,*

$$\lim_{n\to\infty} \Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i) = v_{\theta_i}(s_i; c^*(\sigma))$$

almost surely in Ξ'*.*

Proof. See the Appendix. □

The intuition of the proof is as follows. The assumption that $\sigma^n(\xi;\boldsymbol{\alpha})$ converges to σ implies, for a given ω, that the probability that a randomly chosen player votes for A converges to $\kappa(\sigma|\cdot)$. By standard asymptotic arguments, the proportion of votes for A becomes concentrated around $\kappa(\sigma|\omega)$. So, for states where $\kappa(\sigma \mid \omega) > \rho$, the probability that the outcome is A converges to 1. Similarly, for states where $\kappa(\sigma \mid \omega) < \rho$, the probability that the outcome is A converges to 0. Finally, we cannot determine what happens to the probability of choosing A for boundary states such that $\kappa(\sigma \mid \omega) = \rho$, but this is irrelevant since, by the assumption that σ is increasing, there is, at most, one (measure zero) boundary state.

The main challenge of the proof of Lemma 1 is being able to apply Lemma 1.1 by first showing that, indeed, players' equilibrium strategies are increasing in games with sufficiently many players. This challenge would not arise if voters were playing Nash equilibrium since, under assumptions A1-A3, players' strategies would always be increasing. What is different in our setting is that players' beliefs about which alternative is best does not depend only on a player's signal and the strategies of other players, but also on a player's own strategy. To understand the main issue, fix a player and a signal and suppose that she votes for A with probability close to 1. Then, most often, A is the outcome of the election whenever at least $k - 1$ or more of the other players have voted for A. Now suppose that the player votes for B with probability close to 1. Then, most often, A is the outcome of the election whenever at least k or more of the other players have voted for A. If players choose nondecreasing strategies, by MLRP, the second event conveys more favorable information about A. Therefore, the difference in expected payoffs will be decreasing in a player's own strategy. This effect goes in a direction that is opposite from the effect that a higher signal makes voting for A more desirable. The key of the next result is that for n sufficiently large, the second effect dominates the first.

Lemma 1.2. There exists $\overline{\varepsilon}$ such that for all $\varepsilon < \overline{\varepsilon}$: If σ is a limit ε-equilibrium, then it is increasing.

Proof. See the appendix. □

The key of the proof is to show that the probability that a player becomes pivotal goes to zero as n increases. Since each alternative is chosen with non-vanishing probability, the effect of a player's own strategy on her own learning must then eventually vanish and become dominated by the effect of her signal (provided a uniform version of the strict MLRP holds). The proof that players become pivotal with vanishing probability relies on the assumption that there is a payoff perturbation that bounds away from zero the probability that each individual player votes for A and B. The randomness in players' votes allows us to apply the central limit theorem to show that the proportion of players that votes for A has a limiting distribution that is continuous, and, hence, the probability that there is any specific number of votes for A must go to zero.[14]

Proof of Lemma 1: Let $\varepsilon \leq \overline{\varepsilon}$, where $\overline{\varepsilon}$ is defined by Lemma 1.2. Suppose that σ is a limit ε-equilibrium with corresponding ε-equilibrium mapping $\boldsymbol{\alpha}$ and convergence in a set Ξ'. By Lemma 1.2, σ is increasing. Therefore, all the hypotheses of Lemma 1.1 are satisfied, implying that σ is a $c^*(\sigma)$-cutoff limit ε-equilibrium, where $\kappa(\sigma \mid c^*(\sigma)) = \rho$. In addition, Lemma 1.1 implies that $\lim_{n\to\infty} \Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i) = v_{\theta_i}(s_i; c^*(\sigma))$ a.s.-Ξ'

[14]It is easy to see how the result that the probability of being pivotal vanishes would fail if the variance were zero: For example, suppose that n is even, voting is by majority rule, and half of the players vote for A and half vote for B. Then, each player is pivotal with probability 1, for all n.

and, by continuity of F_{θ_i} (A4), that $\lim_{n\to\infty} F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i)) = F_{\theta_i}(v_{\theta_i}(s_i; c^*(\sigma)))$ a.s.-Ξ'. Therefore, there exists n_ε such that for all $n \geq n_\varepsilon$, all i, s_i

$$\begin{aligned} |\alpha_i^n(\xi)(s_i) - F_{\theta_i}(v_{\theta_i}(s_i; c^*(\sigma)))| &\leq |\alpha_i^n(\xi)(s_i) - F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i))| \\ &\quad + |F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i)) - F_{\theta_i}(v_{\theta_i}(s_i; c^*(\sigma)))| \\ &\leq 2\varepsilon \end{aligned}$$

a.s.-Ξ', where for the first term in the RHS, we have used the fact that α is an ε-equilibrium mapping. Moreover, the previous inequality and equation (3) imply that for all $n \geq n_\varepsilon$, all θ, s_θ,

$$|\sigma_\theta^n(\xi; \boldsymbol{\alpha})(s_\theta) - F_\theta(v_\theta(s_\theta; c^*(\sigma)))| \leq 2\varepsilon.$$

Finally, the previous result and the fact that $\lim_{n\to\infty} \sigma^n(\xi; \boldsymbol{\alpha}) = \sigma$ for all $\xi \in \Xi'$ imply that there exists $n'_\varepsilon \geq n_\varepsilon$ such that for $n \geq n'_\varepsilon$, all θ, s_θ,

$$\begin{aligned} |\sigma_\theta(s_\theta) - F_\theta(v_\theta(s_\theta; c^*(\sigma)))| &\leq |\sigma_\theta(s_\theta) - \sigma_\theta^n(\xi; \boldsymbol{\alpha})(s_\theta)| \\ &\quad + |\sigma_\theta^n(\xi; \boldsymbol{\alpha})(s_\theta) - F_\theta(v_\theta(s_\theta; c^*(\sigma)))| \\ &\leq 3\varepsilon. \end{aligned}$$

Lemma 1 then follows by letting $\gamma(\varepsilon) = 3\varepsilon$. □

To conclude this section, we use Lemma 1 to show that the set of limit equilibria has a convenient characterization.

Theorem 1. *If σ is a limit equilibrium, then it is an increasing, c^*-cutoff limit equilibrium, where c^* solves equation (6), and equation (5) is satisfied for all $\theta \in \Theta$ and $s_\theta \in S_\theta$. If, on the other hand, σ satisfies equation (5) for all $\theta \in \Theta$ and $s_\theta \in S_\theta$, where $c^*(\sigma) = c^* \in (-1, 1)$ solves equation (6), then σ is a limit equilibrium.*

Proof. Part 1. Let σ be a limit equilibrium, so that σ is a limit ε-equilibrium for all $\varepsilon > 0$. Lemma 1 implies that (a) σ is increasing, (b) σ is a $c^*(\sigma)$-cutoff limit ε-equilibrium for all $\varepsilon > 0$, where equation (7) is satisfied, and (c) for all $\bar{\varepsilon} \geq \varepsilon > 0$, for all $\theta \in \Theta$ and $s_\theta \in S_\theta$, $|\sigma_\theta(s_\theta) - F_\theta\left(v_\theta\left(s_\theta; c^*(\sigma)\right)\right)| \leq \gamma(\varepsilon)$. Since the LHS of the inequality in part (c) does not depend on ε and $\gamma(\varepsilon) \to 0$ as $\varepsilon \to 0$, then $|\sigma_\theta(s_\theta) - F_\theta(v_\theta(s_\theta; c^*(\sigma)))| = 0$, thus establishing equation (5). Equation (6) follows by replacing equation (5) into equations (4) and (7).

Part 2. Consider the strategy mapping $\boldsymbol{\alpha}$ defined by letting players of type θ always play σ_θ–i.e., $\alpha_i(\xi)(s_i) = \sigma_{\theta_i}(s_{\theta_i})$ for all ξ, all i. First, note that $\sigma^n = \sigma$ converges trivially to σ, and σ is increasing because it satisfies equation (5) and by A4 and claim 2.1(ii) (for the latter, see the proof of Lemma 2 in the appendix). Therefore, we can follow the proof leading to equation (15) in the Appendix to obtain that $\lim_{n\to\infty} P^n(\xi)(o = A|\omega) = 1\{\omega < c^*\}$ a.s.-Ξ. The dominated convergence theorem and the fact that $c^* \in (-1, 1)$ implies that $\lim_{n\to\infty} P^n(\xi)(o = A) \in (0, 1)$, and, thus, $\boldsymbol{\alpha}$ is Ξ-asymptotically interior. Therefore, we can apply Lemma 1.1 to obtain $\lim_{n\to\infty} \Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i) = v_{\theta_i}(s_i; c^*)$ a.s.-Ξ. By continuity of F_{θ_i} (A4), $\lim_{n\to\infty} F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i)) = F_{\theta_i}(v_{\theta_i}(s_i; c^*))$. Therefore, for $\varepsilon > 0$, there exists a n_ε such that for $n \geq n_\varepsilon$, all i, s_i

$$\begin{aligned}|\alpha_i^n(\xi)(s_i) - F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i))| &= |\sigma_{\theta_i}(s_i) - F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i))| \\ &= |F_{\theta_i}(v_{\theta_i}(s_i; c^*)) - F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i))| < \varepsilon\end{aligned}$$

a.s.-Ξ. □

5.2 Vanishing perturbations

While the perturbations may have a natural interpretation in some contexts, we now consider sequences of equilibria where the perturbations vanish. We index games by a parameter η that affects the distribution F^η from which perturbations are drawn.

Definition 7. A family of perturbations $\{F^\eta\}_\eta$, where $F^\eta = \{F^\eta_\theta\}_{\theta\in\Theta}$, is *feasible* if for all $\theta \in \Theta$ and η: assumption A4 is satisfied and

$$\lim_{\eta\to 0} F^\eta_\theta(v) = \begin{cases} 0 & \text{if } v < 0 \\ 1 & \text{if } v > 0 \end{cases}$$

By Theorem 1, all limit equilibria can be characterized by the same cutoff, which solves equation (6). Let c^η denote the *limit equilibrium cutoff* that solves equation (6) when perturbations are drawn from F^η.

Definition 8. c is a *perfect limit equilibrium cutoff* if it is the limit of a sequence of limit equilibrium cutoffs $\{c^\eta\}$ for some feasible family $\{F^\eta\}_\eta$.

The final result of this section characterizes the set of perfect equilibrium cutoffs. Define

$$c_\theta(s_\theta) \equiv \arg\min_{c \in \Omega} \left| v_\theta\left(s_\theta; c\right) \right|, \tag{8}$$

and note that there is a unique solution $c_\theta(s_\theta)$ that is decreasing in s_θ (because Ω is compact, $v_\theta(s_\theta; \cdot)$ is continuous, and strict MLRP holds). For each cutoff outcome $c \in \Omega$,

$$\overline{\kappa}(c) \equiv \sum_{\theta \in \Theta} \phi_\theta q_\theta \left(\{s \in S_\theta : c_\theta(s) < c\} \mid c\right)$$

may be interpreted as the proportion of players that vote for A conditional on state c, as the perturbation vanishes.[15]

Lemma 2. $\overline{\kappa} : \Omega \to [0, 1]$ is weakly increasing and satisfies

$$\overline{\kappa}(c) = \left\{ \begin{array}{c} 0 \\ 1 \end{array} \right\} \text{ for } c \left\{ \begin{array}{l} < \min_\theta c_\theta(s_\theta^H) \\ > \max_\theta c_\theta(s_\theta^L) \end{array} \right\},$$

where $-1 < \min_\theta c_\theta(s_\theta^H) < \max_\theta c_\theta(s_\theta^L) < 1$.

Proof. See the Appendix. □

Figure 4 depicts the function $\overline{\kappa}$ for two different sets of primitives of a game. The function in panel (a) is strictly increasing and corresponds to an example with only one type, while the function in panel (b) has a flat segment and corresponds to an example with two types. In panel (a), the perfect equilibrium cutoff is given by the state c^* where $\overline{\kappa}$ intersects the voting rule ρ. In panel (b), the perfect equilibrium cutoff lies in the segment where $\overline{\kappa}$ intersects ρ; the particular point in the segment depends on the particular family of perturbations that we take.

[15] The interpretation is correct unless c is one of the cutoffs $c_\theta(s_\theta)$ for some θ, s_θ.

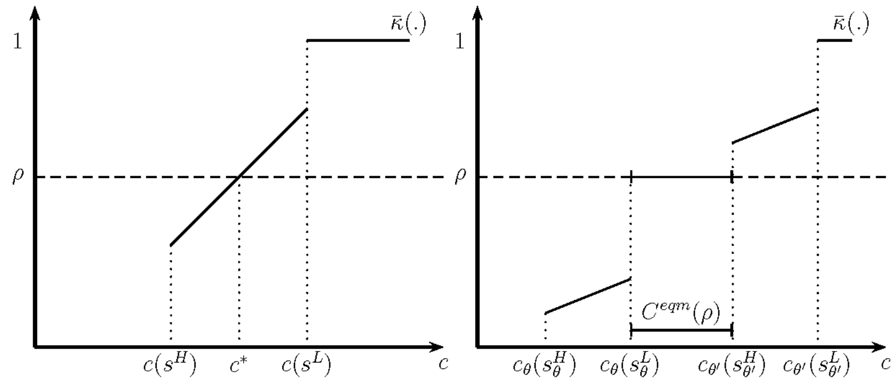

Figure 4: Characterization of perfect equilibrium cutoffs.

Theorem 2. *For a game with voting rule ρ, the set of perfect equilibrium cutoffs is given by* [16]

$$C^{eqm}(\rho) \equiv \left[\inf_c \left\{\overline{\kappa}(c) \geq \rho\right\}, \sup_c \left\{\overline{\kappa}(c) \leq \rho\right\}\right].$$

Proof. See the Appendix. □

6 Information aggregation

In this section, we consider a voting stage game with a large number of players where the perturbations vanish. We apply the results in Section 5 to obtain necessary and sufficient conditions for information aggregation and to characterize optimal voting rules. In particular, we provide a new rationale for majority voting. We then present examples that illustrate the results and provide additional insights into the conditions for information aggregation. Finally, we discuss how our results extend to the case where non-strategic players coexist with Nash players.

We carry out our analysis from the perspective of a social planner who wants to maximize aggregate welfare. Let $W(c)$ denote welfare when the outcome follows a cutoff rule c, so that alternative A is chosen for $\omega > c$ and B is chosen for $\omega < c$. We assume that $W(c)$ is single-peaked at $c = 0$ and strictly decreases as c either increases or decreases away from $c = 0$.[17]

[16] By convention, let $\sup \emptyset = -\infty$ and $\inf \emptyset = +\infty$.

[17] The welfare function $W(\cdot)$ is fairly general and consistent, for example, with the objective of maximizing a weighted average of players' utility. The assumption that $c = 0$ is the optimal cutoff

Definition 9. A voting rule ρ^* is *optimal* if there exists $c^* \in C^{eqm}(\rho^*)$ such that

$$W(c^*) \geq W(c)$$

for all $c \in \cup_{0<\rho<1} C^{eqm}(\rho)$. A voting rule ρ^* *aggregates information* if $0 \in C^{eqm}(\rho^*)$. *Information is said to be aggregated* if there exists a voting rule ρ^* that aggregates information.

Feddersen and Pesendorfer (1997) show that if the solution concept is Nash equilibrium and if the planner's preferences coincide with the preferences of the median (or any other percentile) voter, then the first-best outcome can be achieved with majority voting rule (or the corresponding percentile voting rule).[18] In our context, information may or may not be aggregated, depending on the primitives. The next result, which follows immediately from Theorem 2 and the characterization of $\overline{\kappa}$ in Lemma 2, provides necessary and sufficient conditions on the primitives such that there exists a voting rule that aggregates information.

Proposition 1. *Information is aggregated by non-strategic voters if and only if*

$$\min_\theta c_\theta(s_\theta^H) \leq 0 \leq \max_\theta c_\theta(s_\theta^L). \tag{9}$$

What makes information aggregation difficult is that players' beliefs do not depend on their equilibrium strategies *once we assume* that the outcome is the first-best outcome. In contrast, in a Nash equilibrium, beliefs depend on the event that a player is pivotal; even conditional on the first-best outcome, the pivotal event conveys information that depends on players' equilibrium strategies.

To see intuitively why (9) is necessary, suppose that $\max_\theta c_\theta(s_\theta^L) < 0$, as in Figure 5(a). If information were aggregated, then even after observing their lowest signal, all

is only for simplicity; the important assumption is that the optimal cutoff is interior.

[18] Feddersen and Pesendorfer (1997) state their main result in terms of what they call full information equivalence, meaning that for any voting rule ρ, the (Nash equilibrium) outcome of an election coincides with the outcome that would be chosen by the ρ-median voter if the state were known by all voters. In our context, full-information equivalence need not hold; therefore, we focus on finding rules that achieve the "best" outcome–hence the need to introduce the notion of a planner. Of course, Proposition 1 can be reinterpreted as providing conditions such that full information equivalence obtains given rule ρ by replacing the optimal cutoff 0 with the cutoff of the ρ-median voter.

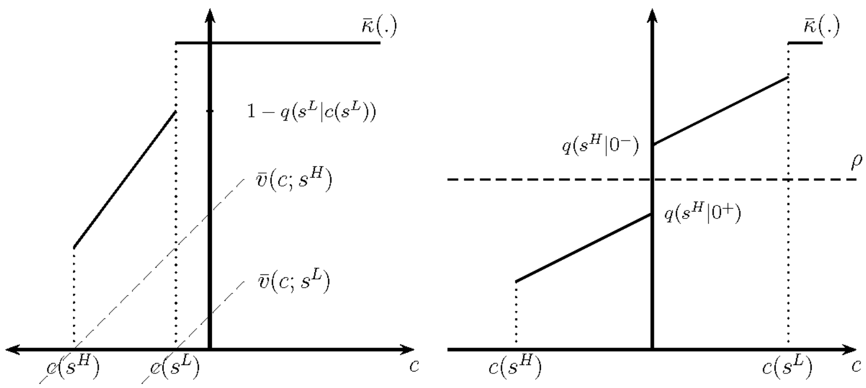

Figure 5: Information aggregation and optimal voting rules.

types would prefer to vote for A. But the fact that no one would vote for B contradicts the assumption that information is aggregated in the first place.

The next result, also an immediate implication of Theorem 2 and Lemma 2, provides a characterization of optimal voting rules. Let $\theta^0 \equiv \arg\max_\theta c_\theta(s_\theta^L)$ and $\theta_0 \equiv \arg\min_\theta c_\theta(s_\theta^H)$.

Proposition 2. *Suppose that information is not aggregated. If* $\max_\theta c_\theta(s_\theta^L) < 0$*, then voting rule* ρ *is optimal if and only if* $\rho \geq 1-\phi_{\theta^0}\cdot q_{\theta^0}\left(s_{\theta^0}^L \mid c_{\theta^0}(s_{\theta^0}^L)\right)$*; if* $\min_\theta c_\theta(s_\theta^H) > 0$*, then* ρ *is optimal if and only if* $\rho \leq \phi_{\theta_0}\cdot q_{\theta_0}\left(s_{\theta_0}^H \mid c_{\theta_0}(s_{\theta_0}^H)\right)$.

Suppose that information is aggregated. Then, voting rule ρ *is optimal if and only if* $\rho \in [\lim_{c\to 0^-}\overline{\kappa}(c), \lim_{c\to 0^+}\overline{\kappa}(c)]$.

To understand Proposition 2, consider the case where $\max_\theta c_\theta(s_\theta^L) < 0$, so that information is not aggregated because, if it were, everyone would prefer to vote for A, irrespective of their signal (see Figure 5(a)). How can we provide incentives so that some type votes for B with positive probability? Clearly, we do so by having the committee occasionally make a mistake and choose A in states of the world where B would have been best; such mistakes make B more attractive to players. But mistakes carry a welfare cost. The lowest level of this mistake that still provides incentives for some type to play B is the mistake that makes the type with the highest $c_\theta(s_\theta^L)$, defined as type θ^0, indifferent between A and B when she observes her lowest signal. Given such indifference, there is at least a proportion $1-\phi_{\theta^0}\cdot q_{\theta^0}\left(s_{\theta^0}^L \mid c_{\theta^0}(s_{\theta^0}^L)\right)$ of players who would vote for A conditional on $c_{\theta^0}(s^L_{\theta^0})$ being an equilibrium cutoff. But then, the voting rule must be higher than the previous proportion if B is to be the outcome with positive probability. In addition, voting rules that require a lower proportion to choose A also require a larger mistake in order to induce more people to vote for B, so that both A and B are chosen in equilibrium. Since larger mistakes are associated with lower welfare, such voting rules are not optimal.

Our final result provides a novel justification for optimality of majority rule: If information is sufficiently accurate, then majority rule is optimal in symmetric settings where there is only one type of player.

Definition 10. Information is *sufficiently accurate* if there exist signals $s' \neq s$ such that $q(s' \mid \omega) > 1/2$ for $\omega > 0$ and $q(s \mid \omega) > 1/2$ for $\omega < 0$.[19]

The notion of signals being sufficiently accurate can be related to Condorcet's initial praise for majority rule. Condorcet (1785) argued that, if each player votes for the right alternative with probability greater than one-half, then, as the number of players increases, the probability that the committee makes the right decision goes to 1. Translated to the voting context, the behavioral assumption in Condorcet's result is true whenever signals are sufficiently accurate and players vote for A after observing signal s' and vote for B given s. In our case, voting behavior is derived endogenously in equilibrium, and it is not necessarily true that players vote in the previous manner or that information gets aggregated. Nevertheless, majority rule is still optimal.

Proposition 3. *Consider a symmetric voting game where information is sufficiently accurate. Then, majority rule is optimal.*

Proof. Strict MLRP and the assumption that information is sufficiently accurate imply that the signals identified in Definition 10 are the highest $s^H = s'$ and lowest $s^L = s$ signals, respectively. First, consider the case where information is aggregated, so that 0 is a perfect cutoff equilibrium and Proposition 1 implies $c(s^H) \leq 0 \leq c(s^L)$.

[19] A5(iii) (continuity of q_θ) restricts S to contain only two signals if information is sufficiently accurate. However, Proposition 3 does not rely on A5(iii) (indeed, our other results easily extend if q_θ is discontinuous–e.g. the characterization of the cutoff in equation (6) needs to be slightly modified).

Let $c > 0 \geq c(s^H)$: then, $\overline{\kappa}(c) \geq q(s^H \mid c) > 1/2$, where the inequality follows from Definition 10 and the fact that $c > 0$. Similarly, let $c < 0 \leq c(s^L)$: then, $\overline{\kappa}(c) \leq 1 - q(s^L \mid c) < 1/2$, where the inequality follows from Definition 10 and the fact that $c < 0$. Proposition 2 then implies that $\rho = 1/2$ is optimal.

Finally, consider the case where information is not aggregated. If $c(s^H) > 0$, then, by Definition 10, $q(s^H \mid c(s^H)) > 1/2$. Proposition 2 then implies that $\rho = 1/2$ is optimal. Similarly, if $c(s^L) < 0$, then, by Definition 10, $1 - q(s^L \mid c(s^L)) < 1/2$. Proposition 2 then implies that $\rho = 1/2$ is optimal. □

6.1 Examples

The following examples illustrate Propositions 1-3 and provide additional insights into how the payoff and information structure relates to information aggregation. For simplicity, we discuss only examples where all players are symmetric (i.e., there is only one type).[20]

First, suppose that

$$\inf_{\omega>0} u(A,\omega) > \sup_{\omega<0} u(B,\omega), \tag{10}$$

so that alternative A dominates B when restricted to states of the world where each alternative is best. Then, $v(s;0) > 0$ for all s, implying that $c(s^L) < 0$ and, therefore, by Proposition 1, that information cannot be aggregated.[21]

For the remainder of this section, we consider a less extreme example where information aggregation is determined not only by the relative payoffs of adopting the right alternative, but also by the informativeness of the signals. The state $\omega \in [-1,1]$ is drawn from a uniform distribution and there are two signals, $\{s^L, s^H\}$, with

$$q(s^H|\omega) = \begin{cases} (0.5 + r_1\omega)^{1/r_2} & \text{if } \omega > 0 \\ (0.5 + r_1\omega)^{r_2} & \text{if } \omega < 0 \end{cases}. \tag{11}$$

[20] Our results yield additional insights when players are asymmetric. For example, information *may* be aggregated in a status quo setup when players strongly disagree about the states in which one alternative is better than the other. Thus, diversity of preferences may facilitate information aggregation. In addition, optimal voting rules will be biased against the preferences of the largest types. But if types with opposite preferences are similar in size, majority rule may again be optimal.

[21] An example that satisfies (10) is the case where B is a status quo option with a payoff that does not depend on the state of the world–i.e., $u(A,\omega) > u(B)$ for all $\omega > 0$.

Utility functions are

$$u_A(\omega) = \begin{cases} \omega^3 & \text{if } \omega \geq 0 \\ \omega^3 - h & \text{if } \omega < 0 \end{cases}$$

and $u_B(\omega) = -.5\omega^3$. Hence, alternative A does better than B, on average, but (10) does not hold.

We will vary the parameters $r_1 \in (0, 0.5)$, $r_2 \geq 1$ and $h \geq 0$ in order to emphasize different points. Suppose that the social planner has the same preferences as the players, so that the first-best cutoff is $c^* = 0$ and first-best welfare is consequently given by $W^{FB} = W(0)$. The (percentage) loss function $L(c) = (W^{FB} - W(c))/W^{FB}$ measures the percentage by which welfare deviates from the first best.

(i) *Correct payoffs and informativeness of signals.* Let $r_2 = 1$, so that $q(s^H \mid \cdot)$ is linear. At one extreme, $r_1 \approx 0$, and the signal is almost uninformative about the state. Since

$$E\left(u(A, \omega) \mid \omega \geq 0\right) > E\left(u(B, \omega) \mid \omega \leq 0\right), \tag{12}$$

information cannot be aggregated. At the other extreme, $r_1 \approx .5$ and signals are fairly informative. Conditional on $\omega > 0$, signal s^L puts a larger weight on states near 0; conditional on $\omega < 0$, signal s^L puts a higher weight on states near -1. Therefore, we may expect

$$E\left(u(A, \omega) \mid \omega \geq 0, s^L\right) < E\left(u(B, \omega) \mid \omega \leq 0, s^L\right),$$

implying that $c(s^L) > 0$ and, by Proposition 1, that information gets aggregated. In fact, there exists $r_1^* = .41$, which is the solution of $c(s^L, r_1^*) = 0$ (see equation 8), such that: for $r_1 < r_1^*$, $c(s^L, r_1) < 0$ and information is not aggregated; for $r_1 > r_1^*$, $c(s^L, r_1) > 0$ and information is aggregated. This example suggests that information aggregation obtains, provided that payoffs from adopting the right alternative are not too far from each other, that these payoffs vary in intensity depending on the state, and that there are signals that detect this variation.

(ii) *Optimal voting rules.* Let's continue to suppose that $r_2 = 1$ and let's now fix $h = 0$. Consider, first, the case where $r_1 < r_1^*$, so that information is not aggregated.[22] Figure 5(a) illustrates that the best possible equilibrium outcome is $c(s^L) < 0$, and this outcome is obtained with voting rules $\rho \geq 1 - q(s^H \mid c(s^L))$. In particular, (11)

[22]Note that h does not affect the threshold of information aggregation, r_1^*.

and the fact that $c(s^L) < 0$ imply that majority rule, $\rho = 1/2$, aggregates information.

Consider, next, the case where $r_1 > r_1^*$, so that information is aggregated. Since $\overline{\kappa}$ is continuous and $\overline{\kappa}(0) = q(s^H \mid 0) = 1/2$, Proposition 2 (see, also, Figure 5(b)) implies that majority rule is the unique optimal voting rule. Taken together, these two cases illustrate optimality of majority rule in symmetric environments (Proposition 3).

Next, we show that choosing the wrong voting rule can substantially reduce welfare in those cases where there exists a rule that aggregates information. By Theorem 2 (see, also, Figure 5(a)), the worst equilibrium outcome is given by $c(s^H) < 0$. We now compute (loss of) welfare under this worst outcome for the two extreme cases $r_1 \approx 0$ and $r_1 = 0.5$. In the first case, the signal is not informative and $c(s^H) \approx c(s^L) \approx -.33$; therefore, all voting rules lead to similar equilibrium welfare loss of $L(-.33) = .26$, or 26% of the first-best welfare. In the case where $r_1 = 0.5$, we obtain $c(s^H) = -.63$ and $L(-.63) = .95$, so that a welfare loss of 95% results from choosing the worst voting rules (compared to no welfare loss from choosing the voting rule that aggregates information).

(iii) *Type I errors.* One may conjecture that in cases where information is not aggregated, a very large payoff penalty for errors translates into a very high equilibrium cost of making wrong decisions. Nevertheless, we show that any effect of a larger type I error gets mitigated in equilibrium. The idea is that, by making mistakes costlier, a larger type I error makes it easier to provide incentives to those who obtain the lowest signal to vote for B. Thus, a higher cost of making mistakes is mitigated by a corresponding lower probability of making mistakes in equilibrium. To illustrate, suppose that $r_2 = 1$ and $r_1 = .05$. Then, $L(c(s^L; h = 0)) = .23$ and $\lim_{h\to\infty} L(c(s^L; h)) = .29$. Hence, despite the cost of the type I error going to infinity, welfare loss in an optimal equilibrium increases only from 23% to 29%.

(iv) *Between vs. within informativeness.* We compare two notions of informativeness of a signal. First, fix $r_1 \approx 0$ and note that as r_2 increases, the signals become increasingly good at distinguishing *between* the events that A is best and B is best–i.e., $\{\omega > 0\}$ and $\{\omega < 0\}$; in the limit as r_2 approaches infinity, the signals become fully revealing about which alternative is best. Second, fix $r_2 = 1$ and note that as r_1 increases, the signals are never fully revealing, but, *within* each of the events $\{\omega > 0\}$ and $\{\omega < 0\}$, they increasingly distinguish the high from the low states. Above, we showed that in this second case, there is a cutoff r_1^* above which information is aggregated. We now show that in the first case, even for very large values of r_2, information fails to aggregate. To see this, let $r_1 \approx 0$ and take $r_2 \to \infty$. Then, $q(s_L|\omega) \approx 1$ for $\omega < 0$ and $q(s_L|\omega) \approx 0$ for $\omega > 0$; within each of these two events the signal function is almost flat and, therefore, pretty uninformative. Therefore, $E\left(u(A,\omega)|\omega > 0, s_L\right) \approx E\left(u(A,\omega)|\omega > 0\right)$ and $E\left(u(B,\omega)|\omega < 0, s_L\right) \approx E\left(u(B,\omega)|\omega < 0\right)$. Equation (12) then implies that information cannot be aggregated.[23] This example reinforces the point made in (i) above: For information aggregation to obtain, the key is not so much to have signals that are very good at distinguishing whether A or B is the right alternative, but, rather, to have signals that sufficiently distinguish between states where an alternative is best by a wide margin and states where it is best by a narrow margin.

6.2 Coexistence of non-strategic and Nash players

We now illustrate how our results extend in the presence of a small fraction of Nash players, who both understand the selection problem and can perfectly account for it.

First, consider a case where information is not aggregated in the presence of non-strategic players, as in Figure 6(a). If a fraction $\gamma \approx 0$ of players is Nash and the remaining fraction $1-\gamma$ is non-strategic, the $\overline{\kappa}(c)$ function shifts proportionally downward by $(1-\gamma)$ for $c < 0$ and remains at 1 for $c > 0$. The reason is that non-strategic players behave as usual, but Nash players now vote conditional on the belief that they are pivotal. Being pivotal at a hypothetical cutoff equilibrium $c < 0$ implies that they can almost perfectly infer that the state is lower than zero; hence, for $c < 0$, Nash players vote for B irrespective of their signal. Similarly, for $c > 0$, Nash players always vote for A. The implications are the following. For most voting rules ρ, equilibrium with non-strategic players is robust to a small introduction of Nash players. However, for rules $\rho > (1-\gamma)$, equilibrium shifts from $c(s^L)$ to $c^* = 0$. We know this is true because it is a particular result in Feddersen and Pesendorfer (1997): We can interpret the non-strategic players as a large (exogenous) fraction of partisans who always vote B; a small fraction of (Nash) players who vote informatively is then sufficient to aggregate information. Therefore, rules $\rho > (1-\gamma)$ now aggregate information. This result, however, is weaker than that obtained by

[23] The above is true if $r_1 = 0$; for $r_1 > 0$ but small, r_2 has to be substantially large for information to be aggregated.

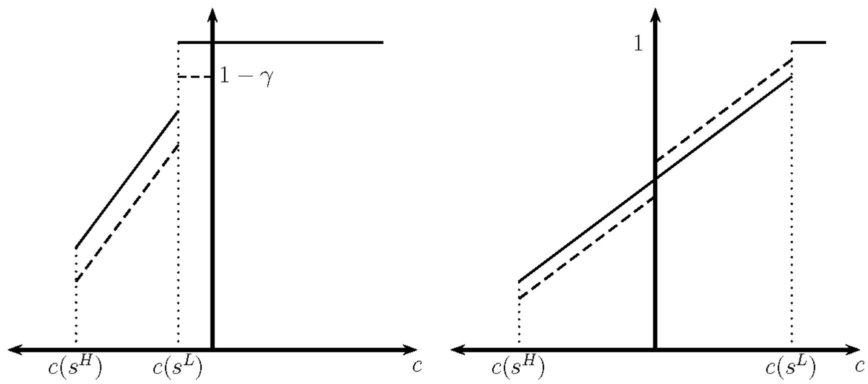

Figure 6: Coexistence of Nash and non-strategic players.

Feddersen and Pesendorfer (1997) when all players are Nash: When both Nash and non-strategic players coexist, their full information equivalence result holds only for rules $\rho > (1-\gamma)$, rather than for all voting rules.[24] By a similar argument, if $c(s^H) > 0$, then information is aggregated for rules $\rho < \gamma$. If the planner is uncertain about whether $c(s^L) < 0$ or $c(s^H) > 0$, then majority rule may remain optimal.

Second, consider a case where information is aggregated in the presence of non-strategic players, as in Figure 6(b). Again, with a fraction γ of Nash players, the $\overline{\kappa}$ function will shift downwards for $c < 0$ and upwards for $c > 0$. In particular, the figure shows that the result that majority rule is optimal in symmetric settings with sufficiently accurate signals remains true in the presence of Nash players.

7 Conclusion

We have studied the information-aggregation properties of group decision-making when people learn in a decentralized fashion and fail to account for sample selection issues. We did so by applying the notion of a behavioral (naive) equilibrium to a voting game and by characterizing equilibria as the number of players becomes large. We provided necessary and sufficient conditions in order for information to be aggregated, showing that biases at the individual level may not necessarily disappear

[24] Again, the key intuition is that the behavior of the partisans (i.e., non-strategic players) is now exogenous and will not adjust in the presence of different rules (beyond what is determined by the original $\overline{\kappa}$ function).

in large populations. We also characterized optimal voting rules and provided a new rationale for optimality of majority voting. Overall, a more nuanced view emerges about the benefits of using elections or committees in order to aggregate information.

Feddersen and Pesendorfer (1997) sparked a literature that maintains the Nash assumption but qualifies results on aggregation when some of the assumptions in their benchmark model are relaxed (e.g., costly information acquisition: Persico (2004), Martinelli (2006), Oliveros (2007), Gershkov and Szentes (2009); costly voting: Krishna and Morgan (2008)). In further work, it would be interesting to study the robustness of these results to our alternative behavioral assumption.

8 Appendix

8.1 Limit equilibrium

We let $\kappa_i^n(\xi \mid \omega) \equiv P^n(x_i = A \mid \omega)$ (we use $\kappa_{i,\omega}^n$ when ξ is omitted) be the probability that player $i = 1, ..., n$ votes for A conditional on the state being ω, and let $\kappa^n(\xi \mid \omega) \equiv \frac{1}{n}\sum_{i=1}^n \kappa_i^n(\xi \mid \omega)$ (we use κ_ω^n when ξ is omitted) be the average over all players. We also omit $\boldsymbol{\alpha}$ from the notation: $P^n(\xi) \equiv P^n(\boldsymbol{\alpha}(\xi))$ and $\sigma_\theta^n(\xi)$ denotes the average strategy profile of type θ.

8.1.1 Proof of Lemma 1.1

Recall that to show this lemma we assume that: (a) α is Ξ' asymptotically interior, (b) $\lim_{n\to\infty}\sigma_\theta^n(\xi) = \sigma_\theta$ a.s. in Ξ', and (c) σ is increasing. The proof relies on the following claims.

Claim 1.1.1: $\kappa(\sigma \mid \cdot)$ is increasing and therefore $\{\omega : \kappa(\sigma \mid \omega) = \rho\}$ is either empty or a singleton.

Proof. We show that $\kappa_\theta(\sigma \mid \cdot)$ is increasing given that σ_θ is increasing. First note that by Bayes theorem and A5(i)(ii), for all $\omega' > \omega$, for all θ, and $s'_\theta > s_\theta$

$$\frac{q_\theta(s'_\theta|\omega')}{q_\theta(s'_\theta|\omega)} > \frac{q_\theta(s_\theta|\omega')}{q_\theta(s_\theta|\omega)} \iff \frac{g_\theta(\omega'|s'_\theta)}{g_\theta(\omega'|s'_\theta)} > \frac{g_\theta(\omega|s_\theta)}{g_\theta(\omega|s_\theta)}.$$

(where g_θ is the pdf of ω given s_θ). Therefore, A3 implies the RHS. Moreover, by Proposition 1 in Milgrom (1981a), $\sum_{s<s'} q_\theta(s|\omega')$ strictly dominates (in a first order

stochastic sense) $\sum_{s<s'} q_\theta(s|\omega)$.

Note also that, casting $S_\theta = \{s_\theta^1, ..., s_\theta^{S_\theta}\}$, it follows that

$$\sum_{s_\theta \in S_\theta} \sigma_\theta(s_\theta) q_\theta(s_\theta|\omega) = \sum_{i=1}^{S_\theta} A_\theta(s_\theta^i) \left(\sum_{s \leq s_\theta^i} q_\theta(s|\omega) \right),$$

where $A_\theta(s_\theta^i) = \sigma_\theta(s_\theta^i) - \sigma_\theta(s_\theta^{i+1})$ and $A_\theta(s_\theta^{S_\theta}) = \sigma_\theta(s_\theta^{S_\theta})$. Hence

$$\sum_{s_\theta \in S_\theta} \sigma_\theta(s_\theta)\{q_\theta(s_\theta|\omega') - q_\theta(s_\theta|\omega)\} = \sum_{i=1}^{S_\theta - 1} A_\theta(s_\theta^i) \left(\sum_{s \leq s_\theta^i} q_\theta(s|\omega') - \sum_{s \leq s_\theta^i} q_\theta(s|\omega) \right).$$

Since σ is nondecreasing, $A_\theta(s_\theta^i) < 0$, then the expression above is strictly positive. Since $\phi(\theta) > 0$ all θ, the desired result follows from the construction of κ. □

Claim 1.1.2: For all $\omega \in \Omega$, $\lim_{n\to\infty} \kappa^n(\xi \mid \omega) = \kappa(\sigma \mid \omega)$ a.s. in Ξ'.

Proof. First, note that

$$\begin{aligned} \kappa^n(\xi \mid \omega) &= \frac{1}{n} \sum_{i=1}^{n} \sum_{\theta \in \Theta} \sum_{s \in S_\theta} q_\theta(s|\omega) 1\{\theta_i(\xi) = \theta\} \alpha_i^n(\xi)(s) \\ &= \sum_{\theta \in \Theta} \sum_{s \in S_\theta} q_\theta(s|\omega) \left\{ \frac{1}{n} \sum_{i=1}^{n} 1\{\theta_i(\xi) = \theta\} \alpha_i^n(\xi)(s) \right\} \\ &= \sum_{\theta \in \Theta} \sum_{s \in S_\theta} q_\theta(s|\omega) \left\{ \sigma_\theta^n(\xi)(s) \times \left(\frac{1}{n} \sum_{i=1}^{n} 1\{\theta_i(\xi) = \theta\} \right) \right\} \\ &\to \sum_{\theta \in \Theta} \sum_{s \in S_\theta} q_\theta(s|\omega) \sigma_\theta(s) \phi(\theta) = \kappa(\sigma \mid \omega), \end{aligned}$$

where convergence is a.s. in Ξ' and follows from (i) the assumption that $\lim_{n\to\infty} \sigma_\theta^n(\xi) = \sigma_\theta$ a.s. in Ξ', (ii) the strong law of large numbers applied to $\frac{1}{n}\sum_{i=1}^{n} 1\{\theta_i(\xi) = \theta\}$, and (iii) the fact that $1\{\cdot\}$ and σ_θ^n are uniformly bounded. □

Claim 1.1.3:

$$\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega) = \begin{cases} 0 & \text{if } \rho > \kappa(\sigma \mid \omega) \\ 1 & \text{if } \rho < \kappa(\sigma \mid \omega) \end{cases} \text{ a.s. in } \Xi'$$

Proof. It follows that

$$P^n(\xi)(o = A|\omega) = \Pr\left(n^{-1}\sum_{i=1}^{n} 1\{x_i^n = A\} \geq \rho \mid \omega\right)$$
$$= \Pr\left(n^{-1/2}\sum_{i=1}^{n}(1\{x_i^n = A\} - \kappa_i^n(\xi \mid \omega)) \geq \sqrt{n}(\rho - \kappa^n(\xi \mid \omega)) \mid \omega\right) \quad (13)$$

Moreover, by the Markov inequality,

$$\Pr\left(n^{-1/2}\left|\sum_{i=1}^{n}(1\{x_i^n = A\} - \kappa_i^n(\xi \mid \omega))\right| \geq \sqrt{M} \mid \omega\right) \leq (nM)^{-1}\sum_{i=1}^{n} E\left[(1\{x_i^n = A\} - \kappa_i^n(\xi \mid \omega))^2 \mid \omega\right]$$
$$\leq 4M^{-1} \quad (14)$$

goes to zero as $M \to \infty$.

Suppose that $\rho > \kappa(\sigma \mid \omega)$. By Claim 1.1.2, $\sqrt{n}\,(\rho - \kappa^n(\xi \mid \omega)) \to \infty$ a.s. in Ξ'. Therefore, by equations (13) and (14), $\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega) = 0$ a.s. in Ξ'. Similarly, if $\rho < \kappa(\sigma \mid \omega)$ then $\sqrt{n}\,(\rho - \kappa^n(\xi \mid \omega)) \to -\infty$ and $\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega) = 1$ a.s. in Ξ'. □

Proof of Lemma 1.1. First, Claim 1.1.3 and the facts that $\kappa(\sigma \mid \cdot)$ is increasing (Claim 1.1.1) and continuous (by A5(iii)) imply that there exists $c^*(\sigma) \in [-1, 1]$ such that $c^*(\sigma) \in \arg\min_{c\in[-1,1]} |\kappa(\sigma \mid c) - \rho|$ and

$$\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega) = 1\{\omega > c^*(\sigma)\} \quad \text{a.s. in } \Xi. \quad (15)$$

Suppose that $c^*(\sigma) = 1$. Then $\lim_{n\to\infty} P^n(\xi)(o = A) = 0$ a.s. in Ξ', therefore contradicting that α is asymptotically interior. A similar argument rules out $c^*(\sigma) = -1$. Therefore, $c^*(\sigma) \in (-1, 1)$, implying that $\boldsymbol{\alpha}$ is Ξ'-asymptotically $c^*(\sigma)$-cutoff and that $\kappa(\sigma \mid c^*(\sigma)) = \rho$.

Second, note that,

$$P^n(\xi)(o = A \mid \omega) = \sum_{s_i \in S_{\theta_i}} P^n(\xi)(o = A \mid \omega, s_i) q_{\theta_i}(s_i|\omega),$$

hence, under A5(ii), for any $\omega \in \Omega$ such that $\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega) = 0(= 1)$, it must be the case that $\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega, s_i) = 0(= 1)$ for all $s_i \in S_{\theta_i}$.

Therefore, a.s. in Ξ',

$$\begin{aligned}
\lim_{n\to\infty} E_{P^n(\xi)}\left(u_{\theta_i}(A,\omega) \mid o=A, s_i\right) &= \lim_{n\to\infty} \frac{\int_\Omega P^n(\xi)\left(o=A \mid \omega, s_i\right) q_{\theta_i}(s_i \mid \omega) u_{\theta_i}(A,\omega) G(d\omega)}{\int_\Omega P^n(\xi)\left(o=A \mid \omega, s_i\right) q_{\theta_i}(s_i \mid \omega) G(d\omega)} \\
&= \frac{\int_\Omega \lim_{n\to\infty} P^n(\xi)\left(o=A \mid \omega, s_i\right) q_{\theta_i}(s_i \mid \omega) u_{\theta_i}(A,\omega) G(d\omega)}{\int_\Omega \lim_{n\to\infty} P^n(\xi)\left(o=A \mid \omega, s_i\right) q_{\theta_i}(s_i \mid \omega) G(d\omega)} \\
&= \frac{\int_\Omega 1\{\omega > c^*(\sigma)\} q_{\theta_i}(s_i|\omega) u_{\theta_i}(A,\omega) G(d\omega)}{\int_\Omega 1\{\omega > c^*(\sigma)\} q_{\theta_i}(s_i|\omega) G(d\omega)} \\
&= E\left(u_{\theta_i}(A,\omega) \mid \omega \geq c^*(\sigma), s_i\right),
\end{aligned}$$

where the expectation is well-defined because A5(ii) and the fact that $\boldsymbol{\alpha}$ is asymptotically interior imply that the denominator is greater than zero, where the second line follows from the dominated convergence theorem (since u_i is assumed to be uniformly bounded), where the third line follows from Claim 1.1.3, and where the last line uses A1 to replace $1\{\omega > c^*(\sigma)\}$ by $1\{\omega \geq c^*(\sigma)\}$. $\square$

8.1.2 Proof of Lemma 1.2

Throughout the proof let Ξ' be the set in Definition 5 and fix $\xi \in \Xi'$ and a strategy mapping $\overline{\boldsymbol{\alpha}}$ such that 1.-3. in Definition 5 are satisfied. We drop ξ and $\overline{\alpha}$ from the notation, let $P^n \equiv P^n\left(\overline{\boldsymbol{\alpha}}(\xi)\right)$ and, for each strategy α_i^n, let $P_{\alpha_i}^n \equiv P^n\left(\alpha_i^n, \overline{\alpha}_{-i}^n(\xi)\right)$. The proof relies on the following claims; the proofs of the first three claims appear at the end of this section.

Claim 1.2.1: For all $\delta > 0$ and $\omega \in \Omega$, there exits $n_{\delta,\omega}$ such that for all $n \geq n_{\delta,\omega}$,

$\left|P_{\alpha_i}^n\left(o=A \mid \omega, s_i\right) - P_{\alpha_i'}^n\left(o=A \mid \omega, s_i'\right)\right| < \delta$ uniformly over $i, s_i, s_i', \alpha_i^n, \hat{\alpha}_i^n$.

Claim 1.2.2: For all $\delta > 0$ there exist n_δ such that for all $n \geq n_\delta$, $|\Delta_i(P^n, s_i) - \Delta_i(P_{\alpha_i}^n, s_i)| < \delta$ uniformly over i, s_i, α_i^n.

Claim 1.2.3: There exists $c > 0$ and n_c such that for all $n \geq n_c$, $\Delta_i\left(P_{\alpha_i}^n, s_i'\right) - \Delta_i\left(P_{\alpha_i}^n, s_i\right) \geq c$ for all i and $s_i' > s_i$ such that $\alpha_i^n(s_i') = \alpha_i^n(s_i)$.

Claim 1.2: There exists $c' > 0$ and $n_{c'}$ such that for all $n \geq n_{c'}$, $\Delta_i\left(P^n, s_i'\right) - \Delta_i\left(P^n, s_i\right) \geq c'$ for all i and $s_i' > s_i$.

Proof of Claim 1.2. Fix any α_i^n such that $\alpha_i^n(s_i') = \alpha_i^n(s_i)$. By Claims 1.2.2 and 1.2.3, for all $n \geq \max\{n_c, n_\delta\}$

$$\begin{aligned}\Delta_i\left(P^n, s_i'\right) - \Delta_i\left(P^n, s_i\right) &\geq \left(\Delta_i\left(P^n_{\alpha_i}, s_i'\right) - \delta\right) - \left(\Delta_i\left(P^n_{\alpha_i}, s_i\right) + \delta\right) \\ &\geq c - 2\delta.\end{aligned}$$

The claim follows by setting $\delta = c/4$ and $c' = c/2 > 0$. □

Proof of Lemma 1.2. The definition of ε-equilibrium (equation 1) implies that for all i, $s_i' > s_i$, $n \geq n_\varepsilon$,

$$\begin{aligned}\overline{\alpha}_i^n(s_i') - \overline{\alpha}_i^n(s_i) \geq F_i\left(\Delta_i\left(P^n, s_i'\right)\right) - F_i\left(\Delta_i\left(P^n, s_i\right)\right) - 2\varepsilon. \\ + F_i\left(\Delta_i\left(P^n, s_i\right) + c'\right) - F_i\left(\Delta_i\left(P^n, s_i\right) + c'\right),\end{aligned} \tag{16}$$

where we have added and subtracted the same term to the RHS. Let $c' > 0$ be as defined in Claim 1.2. Since F_i is absolutely continuous, then

$$F_i\left(\Delta_i\left(P^n, s_i\right) + c'\right) - F_i\left(\Delta_i\left(P^n, s_i\right)\right) = \int_{\Delta_i(P^n, s_i)}^{\Delta_i(P^n, s_i) + c'} f_i\left(t\right) dt \geq d \cdot c' \equiv c'' > 0,$$

where the inequality follows from A4. Hence, the sum of the second and fourth terms in the RHS of (16) is at least $c'' > 0$. By Claim 1.2, the sum of the first and last terms in the RHS of (16) is positive. Therefore, for all $i, s_i' > s_i$, $n \geq n_\varepsilon$,

$$\overline{\alpha}_i^n(s_i') - \overline{\alpha}_i^n(s_i) \geq c'' - 2\varepsilon > 0.$$

Since $\sigma_\theta^n(\xi, \alpha)$ are averages of the strategies, then for all $\theta, s_\theta' > s_\theta$, and $n \geq n_\varepsilon$, it follows that $\sigma_\theta^n(s_\theta') - \sigma_\theta^n(s_\theta) \geq c'' - 2\varepsilon$. Since $\lim_{n\to\infty} \sigma^n = \sigma$, then it follows that $\sigma_\theta(s_\theta') - \sigma_\theta(s_\theta) \geq c'' - 2\varepsilon > 0$, thus establishing that limit ε-equilibrium are increasing as long as $0 < \varepsilon < \overline{\varepsilon} \equiv c''/2 > 0$. □

Proof of Claim 1.2.1. The proof is divided into 3 steps.

Step 1. We first show that the probability of being pivotal goes to zero; i.e., for all $\omega \in \Omega$, for all i, $\lim_{n\to\infty} Piv_\omega^n = 0$, where

$$Piv_\omega^n \equiv P_1^n\left(o = A \mid \omega\right) - P_0^n\left(o = A \mid \omega\right),$$

where the "1" and "0" are understood as vectors of the same dimension as α_i.

By simple algebra,

$$Piv_\omega^n = P^n\left(\frac{n}{\sqrt{n-1}}K_\omega^n + \frac{\kappa_{i\omega}^n - 1}{V_\omega^n\sqrt{n-1}} \leq \frac{\sum_{j\neq i} Z_{j\omega}^n}{\sqrt{n-1}} < \frac{n}{\sqrt{n-1}}K_\omega^n + \frac{\kappa_{i\omega}^n}{V_\omega^n\sqrt{n-1}} \mid \omega\right),$$

where $Z_{j\omega}^n \equiv \frac{\{1\{x_j^n = A\} - \kappa_{j\omega}^n\}}{V_\omega^n}$, $V_\omega^n \equiv \sqrt{\frac{1}{n-1}\sum_{j\neq i}\kappa_{j,\omega}^n\left(1-\kappa_{j,\omega}^n\right)}$, and $K_\omega^n \equiv \frac{\rho - \kappa_\omega^n}{V_\omega^n}$. Note that, for a given n, $\{Z_{j\omega}^n\}_{j\neq i}$ are independent, they have zero mean and unit variance. Moreover, by Step 3 below, $\liminf_{n\to\infty} V_\omega^n > 0$, so that

$$\sum_{j\neq i} E\left[\left|\frac{Z_{j\omega}^n}{\sqrt{n-1}}\right|^3\right] \leq \frac{2}{\sqrt{n-1}\left(V_\omega^n\right)^3} \to 0 \; as \; n \to \infty,$$

Hence by Lindeberg-Feller CLT, it follows that, given ω, $\sum_{j\neq i}\frac{Z_{j\omega}^n}{\sqrt{n-1}} \Rightarrow N(0,1)$ as $n\to\infty$.

We divide the remainder of the proof in 3 cases: (a) $\frac{n}{\sqrt{n-1}}K_\omega^n \to -\infty$, (b) $\frac{n}{\sqrt{n-1}}K_\omega^n \to K \in (-\infty,\infty)$ or (c) $\frac{n}{\sqrt{n-1}}K_\omega^n \to \infty$ (if necessary, we take a subsequence that converges, which exists since $(V_\omega^n(\xi))_n$ and $(\kappa_\omega^n(\xi))_n$ are uniformly bounded).

We first explore case (a) (case (c) is symmetrical). Note that, since $\liminf_{n\to\infty} V_\omega^n > 0$, then $\frac{\kappa_{i\omega}^n}{V_\omega^n\sqrt{n-1}} \to 0$. Therefore, $\frac{n}{\sqrt{n-1}}K_\omega^n + \frac{\kappa_{i\omega}^n}{V_\omega^n\sqrt{n-1}} \to -\infty$, so that we can take $n \geq n_{M,\epsilon}$ such that $\sqrt{n}K_\omega^n + \frac{\kappa_{i\omega}^n}{V_\omega^n\sqrt{n-1}} \leq -M$, where $\mathcal{L}_N(-M) < 0.5\epsilon$ (where $\mathcal{L}_N$ is the standard Gaussian cdf) for any ϵ. Therefore, for all $\epsilon > 0$ there exists $n_{\epsilon,\omega}$ such that for all $n \geq \max\{n_{\epsilon,\omega}, n_{M,\epsilon}\}$:

$$Piv_\omega^n \leq P^n\left(\frac{\sum_{j\neq i} Z_{j,\omega}^n}{\sqrt{n-1}} < -M \mid \omega\right) \leq 0.5\epsilon + \mathcal{L}_N(-M) < \epsilon,$$

where the first inequality follows from the fact that $n \geq n_{M,\epsilon}$ and the second follows from CLT and our choice of M.

For case (b) (i.e., K finite) it follows for all $\epsilon > 0$, there exists $n_{\epsilon,\omega}$ such that for all $n \geq \max\{n_{\epsilon,\omega}, n_{\delta,\epsilon}\}$:

$$\begin{aligned} Piv_\omega^n \leq & P^n \left(\frac{n}{\sqrt{n-1}} K_\omega^n - \frac{1}{V_\omega^n \sqrt{n-1}} \leq \frac{\sum_{j \neq i} Z_{j\omega}^n}{\sqrt{n-1}} < \frac{n}{\sqrt{n-1}} K_\omega^n + \frac{1}{V_\omega^n \sqrt{n-1}} \mid \omega \right) \\ \leq & P^n \left(K - \delta < \frac{\sum_{j \neq i} Z_{j\omega}^n}{\sqrt{n-1}} \leq K + \delta \mid \omega \right) \\ \leq & 0.5\epsilon + \mathcal{L}_N \left(K - \delta < \frac{\sum_{j \neq i} Z_{j\omega}^n}{\sqrt{n-1}} \leq K + \delta \right) < \epsilon, \end{aligned}$$

where δ is such that $(V_\omega^n \sqrt{n-1})^{-1} < \delta$ for all $n \geq n_{\delta,\epsilon}$ and $\mathcal{L}_N(K+\delta) - \mathcal{L}_N(K-\delta) < 0.5\epsilon$. The second inequality follows from the CLT. We showed that for any convergent subsequence $(K_\omega^n(\xi))_n$, the associated subsequences of probabilities converge to zero, thus this result must hold for the whole sequence.

Step 2. Note that:

$$\begin{aligned} P_{\alpha_i}^n (o = A \mid \omega, s_i) = & \alpha_i^n(s_i) P_1^n (o = A \mid \omega) + (1 - \alpha_i^n(s_i)) P_0^n (o = A \mid \omega) \\ = & P_0^n (o = A \mid \omega) \\ & + \alpha_i^n(s_i) \left(P_1^n (o = A \mid \omega) - P_0^n (o = A \mid \omega) \right) \\ \equiv & P^n (o = A \mid \omega) + \alpha_i^n(s_i) Piv_\omega^n \end{aligned}$$

Therefore

$$|P_{\alpha_i}^n (o = A \mid \omega, s_i) - P_{\alpha_i'}^n (o = A \mid \omega, s_i') | \leq |\alpha_i^n(s_i) - \hat{\alpha}_i^n(s_i)| \cdot |Piv_\omega^n|.$$

By step 1, it follows that for all $n \geq n_{\delta,\omega}$: $|Piv_\omega^n| \leq \delta$. Since $|\alpha_i^n(s_i) - \hat{\alpha}_i^n(s_i)| \leq 1$ the desired result follows.

Step 3. We now show that for all $\omega \in \Omega$,

$$\liminf_{n \to \infty} \frac{1}{n-1} \sum_{j \neq i} \kappa_{j\omega}^n \left(1 - \kappa_{j\omega}^n\right) > 0. \tag{17}$$

Fix any n and $j \leq n$. By assumption, $\alpha_j^n(s_j) \in [F_j(-2K), F_j(2K)] \subset (0,1)$ for all s_j. Therefore, $0 < \kappa_{j\omega}^n < 1$ for all ω, thus implying equation (17). □

Proof of Claim 1.2.2. We prove that

$$\lim_{n\to\infty}\Big(E_{P^n}\left[u_{\theta_i}(A,\omega)\mid o=A,s_i\right]-E_{P^n_{\alpha_i}}\left[u_{\theta_i}(A,\omega)\mid o=A,s_i\right]\Big)=0;$$

the proof for $o=B$ is similar and therefore omitted. We first show that, for all i,s_i,α_i,

$$E_{P^n_{\alpha_i}}\left[u_{\theta_i}(A,\omega)\mid o=A,s_i\right]=\frac{\int_\Omega P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)q_{\theta_i}(s_i\mid\omega)u_{\theta_i}(A,\omega)G(d\omega)}{\int_\Omega P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)q_{\theta_i}(s_i\mid\omega)G(d\omega)}$$

is well-defined for sufficiently large n. Fix any i. A5(ii) and the fact that $\overline{\boldsymbol{\alpha}}$ is asymptotically interior imply that there exists $\overline{n}$ such that for all $n\geq\overline{n}$, there exists s_i^* such that

$$P^n(o=A,s_i^*)=\int_\Omega P^n(o=A\mid\omega,s_i^*)q_{\theta_i}\left(s_i^*\mid\omega\right)G(d\omega)\geq c>0,$$

which implies that $\int_\Omega P^n(o=A\mid\omega,s_i^*)G(d\omega)\geq c>0$. By Claim 1.2.1, for each s_i,α_i^n, $P^n\left(o=A\mid\omega,s_i^*\right)-P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)$ converges to zero as $n\to\infty$. Since both probabilities are bounded by one, then the dominated convergence theorem implies that $\int_\Omega\big(P^n\left(o=A\mid\omega,s_i^*\right)-P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)\big)G(d\omega)\to 0$ as $n\to\infty$, uniformly over α_i. Therefore, there exists $n_{.5c}$ such that $\sup_{\alpha_i}|\int_\Omega[P^n\left(o=A\mid\omega,s_i^*\right)-P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)]G(d\omega)|<.5c$ for all $n\geq n_{.5c}$. So for all $n\geq\max\bar{n},n_{.5c}\equiv\overline{n}_c$,

$$\int_\Omega P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)q_{\theta_i}\left(s_i\mid\omega\right)G(d\omega)\geq d\int_\Omega P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)G(d\omega)>.5dc>0.$$

Hence, $E_{P^n_{\alpha_i}}\left[u_{\theta_i}(A,\omega)\mid o=A,s_i\right]$ is well defined.

By simple algebra, and letting $\Delta P^n_{\alpha_i}(A,\omega,s_i)\equiv P^n\left(o=A\mid\omega,s_i\right)-P^n_{\alpha_i}\left(o=A\mid\omega,s_i\right)$,

$$\begin{aligned}&\Big|E_{P^n}\left[u_{\theta_i}(A,\omega)\mid o=A,s_i\right]-E_{P^n_{\alpha_i}}\left[u_{\theta_i}(A,\omega)\mid o=A,s_i\right]\Big|\\ &\leq\frac{\left|\int_\Omega\Delta P^n_{\alpha_i}(A,\omega,s_i)q_{\theta_i}(s_i\mid\omega)u_{\theta_i}\left(A,\omega\right)G(d\omega)\right|}{\int_\Omega P^n_{\alpha_i}\left(o=A\mid\omega\right)q_{\theta_i}(s_i\mid\omega)G(d\omega)}\\ &+\frac{\left|\int_\Omega\Delta P^n_{\alpha_i}(A,\omega,s_i)q_{\theta_i}(s_i\mid\omega)G(d\omega)\right|\int_\Omega P^n\left(o=A\mid\omega\right)q_{\theta_i}(s_i\mid\omega)u_{\theta_i}\left(A,\omega\right)G(d\omega)}{\int_\Omega P^n\left(o=A\mid\omega\right)q_{\theta_i}(s_i\mid\omega)G(d\omega)\int_\Omega P^n_{\alpha_i}\left(o=A\mid\omega\right)q_{\theta_i}(s_i\mid\omega)G(d\omega)}\end{aligned}$$

To establish the desired result, it is sufficient to show that each of the two absolute value terms in the numerator of the second and third line converge to zero as $n \to \infty$. However, this result follows by the dominated convergence theorem since $|u_{\theta_i}(A, \omega)| < K$, $q_{\theta_i}(s|\omega) \leq 1$, and pointwise convergence (for each ω) is obtained by Claim 1.2.1. □

Proof of Claim 1.2.3. For each $O \in \{A, B\}$: Let $g^n_{\alpha_i}(\omega \mid O, s_i) \equiv P^n_{\alpha_i}(d\omega \mid o = O, s_i)$ denote the density of ω conditional on $o = O$ and s_i, and let $G^n_{\alpha_i}(\omega \mid O, s_i) \equiv P^n_{\alpha_i}(\{\omega' \leq \omega\} \mid o = O, s_i)$ denote the cdf. Let $\Delta g^n_{\alpha_i}(\omega \mid O, s'_i, s_i) \equiv g^n_{\alpha_i}(\omega \mid O, s'_i) - g^n_{\alpha_i}(\omega \mid O, s_i)$ and $\Delta G^n_{\alpha_i}(\omega \mid O, s'_i, s_i) \equiv G^n_{\alpha_i}(\omega \mid O, s'_i) - G^n_{\alpha_i}(\omega \mid O, s_i)$.

Then

$$\begin{aligned}
\Delta_i \left(P^n_{\alpha_i}, s'_i\right) - \Delta_i \left(P^n_{\alpha_i}, s_i\right) &= \int_\Omega \left(u_{\theta_i}(A, \omega)\Delta g^n_{\alpha_i}(\omega \mid A, s'_i, s_i) - u_{\theta_i}(B, \omega)\Delta g^n_{\alpha_i}(\omega \mid B, {s'}_i, s_i)\right) d\omega \\
&= \int_\Omega \left(u'_{\theta_i}(A, \omega)\Delta G^n_{\alpha_i}(\omega \mid A, s_i, s'_i) - u'_{\theta_i}(B, \omega)\Delta G^n_{\alpha_i}(\omega \mid B, s_i, s'_i)\right) d\omega \\
&\geq \int_{\Omega^n \subset \Omega} u'_{\theta_i}(A, \omega)\Delta G^n_{\alpha_i}(\omega \mid A, s_i, s'_i) d\omega \\
&\geq c_m \cdot c_M \inf_{O \in \{A,B\}, \omega \in \Omega} u'_{\theta_i}(A, \omega) \\
&\equiv c > 0
\end{aligned}$$

for all $n \geq n'$ (where Ω^n, $c_m \cdot c_M > 0$, and n' are all defined in Claim 1.2.3.1 below), where the first line follows by definition, the second by integration by parts (note how the signals are inverted), the third by Claim 1.2.3.1(i) (see below) and the facts that that $u'_{\theta_i}(A, \omega) > 0$ and $u'_{\theta_i}(B, \omega) < 0$ for all ω, the fourth by Claim 1.2.3.1(ii), and the fifth line by the facts that $c_m \cdot c_M > 0$ and $\inf_{\omega \in \Omega} u'_{\theta_i}(A, \omega) > 0$ (because u_{θ_i} is continuously differentiable in a compact set Ω and $u'_{\theta_i}(A, \omega) > 0$ for all ω). □

Claim 1.2.3.1: For all i and $s'_i > s_i$ such that $\alpha^n_i(s_i) = \alpha^n_i(s'_i)$: (i) For all n, $\Delta G^n_{\alpha_i}(\omega \mid O, s_i, s'_i) \geq 0$ for all ω and $O \in \{A, B\}$; (ii) There exists n' and $(\Omega^n)_n$ with $\Omega^n = [l_n, u_n] \subseteq \Omega$ and $\liminf_{n\to\infty} u_n - l_n = \beta_2 > 0$ such that for all $n \geq n'$ and all $\omega^* \in \Omega^n \setminus \{-1, 1\}$,

$$\Delta G^n_{\alpha_i}(\omega \mid A, s_i, s'_i) \geq C_M > 0.$$

Proof of Claim 1.2.3.1. There exists $z > 0$ such that for all n and all $\omega' > \omega$,

$$\begin{aligned}
&g^n_{\alpha_i}(\omega' \mid O, s'_i)g^n_{\alpha_i}(\omega \mid O, s_i) - g^n_{\alpha_i}(\omega' \mid O, s_i)g^n_{\alpha_i}(\omega \mid O, s'_i) \\
&\quad = \frac{P^n_{\alpha_i}(O \mid \omega', s_i)P^n_{\alpha_i}(O \mid \omega, s_i)g(\omega')g(\omega)}{P^n_{\alpha_i}(O, s'_i)P^n_{\alpha_i}(O, s_i)} \left[q_{\theta_i}\left(s'_i \mid \omega'\right) q_{\theta_i}\left(s_i \mid \omega\right) - q_{\theta_i}\left(s_i \mid \omega'\right) q_{\theta_i}\left(s'_i \mid \omega\right)\right] \\
&\quad = z\frac{P^n_{\alpha_i}(O \mid \omega', s_i)P^n_{\alpha_i}(O \mid \omega, s_i)g(\omega')g(\omega)q_{\theta_i}\left(s'_i \mid \omega\right) q_{\theta_i}\left(s_i \mid \omega\right)(\omega' - \omega)}{P^n_{\alpha_i}(O, s'_i)P^n_{\alpha_i}(O, s_i)} \\
&\quad \geq 0
\end{aligned} \tag{18}$$

where the first line uses the fact that $P^n_{\alpha_i}(O \mid \hat{\omega}, s_i) = P^n_{\alpha_i}(O \mid \hat{\omega}, s'_i)$ for all $\hat{\omega}$ (because of conditional independence and the fact that $\alpha^n_i(s_i) = \alpha^n_i(s'_i)$), the second line follows from A3, and the third line follows because $z > 0$ and $\omega' > \omega$. Therefore, it follows from Milgrom (1981, Proposition 1) that, for all n, $\Delta G^n_{\alpha_i}(\omega \mid O, s_i, s'_i) \geq 0$ for all ω.

(ii) From the proof of Claim 1.2.2, there exists n' and $c' > 0$ such that, for all $n \geq n'$,

$$\int_\Omega P^n_{\alpha_i}(o = A \mid \omega, s_i)G(d\omega) \geq c'$$

for all i, α_i, s_i. For $a \in (0, 1)$, let

$$\omega^n_a = min\left\{\omega' : \int_{\omega \leq \omega'} P^n_{\alpha_i}(o = A \mid \omega, s_i)G(d\omega) \geq a \cdot c'\right\} \in \Omega.$$

Fix any $n \geq n'$. Then

$$c'/4 = \int_{\omega^n_{0.25} \leq \omega \leq \omega^n_{0.50}} P^n_{\alpha_i}(o = A \mid \omega, s_i)G(d\omega) \leq G\left(\omega^n_{0.50}\right) - G\left(\omega^n_{0.25}\right).$$

Therefore the fact that G has no mass points (A1) implies that $\omega^n_{0.50} - \omega^n_{0.25} \geq c_L > 0$. A similar argument establishes that $\omega^n_{0.75} - \omega^n_{0.50} \geq c_R > 0$.

Let $\Omega^n = [\omega^n_{0.50} - c_m/2, \omega^n_{0.50} + c_m/2]$, where $c_m \equiv min\{c_L, c_R\} > 0$. Then, $u_n - l_n = c_m > 0$. In addition, fix any $\omega^* \in \Omega^n$. Then, by construction,

$$\int_{\omega < \omega^* - c_m/2} P^n_{\alpha_i}(o = A \mid \omega, s_i)G(d\omega) \geq c'/4 \tag{19}$$

and

$$\int_{\omega > \omega^* + c_m/2} P^n_{\alpha_i}(o = A \mid \omega, s_i)G(d\omega) \geq c'/4. \tag{20}$$

By integrating each side of (18) twice, first with respect to $G(d\omega)$ over $\omega \leq \omega^*$ and second with respect to $G(d\omega')$ over $\omega' > \omega^*$, we obtain

$$\Delta G^n_{\alpha_i}(\omega \mid A, s_i, s_i') = \frac{z}{P^n_{\alpha_i}(A,s_i')P^n_{\alpha_i}(A,s_i)}\times$$

$$\times \int_{\omega'>\omega^*}\int_{\omega<\omega^*} P^n_{\alpha_i}(A \mid \omega', s_i)P^n_{\alpha_i}(A \mid \omega, s_i)g(\omega')g(\omega)q_{\theta_i}(s_i' \mid \omega)\, q_{\theta_i}(s_i \mid \omega)(\omega' - \omega)dG(\omega)dG(\omega')$$

$$\geq z \int_{\omega'>\omega^*+\frac{c_m}{2}}\int_{\omega<\omega^*-\frac{c_m}{2}} P^n_{\alpha_i}(A \mid \omega', s_i)P^n_{\alpha_i}(A \mid \omega, s_i)g(\omega')g(\omega)q_{\theta_i}(s_i' \mid \omega)\, q_{\theta_i}(s_i \mid \omega)(\omega' - \omega)dG(\omega)dG(\omega')$$

$$\geq z \cdot c_m \cdot d^2 \int_{\omega'>\omega^*+\frac{c_m}{2}} P^n_{\alpha_i}(A \mid \omega', s_i)G(d\omega')\int_{\omega<\omega^*-\frac{c_m}{2}} P^n_{\alpha_i}(A \mid \omega, s_i)G(d\omega)$$

$$\geq z \cdot c_m \cdot d^2 \cdot \left(\frac{c'}{4}\right)^2 \equiv c_M > 0,$$

where the first inequality follows from $P^n_{\alpha_i}(A, s_i')P^n_{\alpha_i}(A, s_i) \leq 1$, the second from A5, and the third from (19) and (20).

□

8.2 Vanishing perturbations

8.2.1 Proof of Lemma 2

Claim L2.1 (i) $v_\theta(s_\theta; \cdot)$ is increasing and continuous for all (θ, s_θ); (ii) $v_\theta(\cdot; c)$ is increasing for all $c \in \Omega$.

Proof. (i) Monotonicity of payoffs (A2(i)) and strict MLRP (A3) imply that $v_\theta(s_\theta; \cdot)$ is increasing. For continuity, it is sufficient to show that $E[u_\theta(A, \omega) \mid \omega \geq c, s_\theta] = \int_{\omega \geq c} u_\theta(A, \omega)\frac{G(d\omega|s_\theta)}{1-G(c|s_\theta)}$ is continuous (the result for $E[u_\theta(B, \omega) \mid \omega \leq c, s_\theta]$ is analogous). The result follows from the facts that (by A1 and A2) $\int_{\omega \geq c} u_\theta(A, \omega)G(d\omega \mid s_\theta)$ and $G(c|s_\theta)$ are continuous.

(ii) For $s_\theta' > s_\theta$, $G(\cdot \mid s_\theta)$ strictly dominates (in a first order stochastic sense) $G(\cdot \mid s_\theta')$ by the MLRP (A3). Since $u_\theta(A, \cdot)$ is nondecreasing then $E[u_\theta(A, \omega)|\omega \geq c, \cdot] = \int_{\omega \geq c} u_\theta(A, \omega)\frac{G(d\omega|\cdot)}{1-G(c|\cdot)}$ is nondecreasing; similarly, $E[u_\theta(B, \omega)|\omega \leq c, \cdot]$ is nonincreasing. By A2, one of them holds strictly, thus $v_\theta(\cdot; c)$ is increasing. □

Proof of Lemma 2. Let $S_\theta(c) \equiv \{s \in S_\theta : c_\theta(s) < c\}$ and define $\hat{\kappa}(c \mid \omega) \equiv \sum_{\theta \in \Theta} \phi_\theta q_\theta (S_\theta(c) \mid \omega)$. First, note that $q_\theta (S_\theta(c) \mid \omega)$ is weakly increasing in c, because the fact that $c_\theta(\cdot)$ is monotone implies that the set $S_\theta(c)$ becomes weakly larger as c increases. Second, MLRP and the fact that $S_\theta(c)$ is an interval of the form $\left[s_\theta, s_\theta^H\right]$ for some s_θ imply that $q_\theta\left(S_\theta(c) \mid \omega\right)$ is weakly increasing in ω; the weakly arises because $S_\theta(c)$ may be either $\emptyset$ or S_θ. Finally, if $c < \min_\theta c_\theta(s_\theta^H)$ then $S_\theta(c) = \emptyset$ for all θ and therefore $\hat{\kappa}\left(c \mid \omega\right) = 0$ for all ω. Similarly, if $c > \max_\theta c_\theta(s_\theta^L)$ then $S_\theta(c) = S_\theta$ for all θ and therefore $\hat{\kappa}\left(c \mid \omega\right) = 1$ for all ω. The characterization of $\overline{\kappa}(\cdot)$ then follows because $\overline{\kappa}(c) = \hat{\kappa}(c \mid c)$. Finally, note that $v_\theta\left(s_\theta; 1\right) \geq v_\theta\left(s_\theta^L; 1\right) = u_\theta(A, 1) - Eu_\theta(B, \omega \mid s_\theta^L) > 0$ for all θ, s_θ, where the last inequality follows from A7. By continuity of $v_\theta\left(s_\theta; \cdot\right)$ (Claim L2.1), it follows that $c_\theta(s_\theta^L) < 1$ for all θ. A similar proof establishes that $\min_\theta c_\theta(s_\theta^H) > -1$. □

8.2.2 Proof of Theorem 2

Claim 2.1 $\overline{\kappa}^\eta(\cdot) \equiv \sum_{\theta\in\Theta} \phi(\theta) \sum_{s_\theta \in S_\theta} q_\theta\left(s_\theta \mid \cdot\right) F_\theta^\eta\left(v_\theta\left(s_\theta; \cdot\right)\right)$ is increasing and continuous.

Proof. Continuity of $\overline{\kappa}^\eta(\cdot)$ follows from continuity of $v_\theta\left(s_\theta; \cdot\right)$ (Claim L2.1), F_θ^η (A4), and $q_\theta\left(s_\theta \mid \cdot\right)$ (A5(iii)). To show that $\overline{\kappa}^\eta(\cdot)$ is increasing, it is sufficient to establish that $\hat{\kappa}^\eta(\omega_1, \omega_2) \equiv \sum_{\theta\in\Theta} \phi(\theta) \sum_{s_\theta \in S_\theta} q_\theta\left(s_\theta \mid \omega_1\right) F_\theta^\eta\left(v_\theta\left(s_\theta; \omega_2\right)\right)$ is increasing in ω_1 and ω_2. First, by Claim L2.1 and A4, $F_\theta^\eta\left(v_\theta\left(s_\theta; \omega_2\right)\right)$ is increasing in ω_2. Second, fix ω_2 and let $\sigma_\theta(s_\theta) = F_\theta^\eta\left(v_\theta\left(s_\theta; \omega_2\right)\right)$. Note that by MLRP and because F_θ^η is increasing, then $\sigma_\theta(\cdot)$ is increasing by claim L2.1(ii) and A4. Therefore, we can apply the proof of Claim 1.1.1 to conclude that $\hat{\kappa}^\eta(\cdot, \omega_2)$ is increasing. □

Claim 2.2 $C^{eqm}(\rho) \subset (-1, 1)$.

Proof. Follows immediately from the characterization of $\overline{\kappa}(\cdot)$ in Lemma 2. □

Claim 2.3 Let $\Omega_\theta = \{\omega : c_\theta(s_\theta) = \omega, s_\theta \in S_\theta\}$, where $c_\theta(s_\theta)$ is defined in equation (8). If $\{F^\eta\}$ is feasible, then $\lim_{\eta\to 0} \overline{\kappa}^\eta(\omega) = \overline{\kappa}(\omega)$ for all $\omega \in [-1, 1] \setminus \cup_{\theta\in\Theta} \Omega_\theta$.

Proof. Take any $\omega \in [-1, 1] \setminus \cup_{\theta\in\Theta} \Omega_\theta$. Then for all $\theta \in \Theta$ and all $s_\theta \in S_\theta$ either $v_\theta(s_\theta; \omega) > 0$ or < 0. Thus, for each $\theta \in \Theta$, $\lim_{\eta\to 0} F_\theta^\eta\left(v_\theta\left(s_\theta; \omega\right)\right) = 1\{v_\theta\left(s_\theta; \omega\right) > 0\} = 1\{s_\theta : c_\theta(s) < \omega\}$. So, since S_θ and Θ are finite,

$$\begin{aligned}\lim_{\eta\to 0} \sum_{\theta\in\Theta} \phi(\theta) \sum_{s_\theta \in S_\theta} q_\theta(s_\theta|\omega) F_\theta^\eta\left(v_\theta\left(s_\theta; \omega\right)\right) &= \sum_{\theta\in\Theta} \phi(\theta) \sum_{s_\theta \in S_\theta} q_\theta(s_\theta|\omega) \lim_{\eta\to 0} F_\theta^\eta\left(v_\theta\left(s_\theta; \omega\right)\right) \\ &= \sum_{\theta\in\Theta} \phi(\theta) q_\theta(\{s \in S_\theta : c_\theta(s) < \omega\}|\omega);\end{aligned}$$

hence the desired result follows. □

Claim 2.4 (i) If $C^{eqm}(\rho) = \{c\}$ then for all feasible $\{F^\eta\}$ there exists $\overline{\eta}$ and $\{c^\eta\}$ such that $\overline{\kappa}^\eta(c^\eta) = \rho$ for all $\eta < \overline{\eta}$ and $c^\eta \to c$; (ii) If $c \in C^{eqm}(\rho)$ and $(C^{eqm}(\rho))^o \neq \emptyset$ (non-empty interior) then there exists a feasible $\{F^\eta\}$ such that $\overline{\kappa}^\eta(c) = \rho$ for all η.

Proof. Part (i). Let $c^\eta \equiv \arg\min_{\omega \in [-1,1]} |\overline{\kappa}^\eta(\omega) - \rho|$. Suppose, in order to obtain a contradiction, that $c^\eta = 1$ for all η. Then, by continuity of $\overline{\kappa}^\eta(\cdot)$, $\lim_{\eta \to 0} \overline{\kappa}^\eta(1) < \rho$. Since $c < 1$ (by Claim 2.2), then $1 \notin \cup_{\theta \in \Theta} \Omega_\theta$. Claim 2.3 then implies that $\overline{\kappa}(\omega) \leq \overline{\kappa}(1) < \rho$ for all ω, but then it follows that $c = 1$, thus contradicting Claim 2.2. Therefore, we can rule out $c^\eta = 1$ for η sufficiently low. Similarly, we can rule out $c^\eta = -1$ for η sufficiently low. Therefore, there exists $\overline{\eta}$ such that $c^\eta \in (-1, 1)$ and therefore $\overline{\kappa}^\eta(c^\eta) = \rho$ for all $\eta < \overline{\eta}$. Finally, consider a subsequence of $\{c^\eta\}$ that converges to c^*. It remains to show that $c^* = c$. Suppose, in order to obtain a contradiction, that $c^* > c$ (the case $c < c^*$ is similar). Choose $c' \notin \cup_{\theta \in \Theta} \Omega_\theta$ such that $c < c' < c^*$ (this is possible because $\cup_{\theta \in \Theta} \Omega_\theta$ has only a finite number of elements). Let $\overline{\eta}$ be such that $c^\eta > c'$ for all $\eta < \overline{\eta}$. Then $\overline{\kappa}^\eta(c') < \rho < \overline{\kappa}(c')$ for all $\eta < \overline{\eta}$, but this contradicts Claim 2.3.

Part (ii). Since $(C^{eqm}(\rho))^o \neq \{\emptyset\}$, there exist types $1 \in \Theta$ and $2 \in \Theta$ such that $C^{eqm}(\rho) = \left[c_1(s_1^L), c_2(s_2^H)\right]$, $\overline{\kappa}(c) = \rho$ for all $c \in C^{eqm}(\rho)$, and for all other types θ, $c_\theta(s_\theta) \notin C^{eqm}(\rho)$ for all $s_\theta \in S_\theta$.[25] Therefore, we can partition the type space as follows: $\Theta = \{1\} \cup \{2\} \cup \{\Theta_-\} \cup \{\Theta_+\}$, where $\theta \in \Theta_-$ iff $c_\theta(s_\theta^H) > c_2(s_2^H)$ and $\theta \in \Theta_+$ iff $c_\theta(s_\theta^L) > c_1(s_1^L)$. Fix $c \in (C^{eqm}(\rho))^o \cup \{c_2(s_2^H)\}$ (the proof for $c = c_1(s_1^L)$ is similar and therefore omitted), and note that

$$v_\theta(s_\theta; c) \leq 0 \qquad \text{for all } \theta \in \Theta_- \cup \{2\}, \text{ all } s_\theta \in S_\theta$$
$$v_\theta(s_\theta; c) > 0 \qquad \text{for all } \theta \in \Theta_+ \cup \{1\}, \text{ all } s_\theta \in S_\theta,$$

where the first (weak) inequality holds with equality if and only if $c = c_2(s_2^H)$ and $\theta = 2$.

We construct $\{F^\eta\}$ as follows. Let $z_\theta : S_\theta \to (0, 1)$ be an increasing function and let $z'_\theta : S_\theta \to (0, 1)$ be a decreasing function. Let $F_\theta^\eta(v_\theta(s_\theta; c)) = z_\theta(s_\theta)\eta$ for each $\theta \in \Theta_-$ and $s_\theta \in S_\theta$ as well as for $\theta = 2$ and all $s_2 \neq s_2^H$. In addition, let $F_\theta^\eta(v_\theta(s_\theta; c)) = 1 - z'_\theta(s_\theta)\eta$ for each $\theta \in \Theta_+$ and $s_\theta \in S_\theta$ as well as for $\theta = 1$ and all

[25] The proof of the case where there is more than one type satisfying each of these restrictions is very similar and therefore omitted.

$s_1 \neq s_1^L$. Finally, let $F_1^\eta\left(v_1\left(s_1^L;c\right)\right) = 1 - d_1\eta$, where we leave d_1 and $F_2^\eta\left(v_2\left(s_2^H;c\right)\right)$ unspecified for the moment. It follows that

$$\overline{\kappa}^\eta(c) = \phi(\Theta_-) + (B(c) - A(c))\eta - d_1 q_1(s_1^L|c)\phi(\theta_1)\eta + q_2(s_2^H|c)F_2^\eta\left(v_2\left(s_2^H;c\right)\right)\phi(\theta_2),$$

where $A(c)$ and $B(c)$ are terms that do not depend on η. By the fact that $\overline{\kappa}(c) = \rho$ for all $c \in C^{eqm}(\rho)$, it follows that $\phi(\Theta_-) = \rho$. We now specify $F_2^\eta\left(v_2\left(s_2^H;c\right)\right)$ to be such that $\overline{\kappa}^\eta(c) = \rho$, i.e., $F_2^\eta\left(v_2\left(s_2^H;c\right)\right) = D(c,d_1)\eta$, where $D(c,\cdot)$ is increasing and $\lim_{d_1\to\infty} D(c,d_1) = \infty$ for all c. Therefore, we can find $1 \leq d_1 < \infty$ such that $D(c,d_1) \geq 1$. Pick any such d_1 for our construction. Finally, let $\overline{\eta}$ be small enough such that $F_\theta^\eta\left(v_\theta\left(s_\theta;c\right)\right) \in (0,1)$ for all θ, s_θ. It then follows by construction that $\{F^\eta\}_{\eta<\overline{\eta}}$ is a feasible family of perturbations. □

Proof of Theorem 2. Part 1. Let $\{c^\eta\}$ be a sequence of limit equilibrium cutoffs that converges to c^*. Suppose, in order to obtain a contradiction, that $c^* > \sup_{\omega\in[-1,1]}\{\overline{\kappa}(\omega) \leq \rho\}$. Choose $c' \notin \cup_{\theta\in\Theta}\Omega_\theta$ such that $\sup_{\omega\in[-1,1]}\{\overline{\kappa}(\omega) \leq \rho\} < c' < c^*$ (this is possible because $\cup_{\theta\in\Theta}\Omega_\theta$ has only a finite number of elements). Then $\overline{\kappa}(c') > \rho$ and, by Claim 2.3, $\overline{\kappa}^\eta(c') > \rho$ for all η small enough. Since $\overline{\kappa}^\eta(\cdot)$ is increasing (Claim 2.1) and $c^\eta \to c^* > c'$, it follows that $\overline{\kappa}^\eta(c^\eta) > \rho$ for all η small enough. But this contradicts that c^η is a limit equilibrium cutoff, according to Theorem 1. A similar proof shows that it cannot be the case that $c^* < \inf_c\{\overline{\kappa}(c) \geq \rho\}$.

Part 2. Let $c \in C^{eqm}(\rho)$. By Claim 2.4, there exists a feasible $\{F^\eta\}_{\eta<\overline{\eta}}$ such that $\overline{\kappa}^\eta(c^\eta) = \rho$ and $c^\eta \to c$. Then, by Theorem 1, $\{c^\eta\}$ is a sequence of limit equilibrium cutoffs, so that c is a perfect limit equilibrium cutoff. □

References

[1] Austen-Smith, D. and J. Banks. 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem," American Political Science Review, 90-1, pp. 34-45.

[2] Battaglini, M., Morton, R. and T. Palfrey. Forthcoming "The Swing Voter's Curse in the Laboratory," Review of Economic Studies.

[3] Battigalli, P. 1987. "Comportamento Razionale Ed Equilibrio Nei Giochi e Nelle Situazioni Sociali," unpublished undergraduate dissertation, Bocconi University, Milano.

[4] Bhattacharya, S. 2008. "Preference Monotonicity and Information Aggregation in Elections," working paper.

[5] Condorcet, A. 1785 Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix. Imprimerie Royale, Paris.

[6] Costinot, A. and N. Kartik. 2007. "On Optimal Voting Rules under Homogeneous Preferences," working paper.

[7] Dekel, E., D. Fudenberg, and D. K. Levine. 2004. "Learning to Play Bayesian Games," Games and Economic Behavior, 46: 282-303.

[8] Esponda, I. 2008. "Behavioral Equilibrium in Economies with Adverse Selection," American Economic Review, 98(4): 1269-1291.

[9] Esponda, I. and D. Pouzo. 2010. "Learning Foundation for Equilibrium in Voting Environments," working paper.

[10] Eyster, E. and M. Rabin. 2005. "Cursed equilibrium," Econometrica 73: 1623–1672.

[11] Feddersen, T. and W. Pesendorfer. 1997. "Voting Behavior and Information Aggregation in Elections with Private Information," Econometrica, 65-5, pp. 1029-1058.

[12] Feddersen, T. and W. Pesendorfer. 1998. "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting," American Political Science Review, 92-1, pp. 23-35.

[13] Fudenberg, D. and D. K. Levine. 1993a. "Self-confirming Equilibrium," Econometrica, 61: 523-545.

[14] Gershkov, A. and Balázs S. 2009. "Optimal Voting schemes with Costly Information Acquisition," Journal of Economic Theory, 140(1): 36-68.

[15] Harsanyi, J.C. 1973. “Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points,” Inter. J. of Game Theory 2, 1-23.

[16] Jehiel, P. 2005. “Analogy-based expectation equilibrium,” Journal of Economic Theory, 123:81–104

[17] Jehiel, P. and F. Koessler. 2008. “Revisiting Games of Incomplete Information with Analogy-Based Expectations Equilibrium,” Games and Economic Behavior, 62(2): 533-57.

[18] Jehiel, P. and D. Samet. 2007. “Valuation equilibrium,” Theoretical Economics, 2:163-185.

[19] Krishna, V. and J. Morgan. 2008. “Voluntary Voting: Costs and Benefits,” working paper.

[20] Martinelli, C. 2006. "Would Rational Voters Acquire Costly Information?" Journal of Economic Theory, vol. 129, 225-251.

[21] McLennan, A. 1998 "Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents," American Political Science Review, 92-2, pp. 413-418.

[22] Milgrom, P. 1979. "A Convergence Theorem for Competitive Bidding with Differential Information," Econometrica, vol. 47(3), 679-88.

[23] Milgrom, P. 1981a. "Good News and Bad News: Representation Theorems and Applications," Bell Journal of Economics, vol. 12(2), 380-391.

[24] Milgrom, P. 1981b. “Rational Expectations, Information Acquisition, and Competitive Bidding,” Econometrica, vol. 49(4), 921-943.

[25] Nagel, R. 1995. “Unravelling in Guessing Games: An Experimental Study,” American Economic Review, 85(x): 1313-1326.

[26] Oliveros, S. 2007. "Endogenous information in committees and aggregation of information in large elections," working paper.

[27] Pesendorfer, W. and J. M. Swinkels. 1997. “The Loser’s Curse and Information Aggregation in Common Value Auctions,” Econometrica, Vol. 65, pp. 1247- 1281.

[28] Persico, N. 2004. "Committee Design with Endogenous Information," Review of Economic Studies 71(1), 165-94.

[29] Reny, P. J., and M. Perry. 2006 "Toward a Strategic Foundation for Rational Expectations Equilibrium," Econometrica, 74(5), pp. 1231-69.

[30] Selten, R. 1975. “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” International Journal of Game Theory 4: 25–55.

[31] Stahl, D. O. and P. W. Wilson. 1995. “On Players’ Models of Other Players: Theory and Experimental Evidence,” Games and Economic Behavior, 10(1): 218-254.