Information aggregation, learning, and non-strategic behavior in voting environments*

Ignacio Esponda (NYU Stern)   Demian Pouzo (UC Berkeley)

April 26, 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 what to make of their private information based on previous experience. When the consequences of alternatives that were not chosen in the past (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 present a dynamic framework that provides explicit learning foundations for our solution concept and clarifies the nature of the selection problem and our behavioral assumptions. 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.

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

*We thank Heski Bar-Isaac, Ernesto Dal Bó, Drew Fudenberg, Alessandro Lizzeri, Santiago Oliveros, Wolfgang Pesendorfer, Andy Skrzypacz, and several seminar participants for helpful comments. Esponda: iesponda@stern.nyu.edu; Pouzo: dpouzo@econ.berkeley.edu.

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 learning consequently 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 performance in office of the elected party, they do not observe whether the losing party would have performed better or worse had it been elected into office.[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. Fedderson 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 casted 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). 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 a biased sample. 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).

We provide three main contributions. First, we present a dynamic learning model that clarifies our behavioral assumptions and provides a foundation for (behavioral) equilibrium. The framework is an adaption of the model by Fudenberg and Kreps (1993); players repeatedly face the same voting environment and update their beliefs about the desirability of each of the alternatives by observing how previously chosen alternatives have fared. Our main result is that, when players do not account for the possibility that they learn from a biased sample, a steady state of the dynamic model is an equilibrium of the stage game with naive voters.

Our second contribution is to develop an approach that allows us to characterize all 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 followed by each type of player. 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.

Last, we apply the characterization result to investigate the extent to which information aggregation (i.e. efficiency) obtains in equilibrium with sufficiently many players. On 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

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

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 not be aggregated and that institutional details (e.g. voting rules) matter are in stark contrast to the well-known result, due to Fedderson and Pesendorfer (1997), that information is aggregated when voters play a Nash equilibrium under *any* non-unanimous voting rule. By implicitly assuming that players have correct beliefs about both the consequences of their equilibrium and off-equilibrium (or counterfactual) choices, the Nash solution concept assumes that players can perfectly account for selection. In a decentralized learning problem where counterfactuals are not (perfectly) observed, the issue of selection may be of first-order importance. The difference in results highlights the importance of understanding which behavioral assumptions are appropriate in different contexts. While the Nash assumption is sensible in certain settings (see Section 4.5 for a discussion of alternative behavioral assumptions), our behavioral assumption can always be viewed as a modeling device for understanding how serious the selection problem can be.

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, decentralized learning will lead to inefficient outcomes. Roughly, the intuition is that, if learning were to lead to the efficient outcome where party A is 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 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 three strands of literature: voting (Austen-Smith and Banks (1996), Feddersen and Pesendorfer (1997, 1998)), learning in games (Fudenberg and Levine (1998, 2009)), and information aggregation in both auctions (Milgrom (1979, 1981b), Pesendorfer and Swinkels (2000), Perry and Reny (2006)) and elections (Feddersen and Pesendorfer (1997, 1998)). We deviate from the voting literature by proposing an alternative behavioral assumption that provides a complementary view of the information aggregation question.[2,3] We also adapt the learning setup in Fudenberg and Kreps (1993) to players that observe private information, learn expected payoffs for each of the two alternatives (rather than learning the strategies of other players), and face a selection problem while learning. We explicitly account for the possibility of correlation in strategies and show that the correlation vanishes with time when payoffs are slightly perturbed.[4] Finally, in contrast to most of the information aggregation literature, we characterize all equilibria (not just equilibria in symmetric strategies) and provide both necessary and sufficient conditions for (epsilon) equilibrium.

In Section 2 we present an example that illustrates the motivation for our behavioral assumption, the relationship to other assumptions in the literature, and intuition for some of our results. In Section 3 we present the voting stage game and the notion of a (naive) behavioral equilibrium. In Section 4 we introduce the dynamic setting and use it to interpret and justify the solution concept. In Section 5 we present the setup for games with many players and in Section 6 we characterize equilibrium as the number of players goes to infinity. In Section 7 we apply these results to provide necessary and sufficient conditions for information aggregation and to characterize optimal voting rules. We briefly conclude in Section 8.

2 Motivation and examples

A group of n players chooses between alternatives A or B. Alternative A provides a payoff of 2 in state of the world ω_A and 0 in state ω_B, while B provides a safe payoff of 1 in both states. These payoffs are summarized in Figure 1(a): A is best 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 team adopts A if and only if the proportion of votes in favor of A is higher than some threshold ρ. We later generalize this set-up by allowing for heterogeneity in preferences and information structure among

[2]The seminal contribution by Feddersen and Pesendorfer (1997) sparked a literature that maintains the Nash assumption but qualifies results on aggregation when some of the assumptions in the 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)).

[3]Costinot and Kartik (2007) study an alternative behavioral assumption where voters are level-k thinkers (Stahl and Wilson (1995), Nagel (1995)). They 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. While level-k thinking is best interpreted as an introspective, non-equilibrium concept, our complementary approach is motivated by an explicit learning story. In our case, their robustness result does not obtain and optimal voting rules are generally different depending on whether voters play a Nash or a naive equilibrium.

[4]There is an alternative literature on learning and experimentation by multiple agents (Bolton and Harris (1999), Keller, Rady and Cripps (2005); Strulovici (forthcoming) in a voting context). That literature studies learning in an equilibrium context, while we study learning as a justification for equilibrium.

	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

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.[5] In our example, players would vote for A after observing signal a and would vote 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 to 1 as the group size increases, i.e. information is aggregated. Figure 2(a) illustrates the argument. 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 to $1 - q$ conditional on ω_B. Consequently, if $1 - q < 1/2 < q$ is the proportion required to choose A, then the probability that A is chosen converges to 1 conditional on state ω_A and to 0 conditional on ω_B. This behavioral assumption begs the question of how players learn to make the right choice with probability greater than one-half. In particular, 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 alternative that is best given their private information alone. The reason is that a vote is only relevant 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 fact 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–a simple consequence of the central limit theorem and a comparison of the vanishing tails of two normal distributions. 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

[5]Eyster and Rabin (2005) generalize this behavioral assumption and apply it to several other contexts; see also Jehiel and Koessler (2008).

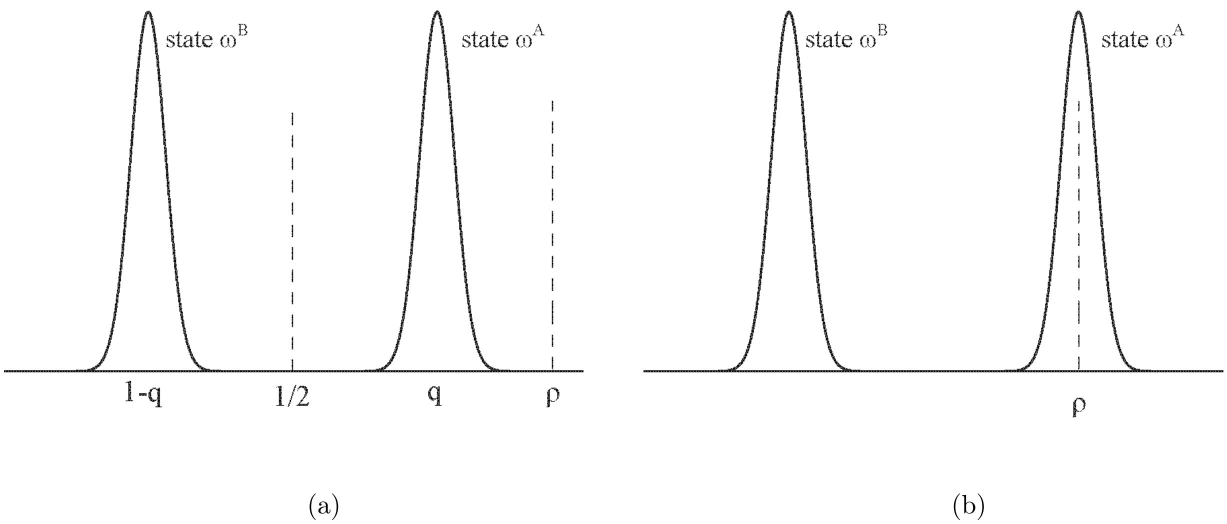

Figure 2: Comparison of sincere, Nash, and naive voting

a Nash equilibrium. 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_B, B is correctly chosen with probability that goes to 1; however, in state ω_A, both A and B are chosen with non-negligible probability. Again, the information that a vote is pivotal suggests that the state is ω_A. But then no one would want to vote for B, contradicting that this case can arise in equilibrium.[6]

The previous reasoning, 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 the hypothetical event that their vote is pivotal. In addition, players must be able to make correct inferences from the event that they are pivotal. In a static context, these inferences may require players to have correct beliefs about the primitives and about the strategies being followed by other players. In a learning context, making correct inferences will be difficult–see Section 4.5 for further discussion.

Our alternative behavioral assumption is motivated by thinking of a learning environment where 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 analogous of sincere voting but when the primitives of the game must be learned. Figure 3 depicts a particular history of past outcomes observed by a player after playing the game for 8 periods. The player remembers the observed signals, the outcome of the election, and the observed payoff in each period. Suppose than in period 9 the player observes signal a. We postulate the following behavior. First, the player forms beliefs about

[6]Alternatively, for the special case of identical preferences, McLennan (1998) shows that the fact that this is a symmetric game of common interest implies that whenever sincere voting leads to information aggregation (e.g. under majority rule), then there is a symmetric Nash equilibrium under which information is aggregated.

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.

the expected benefit of outcome A.[7] These beliefs are given by the average observed payoff obtained from A when the observed signal was a, which in this case is $(0+2+2)/3 = 4/3$.[8] Second, the player votes for the alternative that she believes has the highest expected payoffs; in this case $4/3 > 1$ and she therefore 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 tend to be chosen when they are most likely to be successful-which is to be expected if players are using 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

[7]For simplicity, we assume that the payoff from alternative B is known; the general model allows for learning about both of the alternatives.

[8]The important aspect of the belief formation process is that players consistently estimate observed mean payoffs; in particular, players could also start with a prior and apply Bayesian updating based on observed payoff outcomes.

account for the informational content of other players' actions.[9,10]

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 naive equilibrium either. If it were a naive 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 naive behavior have in common is that beliefs are endogenously restricted by the strategies being followed by all players.

Nash and naive 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, information cannot be aggregated in a naive equilibrium of our example. Suppose that information were close to being aggregated in a naive equilibrium: then players would almost always observe that alternative A has a payoff of 2, thus contradicting that any of them would ever vote in favor of B. In fact, a naive 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 frequently observed for A, thus counterbalancing the high payoff of 2 that is observed when A is chosen in the right state.

While the example illustrates lack of information aggregation in a naive equilibrium, it turns out that there are cases where naive 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.[11] 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 about the nature of the information aggregation problem.

[9]See Kagel and Levin (2002) and Charness and Levin (2006) for experimental evidence of this type of naivete in auction-like contexts, where some players do not take into account the information that can be inferred from being the winner of the auction. In a voting context, Guarnaschelli et al (2000), Ali et al (2008) and Battaglini et al (forthcoming) find that voters sometimes vote against their signal, and interpret it as evidence of strategic voting. In our learning context, naive voters may also vote against their signal. These experiments do not take place in a learning context and therefore cannot provide a direct test of our behavioral assumptions.

[10]See Aragones et al (2005), Al-Najjar (forthcoming), Al-Najjar and Pai (2009), and Schwartzstein (2009) for theoretical foundations of related forms of naivete.

[11]Naturally, the assumption is that counterfactuals are not observed, so 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.

3 Voting framework

3.1 Voting stage game

A group of n players must choose between two alternatives, $x_i \in \mathrm{X}_i = \{A, B\}$. A state ω is drawn from a finite set Ω according to a probability distribution $p \in \Delta(\Omega)$. Player i's utility is

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

where ω is the state of the world, $v_i \in V_i$ is a privately-observed payoff perturbation a la Harsanyi (1973) and Selten (1975), and $o(x) \in \{A, B\}$ is the alternative chosen by the committee given votes $x = (x_1, ..., x_n) \in \mathrm{X} \equiv \prod_{i=1}^{n} \mathrm{X}_i$.

In addition to their idiosyncratic payoff shock v_i, each player also observes a signal s_i from a finite set S_i; let $S \equiv \prod_{i=1}^{n} S_i$. Each signal s_i is drawn independently conditional on the state ω with probability $q_i(s_i \mid \omega) > 0$. After observing their private payoff-shock v_i and signal s_i, players simultaneously submit a vote for either alternative A or B. Votes are aggregated according to a threshold voting rule: alternative A is chosen by the committee if and only if $k > 0$ or more people vote for A; otherwise, alternative B is chosen.

Utility is uniformly bounded, i.e. $|u_i(o, \omega)| < K < \infty$ for all $i = 1, ..., n$, $o \in \{A, B\}$, and $\omega \in \Omega$. Moreover, the perturbation v_i is independently drawn from an absolutely continuous distribution F_i that satisfies $F_i(-2K) > 0$ and $F_i(2K) < 1$.

Let $\mathrm{Y}_i = \{A, B\}^{\#S_i}$ be the set of signal-contingent actions of player i. An *action plan* is a function $\phi_i : V_i \to Y_i$ such that $\phi_i(v_i)(s_i) = A$ for $v_i < -2K$ and $\phi_i(v_i)(s_i) = B$ for $v_i > 2K$, for all $s_i \in S_i$. An action plan indicates player i's signal-contingent action as a function of her payoff perturbation; the restriction is motivated by the bound on utility.

For each action plan ϕ_i, there is an associated (mixed) *strategy* $\alpha_i \in \mathcal{A}_i$, where

$$\mathcal{A}_i = \{\alpha_i : F_i(-2K) \leq \alpha_i(s_i) \leq F(2K)\ \forall s_i \in S_i\}$$

is the set of player i's strategies and

$$\alpha_i(s_i) = \Pr\left(\{v_i : \phi_i(v_i)(s_i) = A\}\right)$$

is the probability that player i votes for A after observing signal s_i. Each strategy profile $\alpha = (\alpha_1, ..., \alpha_n) \in \mathcal{A} \equiv \mathcal{A}_1 \times ... \times \mathcal{A}_n$ induces a distribution over outcomes $P(\alpha) \in \Delta(Z)$, where $Z \equiv X \times S \times \Omega$ and

$$P(\alpha)(x, s, \omega) = p(\omega) \prod_{i=1}^{n} \alpha_i(s_i)^{1\{x_i = A\}} \left(1 - \alpha_i(s_i)\right)^{1\{x_i = B\}} q_i(s_i \mid \omega). \tag{1}$$

Whenever an expectation E_P has a subscript P, the meaning is that the probabilities are taken with respect to the distribution P.

3.2 The role of payoff perturbations

The independent payoff perturbations play two crucial roles in this paper, in addition to the more standard roles played in Harsanyi's (1973) justification of mixed strategies and Selten's (1975) perfection refinement. Harsanyi (1973) showed that a mixed-strategy equilibrium is the limit of equilibria in a perturbed game where payoff perturbations vanish. In contrast, we first fix a game where payoff perturbations are non-zero and justify the notion of equilibrium in that stage game as the outcome of a dynamic learning game. The role of payoff perturbations is to guarantee that behavior in the dynamic game does not converge to *correlated* strategy profiles.

The second main role arises when we consider sequences of stage games where the number of players goes to infinity. Due to the perturbations, each player votes for either alternative with probability that is bounded away from 0 and 1. Hence, as the number of players increases, we can apply the central limit theorem to show that the probability of being pivotal vanishes. This result then allows us to conclude that, under an appropriate ordering of information, equilibrium strategies are increasing in signals when there is a large number of players, which then allows us to characterize equilibrium in large games.

The payoff perturbations also play a combination of the standard Harsanyi (1973) and Selten (1975) roles. After taking the limit with respect to the number of players to infinity, we then take the limit as the perturbations go to zero. We view this exercise as providing an equilibrium prediction in games without payoff perturbations. Moreover, in finite player games, the perturbations provide a refinement on equilibrium beliefs about the desirability of each alternative.

To understand the need for a refinement in the spirit of Selten (1975), consider a setting without payoff perturbations. Then existence of equilibrium (both Nash and our solution concept) is trivial for two reasons. First, there are strategy profiles for which no profitable deviation exists simply because no unilateral deviation may affect the chosen outcome. For example, it is an equilibrium for everyone to vote for the same alternative (no matter how bad this alternative may be) whenever $k < n$. Second, in our context beliefs about the desirability of alternatives will be determined endogenously in equilibrium. But if an alternative is never chosen, then equilibrium beliefs may be arbitrary, hence justifying the decision not to choose the alternative in the first place. Our perturbations imply that each alternative is chosen with strictly positive probability, thus providing the experimentation necessary to pin down beliefs in equilibrium.

The following result formalizes the perfection role discussed above. The proof is straightforward and therefore omitted.

Lemma 1. *Let $\alpha \in \mathcal{A}$. Then, for all i and $s_i \in S_i$:*

(i) $P(\alpha_i', \alpha_{-i})(o = A \mid s_i) - P(\alpha_i'', \alpha_{-i})(o = A \mid s_i) > 0$ *for all* $\alpha_i' > \alpha_i''$.
(ii) $P(\alpha)(o = A \mid s_i) \in (0, 1)$.

Part (i) says that player i's vote affects the outcome of the election; i.e. the probability of being pivotal is strictly greater than zero. Part (ii) says that alternatives A and B are chosen with strictly positive probability, so that in equilibrium beliefs about A and B will not be arbitrary. Hence, payoff perturbations play the role of trembles and rule out the type of trivial equilibria described above.

3.3 Definition of equilibrium

A naive (or, more generally, behavioral) 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. Players choose strategies that are optimal given their beliefs about the consequences of following each possible strategy. In contrast to a Nash equilibrium, these beliefs are not necessarily restricted to be correct, but rather to be consistent with the information feedback players receive. This information is in turn endogenously generated by the equilibrium strategies followed by all players. Our feedback assumption is that players only observe the realized payoff of the alternative that is chosen by the committee, but not the counter-factual payoff of the other alternative.[12] A naive equilibrium requires beliefs to be *naive-consistent*, meaning that this information is not correctly processed by players. In particular, players do not take into account that other players' actions may be correlated with the true state of nature. In this way, we formalize the idea that players do not take into account the informational content of other players' actions, or, equivalently, the sample selection problem.

To gain some intuition for the solution concept, suppose 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 only observed whenever A is chosen, player i will come to observe that, conditional on observing signal s_i, alternative A yields in expectation

$$E_i^A(P(\alpha), s_i) \equiv E_{P(\alpha)}\left[u_i(A, \omega) \mid o = A, s_i\right].$$

Similarly, alternative B yields in expectation

$$E_i^B(P(\alpha), s_i) \equiv E_{P(\alpha)}\left[u_i(B, \omega) \mid o = B, s_i\right].$$

A naive player who observes v_i and s_i believes that expected utility is maximized by voting for

[12]The assumption that counterfactuals are not observed guarantees that players' naive model of the world is consistent with their feedback (see Esponda (2008) for further discussion).

A whenever $\Delta_i(P(\alpha), s_i) - v_i > 0$ and voting for B otherwise,[13] where

$$\Delta_i(P, s_i) \equiv E_i^A(P, s_i) - E_i^B(P, s_i)$$

is well-defined by Lemma 1.

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

$$\alpha_i(s_i) = F_i\left(\Delta_i(P(\alpha), s_i)\right). \tag{2}$$

We refer to $P(\alpha) \in \Delta(Z)$ as a (naive) equilibrium distribution.

In equilibrium, players best respond to beliefs that are endogenously determined by both their own strategy and the strategies of other players and that are consistent with observed equilibrium outcomes. Naive players, however, do not account for the correlation between others' votes and the state of the world (conditional on their own private information). In particular, naive players fail to realize that a change in their own strategies may also affect their beliefs about alternatives A and B.

Theorem 1. *There exists a (naive) equilibrium of the stage game.*

Proof. Let $\Phi : \mathcal{A} \to \mathcal{A}$ be given by $\Phi(\alpha) = \left(F_i\left(\Delta_i(P(\alpha), s_i)\right)_{s_i \in S_i}\right)_{i=1,...,n}$. First, note that $\Phi(\alpha) \in \mathcal{A}$ for all $\alpha \in \mathcal{A}$. Second, $\mathcal{A}$ is a convex and compact subset of a Euclidean space. Third, $P(\cdot)$ is linear, hence continuous, implying that $\Delta_i(P(\cdot), s_i)$ (which is well-defined by Lemma 1) is continuous and, by continuity of F_i, that Φ is continuous. Therefore, by Brouwer's fixed point theorem there exists a fixed point of Φ, which is also an equilibrium of the stage game. □

4 Learning foundation for equilibrium

4.1 A model of learning

We present a dynamic framework in order to clarify and justify the solution concept. A dynamic game is a repetition of the stage game described in Section 3, where the state and the signals are drawn independently across time periods from the same distribution. The main result is that if behavior stabilizes in the dynamic game, then it stabilizes to an equilibrium of the stage game.

We adapt the learning model by Fudenberg and Kreps (1993) to our context. A group of n players play the stage game described in Section 3 for each discrete time period $t = 1, 2, ...$ At time

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

t, the state is denoted by $\omega_t \in \Omega$, the signals by $s_t = (s_{1t}, ..., s_{nt})$, and the votes by $x_t = (x_{1t}, ..., x_{nt})$. The outcome of the election at time t is determined by a threshold voting rule k and denoted by $o_t \in \{A, B\}$. Utility is given by

$$u_i(o_t, \omega_t) + 1\{o_t = B\} v_{it},$$

where v_{it} is the payoff perturbance drawn independently (across players and time) from F_i.

Let $h^t = (z_1, ..., z_{t-1})$ denote the history of the game up to time $t-1$, where $z_t = (x_t, s_t, \omega_t) \in Z$ is the time-t outcome. Let $\mathcal{H}^t$ denote the set of all time-t histories and let $\mathcal{H}$ be the set of infinite histories. At each round of play, players privately collect feedback about past outcomes. For each player i, $Z_i \equiv X \times S_i \times \mathrm{U}_i$ is the set of outcomes that player i may observe at any given period, where U_i is the range of her utility functions $u_i(A, \cdot)$ and $u_i(B, \cdot)$. Let $h_i^t = (z_{i1}, ..., z_{it-1})$ denote player i's private history up to time $t-1$, where $z_{it} = (x_t, s_{it}, u_i(o_t, \omega_t)) \in Z_i$ is the privately observed outcome at time t. Note that payoff perturbations are not part of the history, implicitly assuming that players understand that the perturbations are independent payoff shocks that are unrelated to the learning problem. Let $\mathcal{H}_i^t$ denote the set of all time-t private histories and let $\mathcal{H}_i$ be the set of infinite private histories for player i. By convention, if $t = 1$ then all these sets are singleton sets consisting of the null history.

We complete the specification of the dynamic game by introducing assessment (i.e. belief-updating) and policy rules. An *assessment rule* for player i is a sequence $\mu_i = (\mu_{i1}, ..., \mu_{it}, ...)$ such that $\mu_{it} : \mathcal{H} \to \mathbb{R}^{\#S_i}$ is measurable with respect to the player i's time-t private history. The interpretation is that the s_i-coordinate of $\mu_{it}(h)$, $\mu_{it}(h)(s_i)$, is player i's beliefs, given her private $t-1$-period history in h, about the *difference* in expected utility between alternatives A and B conditional on s_i.

A *policy rule* for player i is a history-dependent sequence of action plans $\phi_i^H = \left(\phi_{i1}^H, ..., \phi_{it}^H, ...\right)$, where $\phi_{it}^H : \mathcal{H} \times V_i \to Y_i \equiv \{A, B\}^{\#S_i}$ is measurable with respect to player i's time-t private history and their time-t payoff perturbation. The interpretation is that $\phi_{it}(h, v_i)(s_i)$ is player i's vote at time t conditional on observing private history h_i^t, perturbation v_i, and signal s_i.

The measurability restrictions on assessment and policy rules imply that players' decisions may depend on the observed payoff outcomes but not on the (unobserved) state of the world, thus capturing the assumption that players do not observe counter-factual payoffs.[14]

Given a *policy rule profile* $\phi^H = \left(\phi_1^H, ..., \phi_n^H\right)$, let $\mathbf{P}^{\phi^H}(\cdot \mid h^t)$ denote the probability distribution over histories conditional on history up to time t, $h^t \in H^t$ –which we can construct by Kolmogorov's extension theorem.

[14]Implicitly, we are also assuming that players do not know the functional form of the utility function nor the structure of the state space; otherwise, players would be able to make inferences about counter-factual payoffs by observing realized payoffs.

4.2 The selection problem: naive assessments

We place the following restrictions on assessment and policy rules.[15]

Definition 2. An assessment rule μ_i is *non-strategic and empirical* if

$$\mu_{it}(h)(s_i) = \frac{\sum_{\tau=1}^{t-1} 1\{o_\tau = A, s_{i\tau} = s_i\} u_{i\tau}}{\sum_{\tau=1}^{t-1} 1\{o_\tau = A, s_{i\tau} = s_i\}} - \frac{\sum_{\tau=1}^{t-1} 1\{o_\tau = B, s_{i\tau} = s_i\} u_{i\tau}}{\sum_{\tau=1}^{t-1} 1\{o_\tau = B, s_{i\tau} = s_i\}} \tag{3}$$

for every $h \in \mathcal{H}$, $s_i \in S_i$, and $t \geq 2$ such that the denominators are greater than zero, where o_τ, $s_{i\tau}$, and $u_{i\tau}$ are time-τ elements of h.

In words, the definition assumes that players believe that the difference in expected payoffs from A and B conditional on an observed signal is given by the *observed* empirical average difference in payoffs–the key here is that only the payoff to the chosen alternative is observed. Hence, players take the information they see at face value. In particular, they make no attempts to account for sample selection that arises because counter-factual payoffs are not observed.

The learning model is completed by assuming that players vote for the alternative that maximizes their current period's *perceived* expected utility.[16]

Definition 3. A policy rule ϕ_i^H is myopic relative to an assessment rule μ_i if for every $h \in \mathcal{H}$, $s_i \in S_i$, and $t \geq 1$,

$$\phi_{it}(h, v_{it})(s_i) = \begin{cases} A & \text{if } \mu_{it}(h)(s_i) - v_{it} \geq 0 \\ B & \text{otherwise} \end{cases}$$

The assumption of myopia is for simplicity. Forward-looking players would have incentives to experiment–even if naive about their assessments–simply to obtain additional information about an alternative they may have not observed enough times. The reason why naive players have incorrect beliefs, however, is not due to the lack of experimentation, but rather to their failure to account for the selection problem (Esponda (2008)). Indeed, even with myopic players there may exist perpetual "experimentation" in the sense that the proportion of time that the committee chooses either alternative is bounded away from zero in the limit as t grows to infinity.

4.3 Stability and equilibrium

Our objective is to relate distributions over outcomes of the dynamic game as $t \to \infty$ to equilibrium distributions over outcomes of the stage game. Define the sequence of random variables $\overline{P}_t : \mathcal{H} \to$

[15] Following Fudenberg and Kreps (1993), we could also define asymptotic versions of these rules.

[16] Implicitly, we are assuming that players believe (correctly) that their vote affects the outcome with a strictly positive probability. Also, the assumption that A is played if indifferent is only for simplicity and may be replaced with any tie-breaking rule.

$\Delta(Z)$, where

$$\overline{P}_t(h)(x,s,\omega) = \frac{1}{t-1}\sum_{\tau=1}^{t-1} 1\left\{x_\tau(h)=x, s_\tau(h)=s, \omega_\tau(h)=\omega\right\}$$

is the frequency distribution over outcomes in the dynamic game. We look at the frequency distribution in order to allow for the possibility that play in the dynamic game is correlated.

We focus attention on frequency distributions that eventually stabilize around a steady state distribution over outcomes. The following definition of stability accounts for the probabilistic nature and possible multiplicity of steady states.

Definition 4. $P^* \in \Delta(Z)$ is a *stable outcome distribution* of the dynamic game under policy rules ϕ^H if for all $\varepsilon > 0$ there exists t_ε such that

$$\mathbf{P}^{\phi^H}\left(\left\|\overline{P}_t - P^*\right\| < \varepsilon \text{ for all } t \geq t_\varepsilon\right) > 0.$$

The definition of stability captures the idea that after a *finite* number of periods there is a strictly positive probability that the frequency distribution over outcomes $\overline{P}_t$ remains forever close to P^*.[17]

Theorem 2. *Suppose that P^* is a stable outcome distribution of the dynamic game under policy rules ϕ^H that are myopic relative to assessment rules μ that are non-strategic and empirical. Then P^* is a (naive) equilibrium distribution of the stage game.*

Theorem 2 is our key justification for focusing on equilibria of the stage game: any profile that is not an equilibrium generates an outcome distribution that is not stable in the dynamic game. In particular, correlated strategy profiles do not generate stable outcome distributions in our environment. As our proof makes clear, this result follows from the assumption that payoff perturbations are independent across players and time.[18]

4.4 Proof of Theorem 2

Throughout the proof, we fix a stable outcome distribution P^* and policy rules ϕ^H that are myopic relative to assessment rules μ that are non-strategic and empirical. The proof compares "strategies"

[17]The statement is true because t_ε is not history-dependent for those histories in the history set that, according to the definition, has positive probability. But note that this definition is weaker than requiring that $\overline{P}_t$ converges to P^* with strictly positive probability.

[18]Of course, a converse to Theorem 2 does not necessarily hold; that is, there may exist equilibria that do not generate stable outcome distributions.

in the dynamic game with strategies in the stage game. To define the former, let the vector-valued random variable $\alpha_t^H = (\alpha_{1t}^H, ..., \alpha_{nt}^H) : \mathcal{H} \to \mathcal{A}_1 \times ... \times \mathcal{A}_n$ denote a time-t strategy profile, where

$$\alpha_{it}^H(h)(s_i) = \int 1 \left\{ \phi_{it}^H (h, v_i) (s_i) = A \right\} dF_i \tag{4}$$

is the probability that player i votes for A when observing signal s_i, conditional on history h^t.

Finally, let $\alpha^* = (\alpha_1^*, ..., \alpha_n^*) \in \mathcal{A}_1 \times ... \times \mathcal{A}_n$ be such that

$$\alpha_i^*(s_i) = F_i \left(\Delta_i(P^*, s_i) \right)$$

is the probability that player i votes A if she optimally responds to beliefs $\Delta_i(P^*, s_i)$.

The proof follows from two claims. In Claim 2.1, we show that stability of P^* implies that beliefs eventually remain close to $\Delta_i(P^*, s_i)$, thus implying that time-t strategies α_t^H eventually remain close to $\alpha^* \in \mathcal{A}_1 \times ... \times \mathcal{A}_n$. The key of the proof is that player's payoff perturbations are independently drawn from an atom-less distribution, implying that if beliefs settle down then strategies must also settle down, not just in an average sense, but actually in a per-period sense. In particular, Claim 2.1 implies that any correlation in players' strategies induced by a common history eventually vanishes. In Claim 2.2, we show that the fact that strategies remain close to α^* implies that $P^* = P(\alpha^*)$, where $P(\cdot)$ was previously defined in (1). Both claims rely on a straightforward generalization of a powerful technical result by Fudenberg and Kreps (1993, Lemma 6.2); this result allows us to apply the law of large numbers in a context where a sequence of random variables is not independently distributed but where the distributions conditional on past history are eventually very close to some common distribution.

Theorem 2 then follows immediately from Claim 2.2, since then for all i and s_i,

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

implying that α^* is an equilibrium of the stage game, and therefore that $P^* = P(\alpha^*)$ is an equilibrium distribution.

Claim 2.1 For all $\varepsilon > 0$ there exists H_ε with $\mathbf{P}^{\phi^H}(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$ there exists $t_{\varepsilon,h}$ such for all $t \geq t_{\varepsilon,h}$, all i, and all s_i, $\left| \alpha_{it}^H(h)(s_i) - \alpha_i^*(s_i) \right| < \varepsilon$ and $\left| \overline{P}_t(h)(z) - P^*(z) \right| < \varepsilon$ for all $z \in Z$.

Proof. By continuity of F_i it suffices to show that for all $\varepsilon > 0$ there exist $\gamma(\varepsilon) > 0$ with $\lim_{\varepsilon \to 0} \gamma(\varepsilon) = 0$ and H_ε with $\mathbf{P}^{\phi^H}(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$ there exists $t_{\varepsilon,h}$ such for all $t \geq t_{\varepsilon,h}$, all i, and all s_i,

$$F_i \left(\Delta_i(P^*, s_i) - \gamma(\varepsilon) \right) \leq \alpha_{it}^H(h)(s_i) \leq F_i \left(\Delta_i(P^*, s_i) + \gamma(\varepsilon) \right) \tag{5}$$

and

$$\left|\overline{P}_t(h)(z) - P^*(z)\right| < \varepsilon \tag{6}$$

or all $z \in Z$.

For each $o \in \{A, B\}$, let

$$Z_{ois_i\omega} \equiv \{(x', s', \omega') \in Z : \ o(x') = o, s_i' = s_i, \omega' = \omega\}$$

and $Z_{ois_i} \equiv \cup_{\omega\in\Omega} Z_{ois_i\omega}$. Let $\mu_{oit}(h)(s_i)$ denote the $o \in \{A, B\}$ term of $\mu_{it}(h)(s_i)$ in expression (3), which can be written as

$$\begin{aligned}\mu_{oit}(h)(s_i) &= \frac{\sum_{\tau=0}^{t-1} 1\{o_\tau = o, s_{i,\tau} = s_i\} u_i(A, \omega_t)}{\sum_{\tau=0}^{t-1} 1\{o_\tau = o, s_{i,\tau} = s_i\}} \\ &= \frac{\sum_{\omega\in\Omega} \overline{P}_t(h)(Z_{ois_i\omega}) u_i(o,\omega)}{\overline{P}_t(h)(Z_{ois_i})},\end{aligned}$$

provided that $\overline{P}_t(h)(Z_{ois_i}) > 0$.

Because P^* is stable, for all $\varepsilon > 0$ there exists t_ε and H_ε^* with $\mathbf{P}^{\phi^H}(H_\varepsilon^*) > 0$ such that for all $h \in H_\varepsilon^*$ and $t \geq t_\varepsilon^*$, equation (6) holds for all $(x, s, \omega) \in Z$. In addition, (6) implies that

$$\left|\overline{P}_t(h)(Z') - P^*(Z')\right| < \varepsilon \times \#Z \tag{7}$$

for all $Z' \subset Z$. Next, let $\eta = \min\left\{\mathbf{P}^{\phi^H}(H_\varepsilon^*), .5K_p\right\} > 0$, where K_p is defined by equation (36) in the appendix; by Lemma 6 in the appendix, for all $h \in H \setminus H^o$ (where H^o has zero measure) there exists $t_{\eta,h}$ such that for all $t \geq t_{\eta,h}$, all $o \in \{A, B\}$,

$$\overline{P}_t(h)(Z_{ois_i}) > K_p - \eta \geq .5K_p. \tag{8}$$

Let $H_\varepsilon = H_\varepsilon^* \cap H \setminus H^o$, and note that by our choice of H_ε^*, $\mathbf{P}^{\phi^H}(H_\varepsilon) > 0$. Therefore, for all $\varepsilon > 0$ there exists H_ε with $\mathbf{P}^{\phi^H}(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$ and $t \geq t_{\varepsilon,h} \equiv max\{t_\varepsilon^*, t_{\eta,h}\}$, all i, s_i

$$|\mu_{it}(h)(s_i) - \Delta_i(P^*, s_i)| \leq \gamma(\varepsilon) \equiv \frac{(\varepsilon \times \#Z) \times (0.5K_p)^2}{2K(1 + \#\Omega) + 0.5(\varepsilon \times \#Z)K_p} \underset{\varepsilon\to 0}{\longrightarrow} 0 \tag{9}$$

where the inequality follows from (7), (8), the facts that $|u_i(o, \omega)| < K$ and $\#\Omega < \infty$, and simple algebra that uses the fact that $E_i^o(P^*, s_i) = \frac{\sum_{\omega\in\Omega} P^*(Z_{ois_i\omega}) u_i(o,\omega)}{P^*(Z_{ois_i})}$. Then (9) and the definition of the policy function imply that

$$\phi_{it}^H(h, v_i)(s_i) = \begin{cases} A & \text{if } v_i \leq \Delta E_i(P^*, s_i) - \gamma(\varepsilon) \\ B & \text{if } v_i > \Delta E_i(P^*, s_i) + \gamma(\varepsilon) \end{cases},$$

so that (5) holds by (4). □

Claim 2.2 $P^* = P(\alpha^*)$

Proof. Note that for each $z \in Z$,

$$\mathbf{P}^{\phi^H}(z_t = z \mid h^t) = P(\alpha_t^H(h))(z).$$

Then Claim 2.1 and the fact that P is continuous in α imply that for all $\varepsilon > 0$ there exists H_ε with $\mathbf{P}^{\phi^H}(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$ there exists $\hat{t}_{\varepsilon,h}$ such that for all $t \geq \hat{t}_{\varepsilon,h}$,

$$\left|\mathbf{P}^{\phi^H}(z_t = z \mid h^t) - P(\alpha^*)(z)\right| < \varepsilon$$

and

$$\left|\overline{P}_t(h)(z) - P^*(z)\right| < \varepsilon \tag{10}$$

for all $z \in Z$.

Then by Lemma 5 in the appendix applied to all the singleton sets of Z,

$$\limsup_{t\to\infty} \overline{P}_t(z) \leq P(\alpha^*)(z) + \varepsilon \quad \text{and} \quad \liminf_{t\to\infty} \overline{P}_t(z) \geq P(\alpha^*)(z) - \varepsilon \tag{11}$$

for all $z \in Z$, almost surely on H_ε.

By the triangle inequality, for any t,

$$||P^* - P(\alpha^*)|| \leq ||P^* - \overline{P}_t(h)|| + ||\overline{P}_t(h) - P(\alpha^*)|| \tag{12}$$

for any $h \in H_\varepsilon$; we pick one $h \in H_\varepsilon$ (outside the measure zero set). By equation (11) the second summand in the RHS is less than ε for all t sufficiently large; by equation (10), the first summand of the RHS is also less than ε for all t sufficiently large. Hence $||P^* - P(\alpha^*)|| \leq \varepsilon$; since this holds for all $\varepsilon > 0$ then we obtain the desired result by taking $\varepsilon \to 0$. □

4.5 Alternative behavioral rules

A player who is sophisticated and understands the selection problem may account for it by conditioning her learning on past periods where her vote was pivotal. *If* her vote is random and independent of the state of the world (conditional on her private information), then the subsample where she is pivotal will not be biased. The reason is that whether she observes the payoff of an alternative or not in those periods where she is pivotal only depends on whether she votes for the alternative or not, which is independent of the state of the world.

There are a few reasons, however, why even a sophisticated player who understands selection may not behave as above. First, the procedure is not applicable where only the outcome of the election is observed and not the number of votes for each alternative. Second, even if the pivotal event were observable, the previous argument relies on a player's vote being independent of other players' votes, conditional on her private information. While this condition may hold in steady state, voting strategies are likely be correlated as players are learning to play the game. Hence, sophisticated players will initially face an endogenous selection bias and the pivotal learning rule will not be optimal along the path of play.[19] Alternatively, the strategy of randomizing their votes over a few periods of time will not necessarily lead to unbiased learning if other players are also concurrently learning to play the game. An exception is where players *coordinate* their randomization, as in the ubiquitous practice of conducting randomized trials to overcome the selection problem. Third, the inferences that a player makes conditional on being pivotal are not suitable for replication, either when facing a similar decision problem with a different group of people (who may have learned to make decisions in a different way) or when facing a decision on her own.

A different kind of behavioral rule may also lead to Nash equilibrium behavior. Suppose that players are pretty unsophisticated and don't have a good understanding of the problem at hand. All they know is that they must choose either A or B every period, and they keep track of their utility (but not of the payoffs of alternatives A and B). Suppose that players then vote for the alternative yielding the highest expected utility. Under certain conditions, this kind of reinforcement learning may lead to (locally) first-best outcomes. For our motivating examples, however, it seems dubious that voters would not seek to learn the expected payoffs of the alternatives. This is particularly true whenever players interact in other settings, either by themselves or as members of others committees, where this knowledge would be useful.[20]

5 Voting framework: a large number of players

We now present the framework for analyzing games where the number of players goes to infinity and introduce the notion of (perfect) limit equilibrium. A limit equilibrium is the limit, as the number of players goes to infinity, of a sequence of naive equilibria of a game with perturbed payoffs. A perfect limit equilibrium is the limit of a sequence of limit equilibria as the perturbation vanishes.

Consider the stage game in Section 3, where we now represent heterogeneity in preferences and information by indexing the utility function u_θ, the precision of the signals q_θ, the set of signals S_θ, the strategy set $\mathcal{A}_\theta$, and the distribution over the perturbation F_θ by a type θ. We assume that θ is randomly drawn from a finite set Θ according to a probability distribution $\phi \in \Delta(\Theta)$. The threshold voting rule is now $k = \rho n$, where $0 < \rho < 1$.

The following assumptions on the primitives are maintained through the remainder of the paper.

C1. For each $\theta \in \Theta$: (i) $u_\theta(A, \cdot)$ is nondecreasing and $u_\theta(B, \cdot)$ is nonincreasing, and one of them

[19] Of course, this bias disappears as the number of players increases and a player becomes negligible; however, the proportion of the sample that can be used to make inference also goes to zero.

[20] One point of this example is that the relationship between sophistication and optimality is not monotone. Our players have a good understanding of the problem they face and are able to learn from the data at hand, but optimality is precluded because they fail to account for selection or (more generally) because they do not perfectly account for it.

holds strictly; (ii) $\sup_{o=\{A,B\},\omega\in\Omega} |u_\theta(o,\omega)| < K < \infty$.

C2. There exists $z > 0$ such that for all i, $\omega' > \omega$, and $s_i' > s_i$:

$$\frac{q_i(s_i'|\omega')}{q_i(s_i'|\omega)} - \frac{q_i(s_i|\omega')}{q_i(s_i|\omega)} = z(\omega' - \omega).$$

C3. $\Omega = [-1, 1]$ and G is an absolutely continuous probability distribution over Ω with density g.

C4. (i)$\inf_\Omega g(\omega) > 0$; (ii) there exists $d > 0$ such that $q_\theta(s_\theta|\omega) > d$ for all θ, s_θ, ω; (iii) for all θ, η, F_θ^η is absolutely continuous and $\inf_{x\in[-2K,2K]} f_\theta^\eta(x) > 0$.

C5. For all θ, $u_\theta(A, \cdot)$ and $u_\theta(B, \cdot)$ are both continuously differentiable.

Assumptions C1 and C2 provide a standard ordering between states, information, and players' preferences, as well as a uniform bound on the utility function.[21] In particular, C2 requires that the strict MLRP (monotone likelihood ration property) holds. For technical reasons we actually need a slight strengthening of strict MLRP: there must be a uniform bound on the rate at which the likelihood ratio changes.

Assumption C3 departs from our previous assumption of a finite state space but its only purpose is to simplify the statement of our characterization result; $\Omega = [-1, 1]$ is chosen for simplicity and can be replaced by any compact, real-valued interval.[22] Assumption C4 requires densities to be uniformly bounded; in particular, "strong signals" (Milgrom 1979) are ruled out. Assumption C5 is for convenience but can be relaxed.

5.1 Limit equilibrium

We now introduce 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 full-support 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.

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

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

[21] It is important for our results that $u_\theta(A, \cdot)$ and $u_\theta(B, \cdot)$ are separately increasing; it does not suffice that their difference is increasing.

[22] Equilibrium can be characterized in terms of a cutoff state above which A is chosen and below which B is chosen with probability going to 1. Assumption C3 allows us to ignore what happens at the cutoff. With a finite state space, we would have the same characterization but we would also have to indicate the probability with which A is chosen at the cutoff state.

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, where $X^{(n)} \equiv \prod_{i=1}^n X_i$ and $S^{(n)} \equiv \prod_{i=1}^n S_i$.

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. This condition 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 5. 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 $s_i \in S_i$,

$$|\alpha_i^n(\xi)(s_i) - F_i\left(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i)\right)| \leq \varepsilon \tag{13}$$

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 restrict attention to studying equilibrium strategies that satisfy this property since we know that the cases where one of the alternatives is chosen with probability 1 always constitute a trivial equilibrium.[23]

Definition 6. A strategy mapping $\boldsymbol{\alpha}$ is $\Xi'-$*asymptotically interior* if

$$\liminf_{n\to\infty} P^n(\boldsymbol{\alpha}(\xi))\,(o = A) > 0 \tag{14}$$

and

$$\limsup_{n\to\infty} P^n(\boldsymbol{\alpha}(\xi))\,(o = A) < 1 \tag{15}$$

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 equilibria can be characterized by this convenient property.

[23]Note that a fixed payoff perturbation does not preclude an alternative from being chosen with probability that goes to 1 as the number of players increases.

Definition 7. 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 follow the same strategy.[24]

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 profile played by players of each type θ in the n-player game with types $(\theta_1,...,\theta_n)$ and strategy profile $\alpha^n(\xi)$. Formally,

$$\sigma^n_\theta(\xi;\boldsymbol{\alpha})(s_i) = \frac{\sum_{i=1}^n 1\left\{\theta_i(\xi) = \theta\right\}\alpha_i^n(\xi)(s_i)}{\sum_{i=1}^n 1\left\{\theta_i(\xi) = \theta\right\}} \in \mathcal{A}_\theta \tag{16}$$

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

Definition 8. 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 9. An average strategy profile $\sigma \in \mathcal{A}^*$ is a *limit equilibrium* if it is a limit ε-equilibrium for all $\varepsilon > 0$.

5.2 Perfection: vanishing perturbations

Up to this point we have defined equilibrium in games where players' payoffs are independently perturbed every period. While in some contexts the perturbations may have a real intepretation,

[24]It will be a result that players of the same type play the same equilibrium strategy, except for at most one signal where players of the same type may mix with different probabilities.

we now consider sequences of equilibria where the perturbation goes to zero. We interpret the limit of equilibria as the perturbation goes to zero as our solution concept for large games without payoff perturbations. To do so, we now index games (and strategies) by a parameter η that affects the distribution F^η from which perturbations are drawn. We assume that F^η satisfies the same assumptions as in Section 3 for all η, and in addition require that perturbations vanish as η goes to zero.

Definition 10. A family of perturbations $\{F^\eta\}_\eta$, where $F^\eta = \{F^\eta_\theta\}_{\theta\in\Theta}$, is *feasible* if it satisfies the following assumptions for all $\theta \in \Theta$ and η: F^η_θ is absolutely continuous, $F^\eta_\theta(-2K) > 0$ and $F^\eta_\theta(2K) < 1$, and

$$\lim_{\eta\to 0} F^\eta_\theta(v) = \begin{cases} 0 & \text{if } v < 0 \\ 1 & \text{if } v > 0 \end{cases} \tag{17}$$

Definition 11. An average strategy profile $\sigma \in \mathcal{A}^*$ is a *perfect limit equilibrium* if there exists a feasible family of perturbations $\{F^\eta\}$ and a sequence $\{\sigma^\eta\}$ of limit equilibria such that $\lim_{\eta\to 0}\sigma^\eta = \sigma$.

6 Characterization of equilibrium in large games

The intuition behind the characterization results can be grasped by thinking about a voting game with a continuum of players.[25] For a given average strategy profile $\sigma \in \mathcal{A}^*$, we interpret

$$\kappa_\theta(\sigma \mid \omega) = \sum_{s_\theta\in S_\theta} q_\theta(s_\theta \mid \omega)\,\sigma_\theta(s_\theta)$$

as the proportion of players of type θ that vote for A conditional on the state being ω, and

$$\kappa(\sigma \mid \omega) = \sum_{\theta\in\Theta} \kappa_\theta(\sigma \mid \omega)\,\phi(\theta)$$

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

$$\Omega(\sigma) \equiv \{\omega \in \Omega : \kappa(\sigma \mid \omega) > \rho\} \tag{18}$$

is the set of states where alternative A is chosen by the committee. We say that an average strategy profile σ has the c-cutoff property if there exists $c \in (-1, 1)$ such that

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

[25]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.

Note that an average profile σ with the c-cutoff property satisfies $0 < p(\Omega(\sigma)) < 1$ because we require the cutoff to be interior, i.e. $c \in (-1, 1)$; hence for such profiles we define

$$v_\theta (s_\theta; \sigma) \equiv E (u_\theta(A, \omega) \mid \omega \in \Omega(\sigma), s_\theta) - E (u_\theta(B, \omega) \mid \omega \notin \Omega(\sigma), s_\theta)$$

to be 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 the state is in $\Omega(\sigma)$ and the payoff of B whenever the state is not in $\Omega(\sigma)$. We can interpret $v_\theta (s; \sigma)$ as type θ's belief about the difference in expected payoff from A and B when players follow average strategy σ.

6.1 Limit equilibrium

In this section we characterize equilibrium for a fixed perturbation structure as the number of players goes to infinity. We first characterize limit ε-equilibrium and then use this result to provide a fixed-point characterization of limit equilibrium.

Lemma 2. *There exists $\bar{\varepsilon} > 0$ and $\gamma(\varepsilon)$ with $\lim_{\varepsilon \to 0} \gamma(\varepsilon)$ such that for all $\varepsilon < \bar{\varepsilon}$: if σ is a limit ε-equilibrium, then*

1. *σ is a c-cutoff limit ε-equilibrium*
2. *σ is increasing and has the c-cutoff property*
3. For all $\theta \in \Theta$ and $s_\theta \in S_\theta$,

$$|\sigma_\theta(s_\theta) - F_\theta (v_\theta (s_\theta; \sigma))| \leq \gamma(\varepsilon) \tag{19}$$

We now provide a discussion and proof of Lemma 2, relegating some of the details to the appendix. The proof relies on the following Lemma.

Lemma 2.1. *Suppose that there exists $\boldsymbol{\alpha}$ and Ξ' with $\Phi(\Xi') > 0$ such that $\boldsymbol{\alpha}$ is Ξ'−asymptotically interior and for all $\xi \in \Xi'$*

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

where σ is increasing. Then σ has the c-cutoff property, $\boldsymbol{\alpha}$ is Ξ'-asymptotically c-cutoff, and for all i, s_i,

$$\lim_{n \to \infty} \Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i) = v_{\theta_i} (s_i; \sigma) \tag{20}$$

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 in $\Omega(\sigma)$, where $\kappa(\sigma \mid \omega) > \rho$, the probability that the outcome is A converges to 1. Similarly, for states such that $\kappa(\sigma \mid \omega) < \rho$, 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 is being able to apply Lemma 2.1 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 a Nash equilibrium, since under assumptions C1 and C2 players' strategies would always be increasing (recall that payoffs are perturbed in a way that no one votes for an alternative with probability 1). What is different in our setting is that players' beliefs about which alternative is best does not only depend on a player's signal and the strategies of other players, but also depends on a player's own strategy. To understand the main issue, fix a player and a signal and suppose 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 players' 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 2.2. There exists $\bar{\varepsilon}$ such that for all $\varepsilon < \bar{\varepsilon}$: 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. Given this result, the effect of a player's own strategy on her own learning must 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 vote 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.[26]

Proof of Lemma 2: Let $\varepsilon \leq \bar{\varepsilon}$, where $\bar{\varepsilon}$ is defined by Lemma 2.2. Suppose that σ is a limit ε-equilibrium with corresponding ε-equilibrium mapping $\boldsymbol{\alpha}$ and convergence in a set Ξ'. By Lemma 2.2, σ is increasing. Therefore all the hypothesis of Lemma 2.1 are satisfied, implying that σ has the c-cutoff property and σ is a c-cutoff limit ε-equilibrium. In addition, Lemma 2.1 implies that $\lim_{n\to\infty} \Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i) = v_{\theta_i}(s_i;\sigma)$ a.s.-Ξ' and, by continuity of F_{θ_i} (assumption C5(i)), that $\lim_{n\to\infty} F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i)) = F_{\theta_i}(v_{\theta_i}(s_i;\sigma))$ a.s.-Ξ'. Therefore, there exists n_ε such that for all $n \geq n_\varepsilon$, all i, s_i

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

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 (16) imply that for all $n \geq n_\varepsilon$, all θ, s_θ,

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

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,

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

Lemma 2 then follows by letting $\gamma(\varepsilon) = 3\varepsilon$. □

To conclude this section, we use Lemma 2 to show that the set of limit equilibria has a convenient and intuitive fixed-point characterization.

Theorem 3. *The following statements are equivalent:*

1. *σ is a limit equilibrium*

[26] 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.

2. σ is increasing, has the c-cutoff property, and for every $\theta \in \Theta$ and $s_\theta \in S_\theta$,

$$\sigma_\theta(s_\theta) = F_\theta\left(v_\theta\left(s_\theta;\sigma\right)\right) \tag{21}$$

Proof. $1 \Rightarrow 2$: Since σ is a limit ε-equilibrium for all $\varepsilon > 0$, Lemma 2 implies that (a) σ is increasing, (b) σ has the cutoff property and (c) for all $\bar{\varepsilon} \geq \varepsilon > 0$, for all $\theta \in \Theta$ and $s_\theta \in S_\theta$, $|\sigma_\theta(s_\theta) - F_\theta\left(v_\theta\left(s_\theta;\sigma\right)\right)| \leq \gamma(\varepsilon)$. Since the LHS of the inequality in part (c) does not depend on ε and $\gamma(\varepsilon) \to 0$, then $|\sigma_\theta(s_\theta) - F_\theta\left(v_\theta\left(s_\theta;\sigma\right)\right)| = 0$.

$2 \Rightarrow 1$: Consider the strategy mapping $\boldsymbol{\alpha}$ defined by letting players of type θ always play σ_θ, i.e., $\alpha_i(\xi)(s_i) = \sigma_\theta(s_\theta)$ for all ξ, all $i : \theta_i = \theta$. First, note that $\sigma^n = \sigma$ converges trivially to σ and σ is increasing by assumption. Second, note that, since σ is increasing, we can follow the proof leading to equation (37) 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 2.1 to obtain $\lim_{n\to\infty} \Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i) = v_{\theta_i}(s_i;\sigma)$ a.s.-Ξ. By continuity of F_{θ_i} (assumption C5(i)), $\lim_{n\to\infty} F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\alpha}(\xi)), s_i)) = F_{\theta_i}(v_{\theta_i}(s_i;\sigma))$. Therefore, for $\varepsilon > 0$, there exists a n_ε such that for $n \geq n_\varepsilon$, all i, s_i

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

a.s.-Ξ. □

6.2 Perfect limit equilibrium

We now characterize equilibrium for vanishing perturbations. By Theorem 3, all limit equilibria have the c-cutoff property, meaning that for states below the cutoff the probability of choosing A goes to 1 and for states below the cutoff it goes to 0. We first characterize equilibrium cutoffs, which suffices for most practical questions–such as those studied in Section 7. For a sequence of limit equilibria with the property that the c-cutoff property is preserved in the limit, we also provide a fixed point characterization of perfect limit equilibrium strategy profiles.

Lemma 3. *Suppose that σ is a perfect limit equilibrium. Then there exists a sequence of limit equilibria $\{\sigma^\eta\}$ that converges to σ and such that for all $\theta \in \Theta$ and $s_\theta \in S_\theta$,*

$$\sigma_\theta(s_\theta) \in \arg\max_{x\in[0,1]}\ x \cdot \left(\lim_{\eta\to 0} v_\theta\left(s_\theta;\sigma^\eta\right)\right).$$

Proof. Let $\{\sigma^\eta\}$ be the sequence of limit equilibria supporting σ; we can always take a subsequence (also denoted by $\{\sigma^\eta\}$) such that $\lim_{\eta\to 0} v_\theta\left(s_\theta;\sigma^\eta\right)$ exists. Suppose, in order to obtain a contradiction, that there exists θ, s_θ such that

$$\sigma_\theta(s_\theta) \notin \arg\max_{x\in[0,1]} \; x \cdot \left(\lim_{\eta\to 0} v_\theta\left(s_\theta;\sigma^\eta\right)\right).$$

Then either (i) $\lim_{\eta\to 0} v_\theta\left(s_\theta;\sigma^\eta\right) > 0$ and $\sigma_\theta(s_\theta) < 1$ or (ii) $\lim_{\eta\to 0} v_\theta\left(s_\theta;\sigma^\eta\right) < 0$ and $\sigma_\theta(s_\theta) > 0$. Suppose (i) –the case (ii) is similar. Then there exists $\delta > 0$ and η_δ such that for all $\eta \leq \eta_\delta$,

$$v_\theta\left(s_\theta;\sigma^\eta\right) \geq \delta > 0. \tag{22}$$

In addition, since $\{F^\eta\}$ is feasible and $\delta > 0$, there exists η_σ such that for all $\eta \leq \eta_\sigma$,

$$F^\eta\left(\delta\right) > \frac{1}{2}\left(1+\sigma_\theta(s_\theta)\right) \tag{23}$$

Therefore, for all $\eta \leq \min\{\eta_\varepsilon, \eta_\delta, \eta_{\varepsilon,\delta}\}$,

$$\begin{aligned}\sigma_\theta^\eta(s_\theta) &= F^\eta\left(v_\theta\left(s_\theta;\sigma^\eta\right)\right)\\ &\geq F_\eta\left(\delta\right)\\ &> \frac{1}{2}\left(1+\sigma_\theta(s_\theta)\right),\end{aligned}$$

where the first line follows from Theorem 3, the second line from (22), and the third line from (23); hence contradicting that $\sigma_\theta^\eta(s_\theta)$ converges to $\sigma_\theta(s_\theta) < 1$. □

Lemma 3 and strict MLRP imply that, for each type, the equilibrium strategy is strictly mixed for at most one signal. However, we cannot rule out that in equilibrium every type chooses a strategy that does not depend on its private information.

Definition 12. An average strategy profile σ is *responsive* if there exists a type θ and signals $s'_\theta, s_\theta \in S_\theta$ such that $\sigma_\theta(s'_\theta) \neq \sigma_\theta(s_\theta)$.

A general fixed-point characterization is precluded by the fact that, for equilibria that are not responsive,

$$\lim_{\eta\to 0} v_\theta\left(s_\theta;\sigma^\eta\right) \neq v_\theta\left(s_\theta;\sigma\right).$$

Nevertheless, we will be able to provide a fixed-point characterization for equilibria that are responsive. For equilibria that are not responsive, we can still characterize the set of equilibrium outcomes.

We next characterize the set of equilibrium *outcomes* for *all* perfect equilibria. We know from Theorem 3 that every limit equilibrium is characterized by a cutoff: as the number of players grows, the probability that the outcome is A goes to 0 for states below the cutoff and goes to 1 for states above the cutoff.

Definition 13. A *perfect equilibrium cutoff* c is the limit of a sequence $\{c^\eta\}$ of cutoffs corresponding to the sequence of limit equilibria $\{\sigma^\eta\}$ for some family of feasible perturbations $\{F^\eta\}$.

Our next objective is to characterize the set of perfect equilibrium cutoffs. To provide some intuition, let $\{c^\eta\}$ be a sequence of cutoffs corresponding to a sequence $\{\sigma^\eta\}$ of limit equilibria and suppose that $c^\eta \to c$; we can always find a convergent subsequence. Let

$$\overline{v}_\theta\left(s_\theta; c\right) \equiv E\left(u_\theta(A,\omega) \mid \omega > c, s_\theta\right) - E\left(u_\theta(B,\omega) \mid \omega < c, s_\theta\right)$$

denotes a players' perceived expected utility conditional on the outcome being driven by a cutoff c.[27] Since σ^η has the c^η-cutoff property, then $v_\theta(s_\theta;\sigma^\eta) = \overline{v}(s_\theta; c^\eta)$ for all η. By continuity of $\overline{v}_\theta(s_\theta;\cdot)$,

$$\lim_{\eta\to 0} v_\theta(s_\theta;\sigma^\eta) = \overline{v}_\theta(s_\theta; c). \tag{24}$$

Moreover, $\overline{v}_\theta(s_\theta;\cdot)$ is increasing by strict MLRP. Therefore, since Ω is compact, for each θ, s_θ there exists a unique

$$c_\theta(s_\theta) = \arg\min_{c\in\Omega} \left|\overline{v}_\theta\left(s_\theta; c\right)\right| \tag{25}$$

that is decreasing in s_θ. In particular, letting s_θ^L be the lowest and s_θ^H the highest signal of type θ, then $c_\theta(s_\theta^H) \le c_\theta(s_\theta^L)$. Let Ω_θ denote the finite set of cutoffs $c_\theta(s_\theta)$, for all $s_\theta \in \Theta$.

For each $c \in \Omega$, let

$$S_\theta(c) \equiv \{s : c_\theta(s) < c\}$$

and

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

When $c \notin \Omega_\theta$, $S_\theta(c)$ denotes the set of signals under which type θ will vote for A and $\overline{\kappa}(c)$ may be interpreted as the proportion of players that vote for A conditional on state c, as the perturbation vanishes.

Lemma 4. $\overline{\kappa} : \Omega \to [0,1]$ is weakly increasing and satisfies

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

[27] Define $\overline{v}_\theta\left(\underline{\omega}, s_\theta\right) \equiv \lim_{c\to\underline{\omega}} \overline{v}_\theta\left(c, s_\theta\right)$ and $\overline{v}_\theta\left(\overline{\omega}, s_\theta\right) \equiv \lim_{c\to\overline{\omega}} \overline{v}_\theta\left(c, s_\theta\right)$.

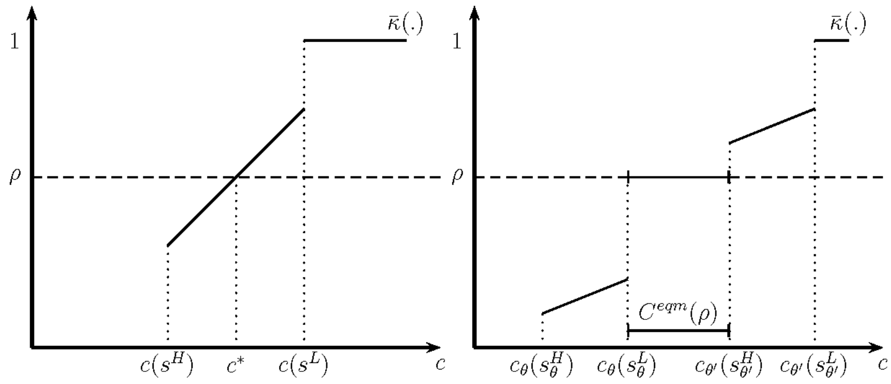

Figure 4: Characterization of perfect equilibrium cutoffs.

Proof. Define

$$\hat{\kappa}\left(c \mid \omega\right) \equiv \sum_{\theta \in \Theta} \phi_\theta q_\theta \left(S_\theta(c) \mid \omega\right). \tag{26}$$

First, note that $q_\theta\left(S_\theta(c) \mid \omega\right)$ 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 result then follows because $\overline{\kappa}(c) = \hat{\kappa}(c \mid c)$. ☐

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 equilibrium cutoff is given by the state c^* where $\overline{\kappa}$ intersects the voting rule ρ. Roughly, the idea is that if the cutoff were at some $c > c^*$, then we would be able to find c' between c^* and c such that the proportion that vote for A remains strictly above ρ. But this violates the assumption of convergence to a cutoff c. In panel (b), there is an entire segment where $\overline{\kappa}$ intersects ρ. We can rule out any state outside this segment being an equilibrium cutoff.

Panels (a) and (b) also distinguish between responsive and non-responsive equilibria. In the game in panel (a), perfect equilibrium strategies must be responsive: the sequence of cutoffs $\{c^\eta\}$ corresponding to the sequence $\{\sigma^\eta\}$ of limit equilibria converges to the cutoff of the perfect equilibrium strategy $\sigma = \lim_{\eta \to 0} \sigma^\eta$. In particular, responsive equilibria are characterized by a unique equilibrium cutoff. In panel (b), perfect equilibrium strategies are non-responsive: while $\kappa(\sigma^\eta \mid \cdot)$ is increasing, in the limit the function becomes flat. Hence, the previous continuity property on equilibrium cutoffs does not hold and we must study the sequence of equilibrium cutoffs rather than rely on properties of the limit of equilibrium strategies.

The next result formalizes the above discussion, provides a converse of the result, and accounts for possible discontinuities in $\overline{\kappa}$.

Theorem 4. *For a game with voting rule ρ, the set of perfect equilibrium cutoffs is given by*[28]

$$C^{eqm}(\rho) = \left[\inf_{c} \left\{ \overline{\kappa}(c) \geq \rho \right\}, \sup_{c} \left\{ \overline{\kappa}(c) \leq \rho \right\} \right].$$

In addition, a perfect limit equilibrium σ is responsive if and only if $\inf_c \left\{ \overline{\kappa}(c) \geq \rho \right\} = \sup_c \left\{ \overline{\kappa}(c) \leq \rho \right\}$.

Proof. See the appendix. □

An immediate corollary of Theorem 4 is that, depending on the primitives of the game, either all perfect equilibria are responsive or all perfect equilibria are not responsive.[29]

When equilibrium strategies are responsive, we also obtain a fixed-point characterization of equilibrium.

Theorem 5. *Suppose that σ is responsive with an interior cutoff $c \in (0,1)$. Then σ is a perfect limit equilibrium if and only if for every $\theta \in \Theta$ and $s_\theta \in S_\theta$,*

$$\sigma_\theta(s_\theta) \in \arg \max_{x \in [0,1]} \; x \cdot v_\theta \left(s_\theta ; \sigma \right). \tag{27}$$

Proof. See the appendix. □

7 Information aggregation

In this section we consider a voting stage game with a large number of players where the perturbation goes to zero. We maintain assumptions C1-C5 and apply the characterization results in Section 6 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 regarding the conditions for information aggregation. Finally, we discuss how our results extend to the case where naive players coexist with Nash players.

Our analysis is carried out from the perspective of a social planner who wants to maximize players' welfare. Let $W(c)$ denote aggregate 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 $W(c') > W(c)$

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

[29] In particular, if there is only one type then $\overline{\kappa}$ is increasing (while its image is in $(0,1)$) and therefore all equilibria are responsive.

for $0 > c' > c$ and $W(c') < W(c)$ for $c' > c > 0$. Thus, $W(c)$ is single-peaked at $c = 0$ and strictly decreases as c either increases or decreases away from $c = 0$.[30]

Definition 14. 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).[31] In our context, information may or may not be aggregated depending on the primitives. The next result, which follows immediately from Theorem 4 and the characterization of $\overline{\kappa}$ in Lemma 4, provides necessary and sufficient conditions on the primitives such that there exists a voting rule that aggregates information.

Proposition 1. *Information is aggregated by naive voters if and only if*

$$\min_{\theta} c_\theta(s_\theta^H) \leq 0 \leq \max_{\theta} c_\theta(s_\theta^L). \tag{28}$$

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

[30] The welfare function $W(c)$ 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 is only for simplicity; the important assumption is that the optimal cutoff is interior.

[31] 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, and we therefore 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.

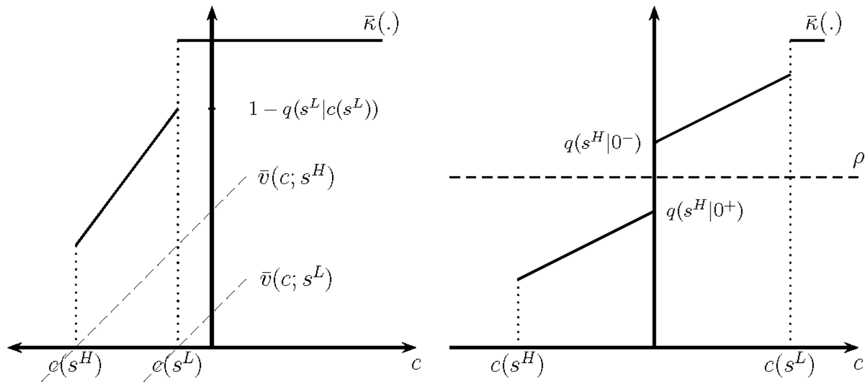

Figure 5: Information aggregation and optimal voting rules.

The next result, also an immediate implication of Theorem 4 and Lemma 4, 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 voting rule* ρ *is optimal if and only if*

$$\rho \leq \phi_{\theta_0} \cdot q_{\theta_0}\left(s_{\theta_0}^H \mid c_{\theta_0}(s_{\theta_0}^H)\right).$$

If information is aggregated, then voting rule ρ *is optimal if and only if*

$$\rho \in \left[\lim_{c \to 0^-} \overline{\kappa}(c), \lim_{c \to 0^+} \overline{\kappa}(c)\right].$$

To understand part of the intuition behind Proposition 2, suppose that $\max_\theta c_\theta(s_\theta^L) < 0$, so that information is not aggregated (see Figure 5(a)). As argued above, the reason why information is not aggregated is that, if it were, everyone would prefer to vote for A, irrespective of their signal. 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 carrry 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, be indifferent between A and B when she observes her lowest signal. Given such indifference, there are 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 $c_{\theta^0}(s_{\theta^0}^L)$ 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 15. Information is *sufficiently accurate* if there exist signals $s \neq s'$ such that

$$q\left(s \mid \omega\right) > 1/2 \quad \text{for } \omega > 0 \tag{29}$$

and

$$q\left(s' \mid \omega\right) > 1/2 \quad \text{for } \omega < 0. \tag{30}$$

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 that satisfy (29) and (30) 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 $c(s^H) \leq 0 \leq c(s^L)$. Let $c > 0 \geq c(s^H)$: then $\overline{\kappa}(c) \geq q(s^H \mid c) > 1/2$, where the inequality follows from (29). 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 (30). 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 (29), $q(s^H \mid c(s^H)) > 1/2$. Proposition 2 then implies that $\rho = 1/2$ is optimal. Similarly, if $c(s^L) < 0$, then, by (30), $1 - q(s^L \mid c(s^L)) < 1/2$. Proposition 2 then implies that $\rho = 1/2$ is optimal. □

7.1 Examples

The following examples illustrate Propositions 1-3 and provide additional insights about how the payoff and information structure relates to information aggregation. For simplicity, we only discuss examples where all players are symmetric (i.e. there is only one type); our results can also be applied to obtain additional insights when players are asymmetric.[32]

First, suppose that

$$\inf_{\omega>0} u(A,\omega) > \sup_{\omega<0} u(B,\omega), \tag{31}$$

so that alternative A dominates B when restricted to states of the world where each alternative is best. Then $\overline{v}(0,s) > 0$ for all s, implying that $c(s^L) < 0$ and therefore, by Proposition 1, that information cannot be aggregated: if it were, then no one would like to vote for B.[33]

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

$$q(s^H|\omega) = \begin{cases} (0.5 + r_1\omega)^{1/r_2} & \text{if } \omega > 0 \\ (0.5 + r_1\omega)^{r_2} & \text{if } \omega < 0 \end{cases}. \tag{32}$$

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 (31) does not hold.

We will vary the parameters $r_1 \in (0, 0.5)$, $r_2 \in [1,\infty)$ 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 and continuous. At one extreme, $r_1 \approx 0$ and the signal is almost uninformative about the state. Since

$$E\left(u(A,\omega) \mid \omega > 0\right) > E\left(u(B,\omega) \mid \omega < 0\right), \tag{33}$$

[32] For example, information aggregation *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.

[33] An example that satisfies (31) 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 $\omega > 0$.

then 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

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

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 25), 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. In the second case, signal s^L puts a larger weight on states such that B is more successful than A (conditional on making correct choices), and therefore makes players willing to vote for B under s^L. This example suggests that information aggregation obtains provided that correct payoffs are not too far from each other, that correct payoffs vary in intensity depending on the state, and that there are signals which detect this variation.

(ii) *Optimal voting rules.* Let's continue to suppose that $r_2 = 0$ and let's now fix $h = 0$. Consider first the case where $r_1 < r_1^*$, so that information is not aggregated.[34] 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, (32) 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 4 (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.* So far the magnitude of the type I error has not played an explicit role. One may conjecture that in cases where information is not aggregated, a large payoff penalty for errors translates into a higher 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

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

lowest signal to be willing 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. 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 (33) then implies that information cannot be aggregated.[35] 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.

7.2 Coexistence of naive 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 somehow perfectly account for it. Of course, as discussed in Section 4.5, the presence of Nash players may or may not be justified depending on the setting.

First, consider a case where information is not aggregated in the presence of naive players, as in Figure 6(a). If a fraction $\gamma \approx 0$ of players is Nash and the remaining fraction $1-\gamma$ is naive, 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 naive 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 naive players is robust to a

[35] 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.

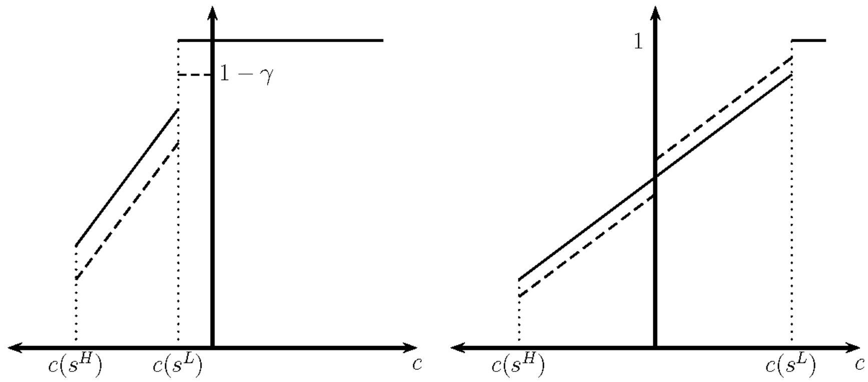

Figure 6: Coexistence of Nash and naive players.

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 case of a result by Fedderson and Pesendorfer (1997): we can interpret the naive players as a large (exogenous) fraction of partisans that 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 Fedderson and Pesendorfer (1997) when all players are Nash: when both Nash and naive players coexist, their full information equivalence result only holds for rules $\rho > (1-\gamma)$, rather than for all voting rules.[36] 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 naive 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.

8 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 provided a learning foundation for the notion of a behavioral equilibrium (Esponda, 2008) applied to voting games, and then used that notion to fully characterize *all* 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 in large populations. We also

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

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.

While we have focused on the benchmark voting context, we hope that our approach leads to further work in both the areas of learning and information aggregation. Our players are in a learning environment where their actions affect what they learn. In several economic contexts, players must make inferences about the primitives of the environment and the actions of other players; disentangling these two sources of uncertainty is likely to present challenges. In particular, players may need to have a model of how other players learn, how other players think that other players learn, and so on. While it is tempting to close the model by making an equilibrium assumption, the purpose of a dynamic learning model is to close the model without such an assumption. This paper constitutes a step in this direction, and this step is tractable because players learn using a mis-specified model that fails to account the informational content of other players' actions.

Finally, we were able to characterize all equilibria of games with sufficiently many players. Our approach relied on studying a perturbed version of the voting game and on defining the notion of an average strategy profile for each type of player. We believe that a similar approach can prove fruitful when studying information aggregation in other contexts and under other solution concepts.

9 Appendix

9.1 Dynamics

The following 2 lemmas are used in the proof of Theorem 2.

Lemma 5. *(cf. Fudenberg and Kreps, Lemma 6.2, 1993) Let $(z_t)_t$ be a sequence of random variables with range on a finite set Z. Fix a set-function $\pi : 2^Z \to [0,1]$ (not necessarily a probability measure) and fix $\varepsilon \in \mathbb{R}$. Let H_ε be a subset of infinite histories such that for all $h \in H_\varepsilon$ there exists $t_{\varepsilon,h}$ such that for all $t \geq t_{\varepsilon,h}$, the distribution of each z_t conditional on $h^t = (z_1, ..., z_{t-1})$, denoted $\pi_t(\cdot \mid h^t)$, satisfies*

$$\max_{Z' \in \mathcal{Z}} \pi_t(Z') - \pi(Z') > -\varepsilon, \tag{34}$$

where $\mathcal{Z} \subset 2^Z$ is a set of subsets of Z.[37]

Then

$$\liminf_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} 1\{z_\tau \in Z'\} \geq \pi(Z') - \varepsilon \tag{35}$$

[37]If H_ε has zero probability, the lemma is taken to be vacuous.

for all $Z' \in \mathcal{Z}$, almost surely on H_ε. Moreover, if (34) is replaced by $\max_{Z'\in\mathcal{Z}} \pi_t(Z') - \pi(Z') < \varepsilon$, *then the conclusion in (35) is replaced by* $\limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} 1\{z_\tau \in Z'\} \leq \pi(Z') + \varepsilon$.

Proof. First note that $\#Z < \infty$ and thus any subset of $\mathcal{Z} \subset 2^Z$ has also finitely many elements. Therefore, it suffices to show the result for any (arbitrary) subset $Z' \in \mathcal{Z}$ since there are only finitely many of them (roughly speak, Z' is what a is in FK 93). Since Z is finite we can order the elements as $(z_1, ...z_{\#Z})$, and WLOG we set the first $\#Z'$ to be the elements of Z'. Just as FK 93, let $(\omega_t)_t$ be an independent sequence of uniform random variables and let $y_t : \Omega \to Z$ be a new random variable.

As in FK 93, we construct $(y_t(\omega_t))_t$ as follows. For $t = 1$, $y_1(\omega_1) = z_m$ iff $\sum_{n=1}^{m-1} \pi_1(z_n) \leq \omega_1 < \sum_{n=1}^{m} \pi_1(z_n)$. For $t = \tau$, let $y_\tau(\omega_\tau) = z_m$ iff $\sum_{n=1}^{m-1} \pi_\tau(z_n|y_1, ..., y_{\tau-1}) \leq \omega_\tau < \sum_{n=1}^{m} \pi_\tau(z_n|y_1, ..., y_{\tau-1})$. Moreover, by construction the probability over h^t coincides with the probability over $(\omega_\tau)_{\tau\leq t}$; we thus can use both interchangeably. In particular, the set of ω for which $y_t(\omega_t) \in Z'$ is the set of $\{\omega : \omega_t \leq \sum_{n=1}^{\#Z'} \pi_t(z_n|y_1, ..., y_{t-1}) = \pi_t(Z'|y_1, ..., y_{t-1})\}$ (recall that Z' consists of the first $\#Z'$ elements in Z).

Under equation 34 the latter set includes the set $\{\omega : \omega_t \leq \pi(Z') - \varepsilon\}$; thus $1\{\omega : \omega_t \leq \pi(Z') - \varepsilon\} \leq 1\{\omega : \omega_t \leq \pi_t(Z'|y_1, ..., y_{t-1})\} = 1\{\omega : y_t(\omega_t) \in Z'\} = 1\{z_\tau \in Z'\}$. Let $\nu_t(r, \omega)$ be the number of times $\omega_t \leq r$. Then

$$\nu_t(\pi(Z') - \varepsilon) \leq \sum_{\tau=1}^{t} 1\{z_\tau \in Z'\}.$$

By the strong law of large numbers, $\lim_{t\to\infty} \nu_t(\pi(Z') - \varepsilon) = \pi(Z') - \varepsilon$ a.s.-H_ε. Therefore it must follow that

$$\liminf_{t\to\infty} t^{-1} \sum_{\tau=1}^{t} 1\{z_\tau \in Z'\} \geq \pi(Z') - \varepsilon.$$

Similarly, under equation 34, the set $\{\omega : \omega_t \leq \pi_t(Z'|y_1, ..., y_{t-1})\}$ is included in the set $\{\omega : \omega_t \leq \pi(Z') + \varepsilon\}$. By a similar argument as before,

$$\limsup_{t\to\infty} t^{-1} \sum_{\tau=1}^{t} 1\{z_\tau \in Z'\} \leq \pi(Z') + \varepsilon.$$

□

Lemma 6. *There exists H' with $\mathbf{P}^{\phi^H}(H') = 1$, such that for all $\eta > 0$ and for all $h \in H'$ there exists $t_{\eta,h}$ such that for all $t \geq t_{\eta,h}$ and all $o \in \{A, B\}$, $\overline{P}_t(h)(Z_{ois_i}) > K_p - \eta$, where*

$$K_p \equiv \min_{i,s_i} \left\{ \sum_{\omega \in \Omega} q_i(s_i \mid \omega) p(\omega) \times \min \left\{ (F_i(-2K))^n, (1 - F_i(2K))^n \right\} \right\}. \tag{36}$$

Proof. By restriction on action plans, for all i, s_i, for all h, and for all t

$$F_i(-2K) \leq \alpha_{it}^H(h)(s_i) \leq F_i(2K).$$

Hence, for all i, s_i, for all h, and for all t,

$$\begin{aligned} P\left(z_t \in Z_{Ais_i} \mid h^t\right) &\geq (F_i(-2K))^n \sum_{\omega \in \Omega} q_i(s_i \mid \omega) p(\omega) \\ &\geq K_p, \end{aligned}$$

and, similarly,

$$P\left(z_t \in Z_{Bis_i} \mid h^t\right) \geq K_p.$$

Let $K_p = \pi(Z_{Ais_i})$ (the case of $Z_{Bis_i}$is analogous and thus omitted); then Lemma5 with $\varepsilon = 0$ and $H_\varepsilon = H$ implies that $\liminf_{t \to \infty} \overline{P}_t(h)(Z_{ois_i}) \geq K_p$ a.s. in H. Therefore, this implies that there exists a $H' \subseteq H$ with $\mathbf{P}^{\phi^H}(H') = 1$ such that for all $\eta > 0$ and all $h \in H'$, there exists a $t_{\eta,h}$ such that for all $t \geq t_{\eta,h}$, $\overline{P}_t(h)(Z_{ois_i}) > K_p - \eta$. □

9.2 Limit equilibrium

Let $x_i^n \in \{A, B\}$ be the vote of agent i when there are n players; thus $\{o(A, x_{-i}^n) = A\} = \left\{\frac{1}{n} \sum_{i=1}^n 1\{x_i^n = A\} \geq \rho - \frac{1}{n}\right\}$. We also let $\kappa_i^n(\xi \mid \omega) \equiv P^n(x_i = A \mid \omega)$ (we also use the simplified notation of $\kappa_{i,\omega}^n$ 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 also use the simplified notation of κ_ω^n when ξ is omitted) be the average over all players. Finally, we 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.

9.2.1 Proof of Lemma 2.1

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

Proof. It suffices to show that $\kappa_\theta(\sigma|\cdot)$ is increasing, which follows from Claim 4.1 below (the assumption that σ_θ is increasing and from MLRP (C2) are sufficient assumptions for this claim). $\square$

Claim 2.1.2: 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. $\square$

Claim 2.1.3:

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

Proof. It follows that

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

Suppose that $\rho > \kappa(\sigma \mid \omega)$. By Claim 2.1.2, $\sqrt{n}\,(\rho - \kappa^n(\xi \mid \omega)) \to \infty$ a.s. in Ξ'. Since $n^{-1/2}\sum_{i=1}^{n}(x_i^n - \kappa_i^n(\xi \mid \omega)$ is asymptotically bounded in probability, then $\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega) = 0$ 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'. $\square$

Proof of Lemma 2.1. First, Claim 2.1.3 and the fact that $\kappa(\sigma \mid \cdot)$ is increasing (Claim 2.1.1) imply that there exists $c \in [-1, 1]$ such that

$$\lim_{n\to\infty} P^n(\xi)(o = A \mid \omega) = 1\{\omega < c\} \quad \text{a.s. in } \Xi. \tag{37}$$

Suppose that $c = -1$. 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 $c = 1$. Therefore, $c \in (0, 1)$ and $\boldsymbol{\alpha}$ is Ξ'-asymptotically c-cutoff.

Second, note that, a.s. in Ξ'

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

where the expectation is well-defined because C4(ii) and the facts 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), and where the third line follows from Claim 2.1.3. □

9.2.2 Proof of Lemma 2.2

Throughout the proof let Ξ' be the set in definition 8 and fix $\xi \in \Xi'$ and a strategy mapping $\overline{\boldsymbol{\alpha}}$ such that (13), (14), and (15) are satisfied and $\lim_{n\to\infty} \sigma(\xi; \overline{\boldsymbol{\alpha}}(\xi)) = \sigma$. To simplify notation, 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, let $P^n_{\alpha_i} \equiv P^n\left(\alpha_i, \overline{\boldsymbol{\alpha}}_{-i}(\xi)\right)$. The proof relies on the following claims; the proofs of the first three claims appear at the end of this section.

Claim 2.2.1: For all $\delta > 0$ and $\omega \in \Omega$, there exits $n_{\delta,\omega}$ such that for all $n \geq n_{\delta,\omega}$,

$$\left| P^n_{\alpha_i}\left(o = A \mid \omega, s_i\right) - P^n_{\alpha_i'}\left(o = A \mid \omega, s_i'\right) \right| < \delta$$

uniformly over $i, s_i, s_i', \alpha_i, \alpha_i'$.

Claim 2.2.2: For all $\delta > 0$ there exist n_δ such that for all $n \geq n_\delta$,

$$\left| \Delta_i\left(P^n, s_i\right) - \Delta_i\left(P^n_{\alpha_i}, s_i\right) \right| < \delta$$

uniformly over i, s_i, α_i.

Claim 2.2.3: There exists $c > 0$ and n_c such that for all $n \geq n_c$

$$\Delta_i \left(P^n_{\alpha_i}, s'_i\right) - \Delta_i \left(P^n_{\alpha_i}, s_i\right) \geq c$$

for all i, all $s'_i > s_i$, and $\alpha_i(s'_i) = \alpha_i(s_i)$.

Claim 2.2.*: 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 2.2..* Fix any α_i such that $\alpha_i(s'_i) = \alpha_i(s_i)$. By Claims 2.2.2 and 2.2.3, 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 2.2. The definition of ε-equilibrium (equation 13) implies that for all i, $s'_i > s_i$, $n \geq n_\varepsilon$,

$$\begin{aligned}\overline{\alpha}^n_i(s'_i) - \overline{\alpha}^n_i(s_i) \geq\ & F_i\left(\Delta_i\left(P^n, s'_i\right)\right) - F_i\left(\Delta_i\left(P^n, s_i\right)\right) - 2\varepsilon. \\ &+ F_i\left(\Delta_i\left(P^n, s_i\right) + c'\right) - F_i\left(\Delta_i\left(P^n, s_i\right) + c'\right),\end{aligned} \tag{38}$$

where we have added and subtracted the same term to the RHS. Let $c' > 0$ be as defined in Claim 2.2.*. Since F_i is absolutely continuous, then

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

where the inequality follows from C4(iii). Hence, the sum of the second and fourth terms in the RHS of (38) is at least $c'' > 0$. By Claim 2.2.*, the sum of the first and last terms in the RHS of

(38) 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 2.2.1. 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,

$$\lim_{n\to\infty} P^n\left(o(A, x_{-i}^n) = A \mid \omega\right) - P^n\left(o(B, x_{-i}^n) = A \mid \omega\right) = 0. \tag{39}$$

By simple algebra,

$$P^n\left(o(A, x_{-i}^n) = A \mid \omega\right) - P^n\left(o(B, x_{-i}^n) = A \mid \omega\right) = P^n\left(\sqrt{n}\frac{\rho - \kappa_\omega^n - \frac{1}{n}}{V_\omega^n} < \frac{1}{\sqrt{n}}\sum_{i=1}^n Z_{i\omega}^n \leq \sqrt{n}\frac{\rho - \kappa_\omega^n}{V_\omega^n} \mid \omega\right), \tag{40}$$

where $Z_{i\omega}^n \equiv \frac{\{1\{x_i^n = A\} - \kappa_{i\omega}^n\}}{V_\omega^n}$, $V_\omega^n \equiv \sqrt{\frac{1}{n-1}\sum_{i=1}^n \kappa_{i,\omega}^n\left(1 - \kappa_{i,\omega}^n\right)}$, and where $V_\omega^n > 0$ by step 3. Note that, for a given n, $Z_{1\omega}^n, ..., Z_{n\omega}^n$ are independent, they have zero mean and unit variance, moreover

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

by step 3. Hence by Lindeberg-Levy CLT, it follows that, given ω: $\sum_{i=1}^n \frac{Z_{i\omega}^n}{\sqrt{n}} \Rightarrow N(0,1)$ as $n \to \infty$.

Let $K_\omega^n \equiv \frac{\rho - \kappa_\omega^n}{V_\omega^n}$. Either (a) $\sqrt{n}K_\omega^n \to -\infty$, (b) $\sqrt{n}K_\omega^n \to K \in (-\infty, \infty)$ or (c) $\sqrt{n}K_\omega^n \to \infty$ (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), so that we can take $n \geq n_{M,\epsilon'}$ such that $\sqrt{n}K_\omega^n \leq -M$, 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 $n \geq \max\{n_{\epsilon',\omega}, n_{M,\epsilon'}\}$:

$$\begin{aligned} P^n\left(\sqrt{n}K_\omega^n - \frac{n^{-1/2}}{V_\omega^n} < \frac{\sum_{i=1}^n Z_{i,\omega}^n}{\sqrt{n}} \leq \sqrt{n}K_\omega^n \mid \omega\right) \leq& P^n\left(-M - \frac{n^{-1/2}}{V_\omega^n} < \frac{\sum_{i=1}^n Z_{i,\omega}^n}{\sqrt{n}} \mid \omega\right) \\ \leq& 0.5\epsilon' + \mathcal{L}_N(-M) = \epsilon' \end{aligned}$$

where the second inequality follows from the fact that $n \geq n_{M,\epsilon'}$ and the last inequality follows from CLT (the first term on the RHS) and our choice of M (the second term on the RHS).

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

$$\begin{aligned} P^n\left(\sqrt{n}K^n_\omega - \frac{1}{V^n_\omega\sqrt{n}} < \frac{\sum_{i=1}^n Z^n_{i,\omega}}{\sqrt{n}} \leq \sqrt{n}K^n_\omega \mid \omega\right) \leq & P^n\left(K - \delta < \frac{\sum_{i=1}^n Z^n_{i,\omega}}{\sqrt{n}} \leq K \mid \omega\right) \\ \leq & 0.5\epsilon' + \mathcal{L}_N\left(K - \delta < \frac{\sum_{i=1}^n Z^n_{i,\omega}}{\sqrt{n}} \leq K\right) = \epsilon', \end{aligned}$$

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

Step 2. Note that:

$$\begin{aligned} P^n_{\alpha_i}\left(o = A \mid \omega, s_i\right) = & \alpha_i P^n\left(o(A, x^n_{-i}) = A \mid \omega\right) + (1-\alpha_i)P^n\left(o(B, x^n_{-i}) = A \mid \omega\right) \\ = & P^n\left(o(B, x^n_{-i}) = A \mid \omega\right) \\ & + \alpha_i\left(P^n\left(o(A, x^n_{-i}) = A \mid \omega\right) - P^n\left(o(B, x^n_{-i}) = A \mid \omega\right)\right) \\ \equiv & P^n\left(o(B, x^n_{-i}) = A \mid \omega\right) + \alpha_i \Delta P^n(\omega). \end{aligned}$$

Therefore

$$|P^n_{\alpha_i}\left(o = A \mid \omega, s_i\right) - P^n_{\alpha'_i}\left(o = A \mid \omega, s'_i\right)| \leq |\alpha_i - \alpha'_i||\Delta P^n(\omega)|.$$

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

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

$$\liminf_{n\to\infty} \frac{1}{n-1}\sum_{i=1}^n \kappa^n_{i,\omega}\left(1 - \kappa^n_{i,\omega}\right) > 0. \tag{41}$$

Fix any n and $i \leq n$. By assumption, $\alpha^n_i(s_i) \in [F_i(-2K), F_i(2K)] \subset (0,1)$ for all s_i. Therefore, $0 < \kappa^n_{i,\omega} < 1$ for all ω, thus implying 41. □

Proof of Claim 2.2.2. We prove that

$$\lim_{n\to\infty}\left(E^A_i\left(P^n, s_i\right) - E^A_i\left(P^n_{\alpha_i}, s_i\right)\right) = 0;$$

the proof for the E_i^B terms is similar and therefore omitted. We first show that, for all i, s_i, α_i,

$$E_i^A\left(P_{\alpha_i}^n, s_i\right) = \frac{\int_\Omega P_{\alpha_i}^n\left(o = A \mid \omega, s_i\right) q(s_i \mid \omega) u_i(A, \omega) G(d\omega)}{\int_\Omega P_{\alpha_i}^n\left(o = A \mid \omega, s_i\right) q(s_i \mid \omega) G(d\omega)}$$

is well-defined for sufficiently large n. Fix any i. Assumption C4(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

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

which implies that $\int_\Omega P^n(o = A \mid \omega, s_i^*) G(d\omega) \geq c > 0$. By Claim 2.2.1, for each s_i, α_i, $P^n\left(o = A \mid \omega, s_i^*\right) - P_{\alpha_i}^n\left(o = A \mid \omega, s_i\right)$ converges to zero as $n \to \infty$. Since both probabilities are bounded by one, then the dominated convergence theorem implies that $\int_\Omega \left(P^n\left(o = A \mid \omega, s_i^*\right) - P_{\alpha_i}^n\left(o = A \mid \omega, s_i\right)\right) G(d\omega) \to 0$ as $n \to \infty$, uniformly over α_i. Therefore, there exists $n_{0.5c}$ such that $\sup_{\alpha_i} \left| \int_\Omega [P^n\left(o = A \mid \omega, s_i^*\right) - P_{\alpha_i}^n\left(o = A \mid \omega, s_i\right)] G(d\omega) \right| < 0.5c$ for all $n \geq n_{0.5c}$. So for all $n \geq \max \bar{n}, n_{0.5c} \equiv \overline{n}_c$,

$$\begin{aligned}\int_\Omega P_{\alpha_i}^n\left(o = A \mid \omega, s_i\right) q\left(s_i \mid \omega\right) G(d\omega) &\geq d \int_\Omega P_{\alpha_i}^n\left(o = A \mid \omega, s_i\right) G(d\omega) \\ &> 0.5dc > 0.\end{aligned}$$

Hence, $E_i^A\left(P_{\alpha_i}^n, s_i\right)$ is well defined.

By simple algebra,

$$\begin{aligned}&\left|E_i^A\left(P^n, s_i\right) - E_i^A\left(P_{\alpha_i}^n, s_i\right)\right| = \\ &\leq \frac{\left|\int_\Omega (P^n\left(o = A \mid \omega\right) - P_{\alpha_i}^n\left(o = A \mid \omega\right)) q(s_i \mid \omega) u_i\left(A, \omega\right) G(d\omega)\right| \int_\Omega P^n\left(o = A \mid \omega\right) q(s_i \mid \omega) G(d\omega)}{\int_\Omega P^n\left(o = A \mid \omega\right) q(s_i \mid \omega) G(d\omega) \int_\Omega P_{\alpha_i}^n\left(o = A \mid \omega\right) q(s_i \mid \omega) G(d\omega)} \\ &+ \frac{\left|\int_\Omega (P^n\left(o = A \mid \omega\right) - P_{\alpha_i}^n\left(o = A \mid \omega\right)) q(s_i \mid \omega) G(d\omega)\right| \int_\Omega P^n\left(o = A \mid \omega\right) q(s_i \mid \omega) u_i\left(A, \omega\right) G(d\omega)}{\int_\Omega P^n\left(o = A \mid \omega\right) q(s_i \mid \omega) G(d\omega) \int_\Omega P_{\alpha_i}^n\left(o = A \mid \omega\right) q(s_i \mid \omega) G(d\omega)}\end{aligned}$$

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 $|u_i(A, \omega)| < K$, $q(s|\omega) \leq 1$, and pointwise convergence (for each ω) is obtained by Claim 2.2.1. □

Proof of Claim 2.2.3. Throughout this proof, let $P_i^n(\omega) \equiv P_{\alpha_i'}^n\left(o = A \mid \omega, s_i'\right) = P_{\alpha_i'}^n\left(o = A \mid \omega, s_i\right)$, where the equality follows by conditional independence and because $\alpha_i'(s_i') = \alpha_i'(s_i)$.

Let also $G(\omega \mid o = A, s_i)$ be the conditional conditional distribution function of ω given $o = A$ and s_i when the number of players is n. Also, we assume that in assumption C1, $u(A, \cdot)$ is the one that is strictly increasing and show the result only for this part of $\Delta_i\left(P^n_{\alpha'_i}(\xi),\cdot\right)$.

Step 1. We first show that, for any given i and $s'_i > s_i$. There exists $(\Omega^n)_n$ with $\Omega^n \subseteq \Omega$ and $\liminf_{n\to\infty} \Pr_G(\Omega^n) \equiv \int_{\Omega^n} G(d\omega) > 0$ such that for all $n \geq n_c$ and all $\omega^* \in \Omega^n \setminus \{-1, 1\}$,

$$G(\omega^* \mid o = A, s_i) - G(\omega^* \mid o = A, s'_i) \geq \beta > 0.$$

As shown in the proof of Claim 2.2.2, for all $n \geq n_c$,

$$\int_\Omega P^n_i(\omega) G(d\omega) \geq c$$

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

$$\omega^n_a = min\left\{\omega : \int_{\omega'\leq\omega} P^n_i(\omega') G(d\omega') \geq a\cdot c\right\} \in \Omega.$$

Fix $n \geq n_c$. Then

$$c/4 = \int_{\omega^n_{0.25}\leq\omega\leq\omega^n_{0.50}} P^n_i(\omega) G(d\omega) \leq G\left(\omega^n_{0.50}\right) - G\left(\omega^n_{0.25}\right).$$

Therefore the fact that G has no mass points (assumption C3) implies that

$$\omega^n_{0.50} - \omega^n_{0.25} \geq d_1 > 0. \tag{42}$$

A similar argument establishes that

$$\omega^n_{0.75} - \omega^n_{0.50} \geq d_2 > 0.$$

Let $\Omega^n = [\omega^n_{0.50} - d_3/2, \omega^n_{0.50} + d_3/2]$, where $d_3 \equiv min\{d_1, d_2\} > 0$. Then for all $\omega^* \in \Omega^n$

$$\int_{\omega<\omega^*-d_3/2} P^n_i(\omega) G(d\omega) \geq c/4 \tag{43}$$

and

$$\int_{\omega>\omega^*+d_3/2} P^n_i(\omega) G(d\omega) \geq c/4. \tag{44}$$

In addition, assumption C4(i) and equation 42 imply that $\Pr_G(\Omega^n) \geq c_g > 0$, so that $\liminf_{n\to\infty} \Pr_G(\Omega^n) > 0$. Next, assumption C2 implies that for $\omega' > \omega$ there exists $z > 0$ such that

$$P^n_i(\omega)P^n_i(\omega')\left(q_i\left(s'_i \mid \omega'\right) q_i\left(s_i \mid \omega\right) - q_i\left(s_i \mid \omega'\right) q_i\left(s'_i \mid \omega\right)\right) = zP^n_i(\omega)P^n_i(\omega')q_i\left(s'_i \mid \omega\right) q_i\left(s_i \mid \omega\right)\left(\omega' - \omega\right).$$

Integrating each side 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

$$G(\omega^* \mid o = A, s_i) - G(\omega^* \mid o = A, s_i') = \int_{\omega' > \omega^*} \int_{\omega \le \omega^*} P_i^n(\omega') P_i^n(\omega) z(\omega' - \omega) q_i\left(s_i' \mid \omega\right) q_i\left(s_i \mid \omega\right) G(d\omega) G(d\omega'). \tag{45}$$

Thus, the desired result follows from

$$\begin{aligned} G(\omega^* \mid o = A, s_i) - G(\omega^* \mid o = A, s_i') &\geq zd^2 \int_{\omega' > \omega^* + d_3/2} \int_{\omega < \omega^\star + d_3/2} P_i^n(\omega') P_i^n(\omega)(\omega' - \omega) G(d\omega) G(d\omega') \\ &\geq z \cdot d_3 d^2 \int_{\omega > \omega^\star + d_3/2} P_i^n(\omega) G(d\omega) \int_{\omega < \omega^\star + d_3/2} P_i^n(\omega) G(d\omega) \\ &\geq z d_3 d^2 (c/4)^2 \equiv \beta > 0, \end{aligned}$$

where the last line follows from (43) and (44).

STEP 2. Note that $\int_\Omega u_i(A, \omega)\left(G(d\omega|o = A, s_i') - G(d\omega|o = A, s_i)\right)$; so by integration by parts, the fact that $G(\omega|o = A, s_i') = G(\omega|o = A, s_i)$ for $\omega \in \{-1, 1\}$, and $u_i'(A, \cdot) > 0$ (assumption C1), MLRP (assumption C2) and similar calculations to the ones in step 1 but over all ω, it follows

$$\begin{aligned} -\int_\Omega u_i(A, \omega)\left(G(d\omega|o = A, s_i') - G(d\omega|o = A, s_i)\right) &\geq \int_{\Omega^n} u_i'(A, \omega)(G(d\omega|o = A, s_i') - G(d\omega|o = A, s_i))d\omega \\ &\geq \beta \int_{\Omega^n} u_i'(A, \omega) d\omega \end{aligned}$$

the last inequality follows from step 1. The proof is thus established by noting that,

$$\int_{\Omega^n} u_i'(A, \omega) d\omega = \int_{\Omega^n} \frac{u_i'(A, \omega)}{g(\omega)} g(\omega) d\omega$$

By assumptions C4(i) and C5, $\int_{\Omega^n} u_i'(A, \omega) G(d\omega) \geq \mathit{const.} \times \Pr_G(\Omega^n)$. Finally, since $\liminf_{n\to\infty} \Pr_G(\Omega^n) > 0$ the desired result follows. □

9.3 Perfect limit equilibrium

9.3.1 Proof of Theorem 4

The proof relies on the following claims.

Claim 4.1 If σ is nondecreasing and responsive, and assumptions C2-C3 hold, then $\kappa(\sigma|\cdot)$ is increasing.

Proof. By Bayes theorem and assumption C3, 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). 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

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

where $A_\theta(s_\theta^i) = \sigma_\theta(s_\theta^{i-1}) - \sigma_\theta(s_\theta^i)$ 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, $A_\theta(s_\theta^i) \leq 0$ and for at least one θ, there exists a s: $A_\theta(s) < 0$. Thus, it is easy to see that the expression above is strictly positive. Since $\phi(\theta) > 0$ all θ, the desired result follows from the construction of κ. □

Claim 4.2 $\bar{v}_\theta(s_\theta; .)$ is continuous for all (θ, s_θ).

Proof. By construction of $\bar{v}_\theta(s_\theta; .)$ is sufficient to show that $E[u_\theta(A, \omega)|\omega > c, s_\theta]$ is continuous (the result for $E[u_\theta(B, \omega)|\omega < c, s_\theta]$ is analogous). It follows that $E[u_\theta(A, \omega)|\omega > c, s_\theta] = \int_{\omega > c} u_\theta(o, \omega) \frac{G(d\omega, s_\theta)}{1 - G(c, s_\theta)}$, so by assumptions C1 and C3, $\int_{\omega > c} u_\theta(o, \omega) G(d\omega, s_\theta)$ and $G(c|s_\theta)$ are continuous; this implies continuity of $\bar{v}_\theta(s_\theta; \cdot)$. □

Claim 4.3 *Suppose σ is a perfect limit equilibrium and let c^* the associated perfect equilibrium cutoff. Then:*

(a) $\hat{\kappa}(c^*|\omega) \leq \kappa(\sigma|\omega)$ for all $\omega \in \Omega$; with equality if $c_\theta(s_\theta) \neq c^*$ for all $\theta \in \Theta$ and all $s_\theta \in S_\theta$.

(b) for all $\epsilon > 0$, $\hat{\kappa}(c^* + \epsilon|\omega) \geq \kappa(\sigma|\omega)$ for all $\omega \in \Omega$.

Proof. Part (a). Since σ is a perfect limit equilibrium, there exists a sequence $(\sigma^\eta)_\eta$ such that $\sigma^\eta \to \sigma$ and, by Theorem 3, a sequence of equilibrium cutoffs $c^\eta \to c^*$. We first show that

$$\sigma_\theta(s_\theta) \geq 1\{\overline{v}_\theta(s_\theta; c^*) > 0\},$$

with equality if there does not exist a $s_\theta \in S_\theta$: $\overline{v}_\theta(s_\theta; c^*) = 0$. Note that $\overline{v}_\theta(s_\theta; c^\eta) \to \overline{v}_\theta(s_\theta; c^*)$ (see Claim 4.2). Thus, if $\overline{v}_\theta(s_\theta; c^*) > 0 (< 0)$ then $\sigma^\eta \to 1(0)$ and if $\overline{v}_\theta(s_\theta; c^*) = 0$ then $\sigma^\eta \in [0, 1]$. Therefore, since by definition of perfect limit limit equilibrium $\sigma^\eta \to \sigma$ it must follow that $\sigma_\theta(s_\theta) \geq$

$1\{\overline{v}_\theta(s_\theta;c^*)>0\}$. Moreover, by this construction it is easy to see that if for all $\theta\in\Theta$ and $s_\theta\in S_\theta$: $|\overline{v}_\theta(s_\theta;c^*)|>0$ then $\sigma_\theta(s_\theta)=1\{\overline{v}_\theta(s_\theta;c^*)>0\}$. Now, note that

$$\kappa(\sigma|\omega)=\sum_\theta\sum_{s_\theta\in S_\theta}\sigma_\theta(s_\theta)q_\theta(s_\theta|\omega)\phi(\theta)\geq\sum_\theta\sum_{s_\theta\in S_\theta}1\{\overline{v}_\theta(s_\theta;c^*)>0\}q_\theta(s_\theta|\omega)\phi(\theta)$$

since $1\{\overline{v}_\theta(s_\theta;c^*)>0\}=1\{c_\theta(s_\theta)<c^*\}$ then the RHS equals $\hat{\kappa}(c^*|\omega)$. If $\sigma_\theta(s_\theta)=1\{\overline{v}_\theta(s_\theta;c^*)>0\}$ for all $\theta\in\Theta$ and $s_\theta\in S_\theta$ then $\kappa(\sigma|\omega)=\hat{\kappa}(c^*|\omega)$.

Part (b). If for all $\theta\in\Theta$ and $s_\theta\in S_\theta$: $|\overline{v}_\theta(s_\theta;c^*)|>0$ then by step 1, $\kappa(\sigma|\omega)=\hat{\kappa}(c^*|\omega)$ and since $\hat{\kappa}(\cdot|\omega)$ is nondecreasing, it is easy to see that the desired result holds. Now, suppose there exists $\theta'\in\Theta$ and $s_{\theta'}\in S_{\theta'}$ such that $\overline{v}_{\theta'}(s_{\theta'};c^*)=0$. By definition,

$$\hat{\kappa}(c^*+\epsilon|\omega)=\sum_\theta\sum_{s_\theta\in S_\theta}1\{c_\theta(s_\theta)<c^*+\epsilon\}q_\theta(s_\theta|\omega)\phi(\theta)=\sum_\theta\sum_{s_\theta\in S_\theta}1\{\overline{v}_\theta(s_\theta;c^*+\epsilon)>0\}q_\theta(s_\theta|\omega)\phi(\theta).$$

Since $\overline{v}_{\theta'}(s_{\theta'};.)$ is increasing, $\overline{v}_{\theta'}(s_{\theta'};c^*+\epsilon)>0$, so following the same calculations as those in step (1), for $\sigma_{\theta'}(s_{\theta'})\leq1\{\overline{v}_{\theta'}(s_{\theta'};c^*+\epsilon)>0\}$; for θ such that $\overline{v}_\theta(s,c^*)>0$, then $\sigma_\theta(s)=1\{\overline{v}_\theta(s;c^*+\epsilon)>0\}$; for θ such that $\overline{v}_\theta(s,c^*)<0$, then $\sigma_\theta(s)\leq1\{\overline{v}_\theta(s;c^*+\epsilon)>0\}$, so $\hat{\kappa}(c^*+\epsilon|\omega)\geq\kappa(\sigma|\omega)$. □

Claim 4.4 *Suppose that σ is a responsive perfect limit equilibrium with supporting sequence σ^η and corresponding cutoffs $\{c^\eta\}$ such that $c^*=\lim_{\eta\to0}c^\eta$ (if necessary we take a subsequence). Then σ has the c^*-cutoff property and*

$$c^*=\inf_\omega\{\overline{\kappa}(\omega)\geq\rho\}=\sup_\omega\{\overline{\kappa}(\omega)\leq\rho\}. \tag{46}$$

Proof. By assumption,

$$\kappa(\sigma^\eta\mid\omega)>(<)\rho\quad\text{for }\omega>(<)c^\eta. \tag{47}$$

Since σ is responsive and nondecreasing (because it is the limit of increasing strategies), then $\kappa(\sigma\mid\cdot)$ is increasing by Claim 4.1. Therefore, there exists c such that

$$\kappa(\sigma\mid\omega)>(<)\rho\quad\text{for }\omega>(<)c. \tag{48}$$

We now show that $c=c^*$, so that σ *has the c^*-cutoff property.* Suppose not, so that $c^\eta\to c^*>c$ (the case $c^*<c$ is similar). Let $\omega\in(c,c^*)$. Then there exists $\overline{\eta}$ such that, for all $\eta\leq\overline{\eta}$, $c^\eta>\omega$ and therefore, by (47), $\kappa(\sigma^\eta\mid\omega)<\rho$. In addition, since $\omega>c$ then, by (48), $\kappa(\sigma\mid\omega)>\rho$. These last two results imply that $\lim_{\eta\to0}\kappa(\sigma^\eta\mid\omega)\neq\kappa(\sigma\mid\omega)$, thus contradicting that $\sigma^\eta\to\sigma$.

Next, for all $\omega<c^*$,

$$\hat{\kappa}(c^*\mid\omega)\leq\kappa(\sigma\mid\omega)<\rho,$$

where the first inequality follows from Claim 4.3(a) and the second from (48). Therefore, since $\hat{\kappa}(\cdot \mid \omega)$ is nondecreasing,

$$\hat{\kappa}(c^* - \varepsilon \mid c^* - \varepsilon) = \overline{\kappa}(c^* - \varepsilon) < \rho \tag{49}$$

for all $\varepsilon > 0$. Moreover, Claim 4.3(b) implies that for all $\varepsilon > 0$,

$$\hat{\kappa}(c^* + \varepsilon \mid \omega) \geq \kappa(\sigma \mid \omega) > \rho$$

for all $\omega > c^*$, where the last inequality follows from (48). Therefore,

$$\hat{\kappa}(c^* + \varepsilon \mid c^* + \varepsilon) = \overline{\kappa}(c^* + \varepsilon) > \rho \tag{50}$$

for all $\varepsilon > 0$. Equation (46) follows from (49) and (50). □

First statement in Theorem 4.
If c is a perfect cutoff equilibrium then $c \in C^{eqm}(\rho)$.

Proof. Fix a family $\{F^\eta\}$ and let $\{c^\eta\}$ be a sequence of cutoffs corresponding to the sequence $\{\sigma^\eta\}$ of limit equilibria, where $\sigma^\eta \to \sigma$ and $c^\eta \to c$. Then

$$\forall \omega > c \ \exists \eta_\omega : \ \forall \eta \leq \eta_\omega \ \kappa(\sigma^\eta \mid \omega) > \rho$$

and

$$\forall \omega < c \ \exists \eta_\omega : \ \forall \eta \leq \eta_\omega \ \kappa(\sigma^\eta \mid \omega) < \rho.$$

Since $\kappa(\cdot \mid \omega)$ is continuous, then for all ω,

$$\lim_{\eta \to 0} \kappa(\sigma^\eta \mid \omega) = \kappa(\sigma \mid \omega). \tag{51}$$

By Claim 4.3, for all ω,

$$\kappa(\sigma \mid \omega) \geq \hat{\kappa}(c \mid \omega),$$

where $\hat{\kappa}(c \mid \omega)$, which is defined in (26), is weakly increasing in c; in fact, the above inequality holds with equality for all $\omega \notin \cup_\theta \Omega_\theta$. In order to obtain a contradiction, suppose that $c > \sup_c \{\overline{\kappa}(c) \leq \rho\}$.

Then there exists $c' < c$ such that $\overline{\kappa}(c') > \rho$. Therefore for all $\eta \leq \eta_{c'}$

$$\begin{aligned}\kappa\left(\sigma \mid c'\right) &\geq \hat{\kappa}\left(c \mid c'\right) \\ &\geq \hat{\kappa}\left(c' \mid c'\right) \\ &= \overline{\kappa}(c') \\ &> \rho \\ &> \kappa\left(\sigma^{\eta} \mid c'\right),\end{aligned}$$

thus contradicting (51). A similar proof establishes that it cannot be the case that $c < \inf_c \{\overline{\kappa}(c) \geq \rho\}$. □

If $c \in C^{eqm}(\rho)$, then c is a perfect cutoff equilibrium.

Proof. We divide the proof into 2 cases.

Case (1): Suppose that $(C^{eqm}(\rho))^o = \{\emptyset\}$ (i.e., empty interior), so that

$$c = \inf_{\omega} \{\omega : \overline{\kappa}(\omega) \geq \rho\} = \sup_{\omega} \{\omega : \overline{\kappa}(\omega) \leq \rho\}. \tag{52}$$

Let $X \subset \Theta \times \cup_\theta S_\theta$ be the set of (θ, s_θ) such that

$$\overline{v}_\theta(s_\theta; c) \leq 0 \tag{53}$$

and let $Y \subset \Theta \times \cup_\theta S_\theta$ be the set of (θ, s_θ) such that

$$\overline{v}_\theta(s_\theta; c) > 0. \tag{54}$$

The proof is by construction. We construct a sequence $\{\sigma^\eta\}$ with the property that the solution to $\kappa(\sigma^\eta \mid c^\eta) = \rho$ is given by $c^\eta \to c$. We then specify a feasible family $\{F^\eta\}$ and show that σ^η is a limit equilibrium given F^η for each η.

We construct the sequence $\{\sigma^\eta\}$ as follows. For each $(\theta, s_\theta) \in X$, let $\sigma^\eta_\theta(s_\theta) = z_\theta(s_\theta)\eta$, where $z_\theta : S_\theta \to (0,1)$ is an increasing function. For each $(\theta, s_\theta) \in Y$, let $\sigma^\eta_\theta(s_\theta) = 1 - z'_\theta(s_\theta)\eta$, where $z'_\theta : S_\theta \to (0,1)$ is a decreasing function. By construction

$$\lim_{\eta \to 0} \sigma^\eta_\theta(s_\theta) = 0 \qquad \text{for all } (\theta, s_\theta) \in X \tag{55}$$

and

$$\lim_{\eta \to 0} \sigma^\eta_\theta(s_\theta) = 1 \qquad \text{for all } (\theta, s_\theta) \in Y. \tag{56}$$

Moreover, σ^η is increasing and therefore there exists a unique sequence of cutoffs $\{c^\eta\}$ such that, for all η, $\kappa(\sigma^\eta \mid \omega) > \rho$ for $\omega > c^\eta$ and $\kappa(\sigma^\eta \mid \omega) < \rho$ for $\omega < c^\eta$; let c^* be the limit of any convergent subsequence and let $\sigma = \lim_{\eta\to 0}\sigma^\eta$. Moreover, there exists θ and $s_\theta \neq s_{\theta'}$ such that $\sigma_\theta(s_\theta) = 0$ and $\sigma_{\theta'}(s_{\theta'}) = 1$; otherwise, since

$$\kappa(\sigma|\omega) = \sum_{\theta\in\Theta}\sum_{s_\theta\in S_\theta} q_\theta(s_\theta|\omega)1\{\bar{v}_\theta(s_\theta;c) > 0\}\phi(\theta) = \sum_{\theta\in\Theta}\sum_{s_\theta\in S_\theta} q_\theta(s_\theta|\omega)1\{c_\theta(s_\theta) < c\}\phi(\theta) = \hat{\kappa}(c|\omega),$$

then $\hat{\kappa}(c \mid \omega)$ would be constant, contradicting that $(C^{eqm}(\rho))^o = \{\emptyset\}$. We have then established that σ is responsive. Equation (52) and Claim 4.4 then imply that $c^* = c$.

Next, we construct the family $\{F^\eta\}_{\eta<\overline{\eta}}$ that makes $\{\sigma^\eta\}_{\eta<\overline{\eta}}$ a sequence of limit equilibria. For each η and θ, we choose F^η_θ such that $F^\eta_\theta(\bar{v}_\theta(s_\theta;c^\eta)) = \sigma^\eta_\theta(s_\theta)$ for all s_θ. Since both $\sigma^\eta_\theta(\cdot)$ and $v_\theta(\cdot, c^\eta)$ are increasing–the first by construction, the latter by MLRP–then F^η_θ can be chosen to be increasing. Since $|\bar{v}_\theta(s_\theta;c^\eta)| < 2K$ and $\sigma^\eta_\theta(s_\theta) \in (0,1)$, we can also choose F^η_θ to satisfy $F^\eta_\theta(2K) < 1$ and $F^\eta_\theta(-2K) > 0$. Finally, (53) through (60) imply that F^η_θ can be chosen to satisfy (17).

Last, fix $\eta \leq \overline{\eta}$: we show that σ^η is a limit equilibrium given F^η. We construct the strategy mapping $\boldsymbol{\alpha}^\eta$ by letting $\alpha^{\eta,n}_i(\xi)(s_i) = \sigma^\eta_{\theta_i}(s_i)$ for all $\xi \in \Xi$, all $n, i \leq n, s_i$, where θ_i is the type of player i. First, note that $\lim_{n\to\infty}\sigma^{\eta,n}(\xi) = \sigma^\eta$ for all $\xi \in \Xi$. Second, the fact that the outcome of the election is given by an interior cutoff $c^\eta \in (-1,1)$ implies that $\boldsymbol{\alpha}^\eta$ is asymptotically interior. Finally, note that

$$\begin{aligned}|\alpha^{n,\eta}_{\theta_i}(s_i) - F^\eta_{\theta_i}(\Delta_i(P^n(\xi), s_i))| \leq & |\sigma^\eta_{\theta_i}(s_\theta) - F^\eta(\bar{v}_{\theta_i}(s_i;c^\eta))| \\ & + |F^\eta_{\theta_i}(\bar{v}_{\theta_i}(s_i;c^\eta)) - F^\eta_{\theta_i}(\Delta_i(P^{n,\eta}(\xi), s_i))|.\end{aligned}$$

The first term in the RHS equals zero. By Lemma 2.1, for any $\varepsilon > 0$, there exists $n_{\eta,\varepsilon}$ such that for all $n \geq n_{\eta,\varepsilon}$the second term in the RHS can be made less than ε. Therefore $\boldsymbol{\alpha}^\eta$ is an ε-equilibrium mapping for all $\varepsilon > 0$.

Case (2): Suppose that $(C^{eqm}(\rho))^o \neq \{\emptyset\}$ (i.e., nonpempty interior). Then there exist θ_1 and θ_2 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 $s_\theta \in S_\theta$. [38] Therefore, we can partition the type space as follows: $\Theta = \{\theta_1\} \cup \{\theta_2\} \cup \{\Theta_-\} \cup \{\Theta_+\}$, where $\theta \in \Theta_-$ iff $c_\theta(s^H_\theta) > c_2(s^H_2)$ and $\theta \in \Theta_+$ iff $c_\theta(s^L_\theta) > c_1(s^L_1)$. Note that, for all $c \in (C^{eqm}(\rho))^o \cup \{c_2(s^H_2)\}$ (the proof for $c = c_1(s^L_1)$ is similar and therefore omitted),

$$\bar{v}_\theta(s_\theta;c) < 0 \qquad \text{for all } \theta \in \Theta_- \cup \{\theta_2\}, \text{ all } s_\theta \in S_\theta \tag{57}$$

and

$$\bar{v}_\theta(s_\theta;c) \geq 0 \qquad \text{for all } \theta \in \Theta_+ \cup \{\theta_1\}, \text{ all } s_\theta \in S_\theta, \tag{58}$$

where the last inequality holds with equality if and only if $c = c_2(s^H_2)$ and $\theta = \theta_2$.

Fix any $c \in (C^{eqm}(\rho))$. The proof is by construction. We construct a sequence $\{\sigma^\eta\}$ with the

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

property that the solution to $\kappa(\sigma^\eta \mid c^\eta) = \rho$ is given by $c^\eta = c$ for all η. Therefore, the sequence of cutoffs $\{c^\eta\}$ trivially converges to c. We then specify a feasible familiy $\{F^\eta\}$ and show that σ^η is a limit equilibrium given F^η for each η.

We construct the sequence $\{\sigma^\eta\}$ as follows. As in case (1), let $\sigma^\eta_\theta(s_\theta) = z_\theta(s_\theta)\eta$ for each $\theta \in \Theta_-$ and $s_\theta \in S_\theta$ as well as for $\theta = \theta_2$ and all $s_2 \neq s_2^H$. In addition, let $\sigma^\eta_\theta(s_\theta) = 1 - z'_\theta(s_\theta)\eta$ for each $\theta \in \Theta_+$ and $s_\theta \in S_\theta$ as well as for $\theta = \theta_1$ and all $s_1 \neq s_1^L$. Finally, let $\sigma_1(s_1^L) = 1 - d_1\eta$, where we leave d_1 and $\sigma_2(s_2^H)$ unspecified for the moment. It follows that

$$\begin{aligned}\kappa(\sigma^\eta|c) &= \sum_\theta \sum_{s_\theta \in S_\theta} q_\theta(s_\theta|c)\sigma^\eta(s_\theta)\phi(\theta) \\ &= \phi(\Theta_-) + (B(c) - A(c))\eta - d_1 q_1(s_1^L|c)\phi(\theta_1)\eta + q_2(s_2^H|c)\sigma_2^\eta(s_2^H)\phi(\theta_2),\end{aligned}$$

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 $\sigma_2^\eta(s_2^H)$ as the strategy that solves $\kappa(\sigma^\eta|c) = \rho$, which is given by

$$\sigma_2^\eta(s_2^H) = D(c, d_1)\eta,$$

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 $\sigma_\theta^{\overline{\eta}}(s_\theta) \in (0, 1)$ for all θ, s_θ. By construction, for $\eta < \overline{\eta}$, $\{\sigma^\eta\}$ is a sequence of *increasing* strategy profiles, all of which are characterized by cutoff c and satisfy:

$$\lim_{\eta\to 0} \sigma_\theta^\eta(s_\theta) = 1 \qquad \text{for all } \theta \in \Theta_+ \cup \{\theta_1\}, \text{ all } s_\theta \in S_\theta \tag{59}$$

and

$$\lim_{\eta\to 0} \sigma_\theta^\eta(s_\theta) = 0 \qquad \text{for all } \theta \in \Theta_- \cup \{\theta_2\}, \text{ all } s_\theta \in S_\theta. \tag{60}$$

Exactly as in case (1), we can then construct the family $\{F^\eta\}_{\eta<\overline{\eta}}$ that makes $\{\sigma^\eta\}_{\eta<\overline{\eta}}$ a sequence of limit equilibria. □

Second statement in Theorem 4

If a perfect limit equilibrium σ is responsive then $\inf_\omega \{\overline{\kappa}(\omega) \geq \rho\} = \sup_\omega \{\overline{\kappa}(\omega) \leq \rho\}$.

Proof. Let $\{\sigma^\eta\}$ be a sequence of limit equilibria that converges to σ. By Theorem 3, there is a corresponding sequence of equilibrium cutoffs $c^\eta \to c^*$ (if not, take a subsequence). By Claim 4.4, $c^* = \inf_\omega \{\overline{\kappa}(\omega) \geq \rho\} = \sup_\omega \{\overline{\kappa}(\omega) \leq \rho\}$. □

If σ is a perfect limit equilibrium and $c = \inf_\omega \{\omega : \overline{\kappa}(\omega) \geq \rho\} = \sup_\omega \{\omega : \overline{\kappa}(\omega) \leq \rho\}$, then σ is responsive.

Proof. Suppose not, then for all θ, $\sigma_\theta(s_\theta) = \sigma_\theta$. By definition of perfect limit equilibrium, there exists sequences of $(\sigma^\eta)_\eta$ such that $\sigma^\eta \to \sigma$; let $(c^\eta)_\eta$ denote the sequence of associated cutoffs, let $c^* \equiv \lim_{\eta\to 0} c^\eta$ (we go to a subsequence if necessary). Thus, by Claim 4.2,

$$\bar{v}_\theta(s_\theta; c^\eta) \to \bar{v}_\theta(s_\theta; c^*).$$

Assume that $c_\theta(s_\theta) \neq c^*$ for all θ and s_θ, then for all θ and $s_\theta \in S_\theta : |\bar{v}_\theta(s_\theta; c^*)| > 0$; the previous expression implies that for all θ and $s_\theta \in S_\theta : |\bar{v}_\theta(s_\theta; c^\eta)| > 0$ for $\eta \leq \bar{\eta}(\theta, s_\theta)$. Hence, $\sigma^\eta \to \sigma = 1\{\bar{v}_\theta(s_\theta; c^*) > 0\}$. Since σ is non-responsive, it ought to be the case that: either $\bar{v}_\theta(s_\theta; c^*) > 0$ for all $s_\theta \in S_\theta$, or $\bar{v}_\theta(s_\theta; c^*) < 0$ for all $s_\theta \in S_\theta$. Moreover, since by Claim 4.3(a) $\kappa(\sigma|\omega) = \hat{\kappa}(c^*|\omega) \equiv \hat{\kappa}(c^*)$, this implies that $\kappa(\sigma^\eta|\omega) \to \hat{\kappa}(c^*)$ uniformly over ω. Hence, $|\kappa(\sigma^\eta|c^\eta) - \hat{\kappa}(c^*)| \leq |\kappa(\sigma^\eta|c^\eta) - \hat{\kappa}(c^*|c^\eta)| \to 0$, thereby implying $\rho = \hat{\kappa}(c^*)$. So, $\bar{\kappa}(c^*) = \hat{\kappa}(c^*) = \rho$; if $c \neq c^*$ then we arrive to a contradiction since under this case $\inf_\omega \{\overline{\kappa}(\omega) \geq \rho\} \neq \sup_\omega \{\overline{\kappa}(\omega) \leq \rho\}$; so we assume $c = c^*$. Let $[c - \delta, c + \delta]$ where $\delta > 0$ is chosen such that $\delta < \min_{\theta, s_\theta} |c - c_\theta(s_\theta)|$ (this distance is positive, because $c_\theta(s_\theta) \neq c$ and there only finitely many (θ, s_θ) – this is also why we take "min" and not "inf"). For any $\tilde{c} \in [c - \delta, c + \delta]$, $S_\theta(\tilde{c}) = S_\theta(c)$ so that $\hat{\kappa}(\tilde{c}) = \hat{\kappa}(c) = \rho$. Therefore, $\bar{\kappa}(\tilde{c}) = \rho$ for all $\tilde{c} \in [c - \delta, c + \delta]$, thus contradicting $\inf_\omega \{\overline{\kappa}(\omega) \geq \rho\} = \sup_\omega \{\overline{\kappa}(\omega) \leq \rho\}$.

Assume now there exists a $(\theta', s'_{\theta'})$ such that $c_{\theta'}(s'_{\theta'}) = c^*$. Hence, by Claim4.3, $\hat{\kappa}(c^* + \epsilon) \equiv \hat{\kappa}(c^* + \epsilon, \omega) \geq \kappa(\sigma|\omega) \geq \hat{\kappa}(c^*|\omega) \equiv \hat{\kappa}(c^*)$ for any $\epsilon > 0$. Doing similar algebra as before, $\hat{\kappa}(c + \epsilon) \geq \kappa(\sigma^\eta|c^\eta) \geq \hat{\kappa}(c)$ for small η and hence $\bar{\kappa}(c^* + \epsilon) \geq \rho \geq \bar{\kappa}(c)$. Since this holds for any $\epsilon > 0$, it ought to be the case that $c = c^*$. By choosing $\delta < \min_{\theta, s_\theta \neq \theta' s'_{\theta'}} |c - c_\theta(s_\theta)|$, it follows $\rho \geq \bar{\kappa}(c + \delta)$, and thus $\sup_\omega \{\omega : \overline{\kappa}(\omega) \leq \rho\} > c$ a contradiction. □

9.3.2 Proof of Theorem 5

Only if. Suppose that σ is a responsive perfect limit equilibrium with supporting sequence σ^η and corresponding cutoffs $\{c^\eta\}$ such that $c^* = \lim_{\eta\to 0} c^\eta$ (if necessary we take a subsequence). Then

$$\begin{aligned}\lim_{\eta\to 0} v_\theta(s_\theta; \sigma^\eta) &= \lim_{\eta\to 0} \overline{v}_\theta(s_\theta; c^\eta) \\ &= \overline{v}_\theta(s_\theta; c^*) \\ &= v_\theta(s_\theta; \sigma),\end{aligned}$$

where the first line follows because σ^η has the cutoff property, the second line by Claim 4.2, and the third line because σ has the c^*-cutoff property (Claim 4.2). The fixed-point property then follows from Lemma 3. □

If. Since σ is responsive and has the cutoff property, then $\bar{v}_\theta(s_\theta; c) = v_\theta(s_\theta; \sigma)$; we can thus construct σ^η as in the proof of Case (1) in Theorem 4 except that for the set of (θ, s_θ) such that $\overline{v}_\theta(s_\theta, c) = 0$

we now set $\sigma_\theta^\eta(s_\theta) = \sigma_\theta(s_\theta)$ for all η. Note that, for η small enough, $\sigma_\theta(s_\theta) < \sigma_\theta^\eta(s'_\theta)$ for $s'_\theta > s_\theta$ and $\sigma_\theta(s_\theta) > \sigma_\theta^\eta(s'_\theta)$ for $s'_\theta < s_\theta$ (for any (θ, s_θ) such that $\overline{v}_\theta(s_\theta, c) = 0$). Therefore, for η small enough, $\sigma_\theta^\eta(.)$ and $\overline{v}_\theta(.; c^\eta)$ are both increasing, and thus we can choose F_θ^η increasing. The proof thus follows in the same way as the proof of Case (1) in Theorem 4. □

References

[1] Ali, Nageeb, Jacob K. Goeree, Navin Kartik, and Thomas Palfrey. 2008. "Information Aggregation in Standing and Ad Hoc Committees," American Economic Review Papers & Proceedings, 98(2): 181-186.

[2] Al-Najjar, Nabil. Forthcoming "Decision Makers as Statisticians: Diversity, Ambiguity and Learning." Econometrica.

[3] Al-Najjar, Nabil and Mallesh Pai. 2009. "Coarse Decision Making." working paper.

[4] Aragones, Enriqueta, Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler. 2005. "Fact-Free Learning." American Economic Review, 95(5): 1355-1368.

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

[6] Battaglini, Marco, Morton, Rebecca and Thomas Palfrey. Forthcoming "The Swing Voter's Curse in the Laboratory." Review of Economic Studies.

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

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

[9] Bolton, P. and C. Harris. 1999. "Strategic Experimentation," Econometrica, 67(2): 349-374.

[10] Charness, G. and D. Levin. 2006. "The Origin of the Winner's Curse: A Laboratory Study." Working paper.

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

[12] Costinot, Arnaud and Navin Kartik. 2007. "On Optimal Voting Rules under Homogeneous Preferences." Working paper.

[13] Dekel, Eddie, Drew Fudenberg, and David K. Levine. 2004. "Learning to Play Bayesian Games." Games and Economic Behavior, 46: 282-303.

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

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

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

[17] 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.

[18] Fudenberg, Drew and David Kreps. 1993. "Learning Mixed Equilibria," Games and Economic Behavior, 5: 320-367.

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

[20] Fudenberg, Drew and David K. Levine. 1993b. "Steady-State Learning and Nash Equilibrium." Econometrica, 61: 547-573.

[21] Fudenberg, Drew and David K. Levine. 1998. The Theory of Learning in Games. The MIT Press.

[22] Fudenberg, Drew and David K. Levine. 2009. "Learning and Equilibrium," Annual Review of Economics 1: 385-420.

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

[24] Guarnaschelli, Serena, McKelvey, Richard and Thomas Palfrey. 2000. "An Experimental Study of Jury Decision Rules." American Political Science Review, v. 94, 2, 407-23.

[25] 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.

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

[27] Kagel, J. H. and D. Levin. 2002 Common Value Auctions and the Winner's Curse, Princeton, NJ: Princeton University Press.

[28] Keller, G., S. Rady, and M. Cripps. 2005. "Strategic Experimentation with Exponential Bandits," Econometrica, 73(1): 39-68.

[29] Krishna, Vijay and John Morgan. 2008. "Voluntary Voting: Costs and Benefits," working paper.

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

[31] 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.

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

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

[34] Milgrom, Paul R. 1981b. "Rational Expectations, Information Acquisition, and Competitive Bidding," Econometrica, vol. 49(4), 921-943.

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

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

[37] Pesendorfer, Wolfgang and J. M. Swinkels. 1997. "The Loser's Curse and Information Aggregation in Common Value Auctions," Econometrica, Vol. 65, pp. 1247- 1281.

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

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

[40] Schwartzstein, Josh. 2009. "Incorrect Beliefs," working paper.

[41] Selten, Reinhard. 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory 4: 25–55.

[42] 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.

[43] Strulovici, Bruno. forthcoming. "Learning While Voting: Determinants of Collective Experimentation," Econometrica.