

Premise-Based versus Outcome-Based Information Aggregation

Geoffroy de Clippel and Kfir Eliaz*

This Version: January 2013

Abstract

A group of rational individuals with common interest need to select one of two outcomes. The optimal decision depends on whether certain premises or pieces of evidence are established as being true, and each member receives a noisy signal of the truth value of the relevant premises. Should the group reach a decision by voting whether each premise is true or false, or should they simply vote on the outcome? We show that for any finite number of individuals, the premise-based voting rule is more efficient in aggregating information than the outcome-based rule. However, generically, the gain from using the premise-based over the outcome-based rule can only be marginal when numerous individuals express independent opinions. Indeed, the outcome-based game is almost always asymptotically efficient.

*Brown University, Department of Economics, Providence, Rhode Island - declippel@brown.edu, kfir.eliaz@brown.edu. We wish to thank Philippe Mongin for introducing us to the topic and Kareen Rozen for useful comments. Special thanks go to Neil Thakral, an exceptional Brown undergrad, for superb research assistance.

1. INTRODUCTION

Many collective decisions share the following features. A group of individuals need to decide which of two possible actions to take. The members of the group agree that the optimal action depends on whether certain premises or pieces of evidence are established as being true. For example, in a criminal trial, a jury (or alternatively, judges in an appeals court) needs to decide whether the defendant is guilty or innocent. The jurors may agree that the defendant should be convicted if certain *premises* are found to be true, e.g., if there are reliable eyewitnesses who saw him at the scene of the crime, if a weapon was found with the defendant's fingerprints, if he could not produce a credible alibi, etc. Another example is that of a tenure decision in academia. The members of the committee may agree that the person should be granted tenure if it is established that the candidate has had an impact on the profession, and/or he has sufficiently many quality publications, and/or he is a good citizen, etc. Similar considerations arise when the representatives of a nation or a group of nations need to decide whether or not to implement sanctions against some country. The decision will depend on whether the individuals members believe that some set of premises are true, e.g., the country in question is developing weapons of mass destruction, it is committing crimes against humanity, it has violated international treaties, and so forth.

The most common procedure for reaching a group decision is to hold a vote. Oftentimes, individual members of the group have only partial or imprecise information on the truthfulness of the relevant premises, hence, a vote may help in aggregating the members' beliefs about the validity of the premises. A natural question that arises is whether the group decision depends on whether one aggregates the individual members' beliefs about the relevant premises or over the action to be taken. This question is at the heart of the "doctrinal paradox" or "discursive dilemma," first introduced and studied by Kornhauser and Sager (1986), Pettit (2001), and Brennan (2001), and which has received an increased interest over the past decade in varied academic literatures including computer science, economics, law, philosophy, and political science (see e.g. List and Puppe (2009) for a survey).

The “paradox” or “dilemma” stems from the fact that the group decision may be different depending on which aggregation method is used: premise-based or outcome-based. To illustrate this, consider the jury example described above. Suppose the jurors agree that the defendant is guilty if and only if two pieces of evidence are established. Each juror has formed an opinion regarding the validity of the relevant piece of evidence, and these opinions are aggregated by a majority vote. Suppose that a third of the jury is convinced that both pieces of evidence are established, a second third is convinced that only the first piece of evidence is established, while the rest is convinced that only the second piece of evidence is established. If members of the jury vote truthfully, then the defendant will go free if the vote is on the final verdict, while both pieces of evidence would pass - resulting in a guilty verdict - if voting on evidence. Beyond this simple example of logical conjunction over two premises, numerous results have been established to show the impossibility of finding aggregation methods that deliver logically consistent judgments. Various applications have been cited in the literature, including those listed above.

These impossibility results motivate the next question to investigate, which is to determine which approach – aggregating opinions about premises versus outcomes – is best. Different criteria are conceivable. The purpose of the present paper is to compare the outcome-based versus the premise-based procedures in terms of their ability to aggregate information in the presence of strategic individuals with common interest. Inspired by the Condorcet Jury theorem, we assume that each individual independently receives some noisy signals regarding the truth value of the premises. Much of the literature on voting and information aggregation has focused on the case of common interest: there is a “truth” out there (e.g., guilty/innocent), as well as an unambiguous optimal action, given the truth. Opinions may vary because of different interpretations or realizations of signals.¹ We, therefore, assume that individuals

¹Some notable examples include Austen-Smith and Banks (1996), Feddersen and Persson (1998), McLennan (1998) and the references therein, and more recently, Persico (2004).

all share the same standard regarding which combinations of premises must lead to a positive decision. As in the doctrinal paradox, we consider two aggregation procedures, whereby individuals submit their opinions regarding either premises or outcomes. A critical difference, though, is that these reports need not be truthful (in the sense of systematically reporting their signals). For instance, an individual is free to report that some premise is positive while he received a null signal, or to report that a positive decision should be taken while his signal is such that a null decision should be taken. Why might he be inclined to do so? There are at least two reasons. First he may be more concerned about incorrectly taking a positive decision versus incorrectly taking the null decision, or vice versa. Second, an individual can influence the truth value of a premise, or the final decision, only when he is pivotal, and his opinion conditional on being pivotal may be different from his signal. We will thus investigate and compare the outcomes that can be supported by Bayesian Nash equilibrium (BNE) strategies in both the premise and the outcome-based games.²

Our first result (Proposition 1) establishes that *any* symmetric BNE of the outcome-based game can be replicated by a symmetric BNE of the premise-based game.³ This holds for any finite group of individuals, for any supermajority rule that is not unanimous and for a large class of common preferences. The result holds independently of the logical connection between premises and outcomes, independently of the probability distribution of the true states, and independently of the probability distributions of the signals conditional on the states.

Proposition 1 implies that for *any* finite group of individuals, gathering opinions about premises is systematically at least as good as gathering opinions about outcomes. Our second contribution is to show that the converse result is *not* true. In particular, there are cases where the ex-ante welfare of a

²We do not model any communication that may occur between the individual members of the group. This direction is left for future research.

³We restrict attention to symmetric BNEs throughout the paper because these are the most natural in our symmetric environment, but all our results remain valid over the set of all BNEs.

symmetric BNE of the premise-based game is strictly larger than the ex-ante welfare of any symmetric BNE in the outcome-based game (see Example 1).

While making collective decisions by gathering opinions regarding premises can lead to a strictly higher ex-ante welfare in the presence of finitely many individuals, our third result is to show that, generically, gains over the outcome-based approach can only be marginal when sufficiently many individuals express independent opinions. Indeed, we prove that the outcome-based game is almost always asymptotically efficient (see Proposition 3).⁴ To the best of our knowledge, this is the most general result concerning the asymptotic efficiency of outcome-based voting. Hence, as a corollary of Proposition 1, the premise-based game is almost always asymptotically efficient, as well. Yet we also provide an alternative sufficient condition for asymptotic efficiency in the premise-based game (see Proposition 2) that is useful to show that there exist (non-generic) cases where the premise-based game is asymptotically efficient, while the outcome-based game is not (see Example 2).

Related Literature

Since we focus on voters with common interest, our paper is naturally related to the seminal work by McLennan (1998). That paper established that ex-ante efficient symmetric strategy profiles must be BNE in games with common interest. We make use of this result in our analysis of asymptotic efficiency (see the proofs of Propositions 2 and 3).

Strategic considerations were first introduced to the literature on the doctrinal paradox by Dietrich and List (2007). Instead of investigating information aggregation, they investigate which of the two rules is more robust against strategic manipulations in a context where individuals want the final decision to be as close as possible to their exogenous opinion (private value setting). In our model, instead, individuals share the common preference of taking a collective decision that is best given the true state, while their opinions are

⁴Our treatment of asymptotic efficiency follows the tradition of virtually all works on information aggregation (e.g., Feddersen and Pesendorfer (1997,1998) and more recently, Ahn and Oliveros (2011a,b) and the references therein).

derived from noisy signals of that state. Their results are then quite different. For instance, truth-telling is a weakly dominant strategy in the outcome-based game of their model, while it need not even be a Bayesian Nash equilibrium in ours.

Bozbay et al. (2011) are interested in characterizing procedures for aggregating reports on the truthfulness of premises for which truth-telling is an efficient BNE. Note that they take a mechanism-design approach, while we study a fixed mechanism (any given super-majority) in two different scenarios (premise vs. outcome-based). Bozbay et al. (2011) focus on the case of two independent signals on two premises, where the disutility from making a mistake is independent of its type (e.g. convicting an innocent versus letting a guilty person go free). Their main results are concerned with a class of mechanisms, referred to as “quota rules”, that precludes the outcome-based voting rule.⁵ These rules are characterized by two numbers representing the percentage of positive votes needed on each of the two premises to make it pass before applying the society standard to determine the final decision.⁶ The authors provide a condition on the parameters of their model that is necessary and sufficient to guarantee the existence of a quota rule for which truth-telling forms an efficient BNE of the premise-based game.⁷

The existence of a quota rule that induces an efficient BNE in truthful strategies may be viewed as a knife-edge result in the following sense: whenever such a rule exists, the pair of premise-specific thresholds must be unique (see their Theorem 4 and Corollary 2). Their results are irrelevant for our model whenever these two thresholds differ. A difficulty with their approach is that the mechanism designer needs to know precisely the parameters of the

⁵The outcome-based procedure violates their independence axiom. It also violates their monotonicity axiom when the society standard is different from classical conjunction or disjunction (or, in their terminology, when individuals have consequentialist preferences of the first type). Results that rely on these axioms cannot apply to our outcome-based game.

⁶Our premise-based game is thus a special case of quota rule where both quotas are equal.

⁷Bozbay et al. also present other results (their Corollary 1, Proposition 1, and Theorem 5) that extend beyond the class of quota rules, and which also apply to outcome-based voting.

model if his objective is to design a quota rule for which truth-telling is an efficient BNE. Furthermore, he would have to change the mechanism whenever a new situation with different parameters occurs. Even though truth-telling need not be a BNE when the same quotas are applied to all premises, there will always be other BNEs. The best BNEs may have some focal attractiveness, and changing the equilibrium play as the underlying parameters of the model change may be viewed as a milder requirement compared to the need of changing the mechanism altogether. Nevertheless, the results of Bozbay et al. may be helpful in deriving conditions on the parameters of a model that are necessary and sufficient for a given mechanism- such as the premise-based or outcome-based voting rule - to admit an efficient truthfull equilibrium.

We now turn our attention to papers that relate only to our asymptotic results. A few authors have investigated the “truth-tracking” properties of outcome-based and premise-based aggregation methods - see List (2005) and Bovens and Rabinowicz (2006) for the earliest results on the topic (see also List (2006, Section 6)). The difference between our approaches is that they assume that individuals report their opinions truthfully (as in the original Condorcet Jury theorem). Accounting for the incentives of strategic individuals leads to fundamentally different asymptotic results. Indeed, the main point of the small truth-tracking literature is that premise-based and outcome-based aggregation methods are not systematically comparable, as one may dominate the other and vice versa as a function of the relative reliability of the signals on each premise.

The systematic study of strategic multi-issue voting and information aggregation has been initiated by Ahn and Oliveros (2011a, 11b). The first paper is less relevant as it focuses on pure private values. The second paper, on the other hand, compares the asymptotic efficiency of a joint trial vs. two separate trials for two defendants in an environment with common values. Given the generality of Ahn and Oliveros’ (2011b) framework, our premise-based game could essentially be seen as a particular case of their joint trial scenario with only a binary decision to be chosen. On the other hand, our outcome-based game bears little resemblance with their split trial scenario. Comparing the

type of asymptotic results we derive,⁸ two differences are worth noting. First, we provide an example where the premise-based game is asymptotically strictly superior to the outcome-based game, while their main result shows that the joint trial game is asymptotically efficient if and only if their split trial game is. Second, the translation of their sufficient condition for asymptotic efficiency in our setting is much less permissive than the ones we derive. In particular, their condition would be far from being generic in the outcome based approach. This point will be further discussed in Section 5.

2. MODEL

There exists a set of K issues or premises that are relevant for some binary decision. Each premise may be either true (denoted by the value 1) or false (denoted by 0). The actual state of nature is thus a vector $\omega \in \{0, 1\}^K$. The relative likelihood of these various states is captured by a probability distribution $\pi \in \Delta(\{0, 1\}^K)$. There are n individuals and each individual i receives only a noisy signal of the actual state - $s_i \in \{0, 1\}^K$. Signals are drawn independently (conditional on the state) across individuals. Let $p : \{0, 1\}^K \rightarrow \Delta(\{0, 1\}^K)$ be the function that describes the relative likelihood of the various signals as a function of the actual state: $p_\omega(x)$ is the probability of receiving the signal $x \in \{0, 1\}^K$ conditional on the fact that the actual state is ω .

The *society's standard*, $f : \{0, 1\}^K \rightarrow \{0, 1\}$, determines which decision to take - “0 = *fail*” vs. “1 = *pass*” - as a function of the validity of the various underlying premises. In order for the problem to be non-trivial, we assume that there exist a state ω such that $f(\omega) = 0$ and another state ω' such that $f(\omega') = 1$. Simple examples include logical conjunction (unanimity) - all the underlying premises must be true in order for the decision to be positive -, logical disjunction - at least one underlying premise must be true in order for the decision to be positive -, and majority - at least $\frac{K+1}{2}$ of the underlying premises must be true in order for the decision to be positive. Obviously, a society's standard may also include more complicated logical relations between

⁸Ahn and Oliveros (2011b) focus exclusively on asymptotic results.

premises and decisions. For instance, the first premise together with a majority of the remaining ones must be true in order for the decision to be positive.

All the standards provided so far happen to be monotone: a positive decision is taken at a state ω whenever there exists a state ω' such that $\omega \geq \omega'$ and a positive decision is taken at ω' . Note that one can think of yet other natural examples where the standard is not monotone. For instance, a country may decide that it needs to expand its military capability only if there is a threat either from the north or the south, but not if intelligence shows that neither the north nor the south have hostile plans, nor if both the north and the south have hostile plans (e.g. because there is no way to defend oneself against a coordinated attack). Our results do not rely on monotonicity, and are valid for any non-trivial standard.

It is assumed that each individual shares the same standard as the society. Individuals have the common objective of trying to take the decision that is right for the actual state. This means that individuals' utilities are defined by comparison between the actual decision and $f(\omega)$: 0 if match, $-q$ if the decision is strictly below $f(\omega)$, and $-(1 - q)$ if the decision is strictly above $f(\omega)$.

A *super-majority voting rule* is a general method to make collective judgments. Given a threshold $\theta \in (\frac{1}{2}, 1)$, a judgement is positive if at least a fraction θ of the population votes in its favor.⁹ Given a θ -majority voting rule, decisions can then be taken via two different natural mechanisms. In a *premise-based* approach, individuals are asked to submit ballots regarding the validity of all premises. The θ -majority rule is then applied premise-by-premise, and the society's standard is applied to determine the final decision. Formally, individual i 's *ballot* in the premise-based game is a vector $b_i \in \{0, 1\}^K$. Let then $\tau^k(b^k)$ be premise k 's revealed truth value: 1 if $\sum_{i \in I} b_i^k \geq \theta n$, and 0 otherwise. The final outcome associated with the profile b of ballots is $f((\tau^k(b^k))_{k=1}^K)$. Individual i 's *strategy* in this game is a function $\beta_i : \{0, 1\}^K \rightarrow \Delta(\{0, 1\}^K)$,

⁹In this paper we take the voting rule to be exogenously given and compare the equilibria of two voting mechanism. An interesting question, left for future research, is how to optimally design the voting rule for each voting procedure.

where $\beta_i(s_i)$ is the probability distribution with which he will decide which ballot to submit, conditional on his signal s_i .

In an *outcome-based* approach, each individual is only asked to submit his opinion regarding the final decision - not the premises. Individual i 's *ballot* in that case is just an element $\hat{b}_i \in \{0, 1\}$. The decision is positive if and only if at least a fraction θ of the ballots are equal to 1. Individual i 's *strategy* in this game is a function $\hat{\beta}_i : \{0, 1\}^K \rightarrow \Delta(\{0, 1\})$. With slight abuse of notation, $\hat{\beta}_i(s_i)$ will denote the probability with which he will vote for the decision to pass, conditional on his signal s_i .

In either game, a profile of strategies form a *Bayesian Nash equilibrium* (BNE) if the action prescribed by the strategy of each player is optimal for each signal he might receive. While not needed in any of our results, we will restrict attention to symmetric BNEs. Given that both games are symmetric, these equilibria are more intuitive and more focal.

3. SUPERIORITY OF THE PREMISE-BASED APPROACH

Proposition 1 *Let $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_n)$ be a symmetric BNE in the outcome-based game. Then there exists a symmetric BNE in the premise-based game, $\beta = (\beta_1, \dots, \beta_n)$, such that for every vector of signal realizations, the strategy profile β induces the same probability distribution over decisions as $\hat{\beta}$.*

Proof: By assumption, there exists a state ω such that $f(\omega) = 0$. For notational simplicity, one may redefine the truth value of premises so that $f(0, \dots, 0) = 0$. The rest of the proof is written under this assumption. We also know by assumption that there exists $\bar{\omega}$ such that $f(\bar{\omega}) = 1$.

For every voter i , and for every signal realization s_i , define $\beta_i(s_i)$ to be a mixed strategy that plays action $\bar{\omega}$ with probability $\hat{\beta}_i(s_i)$ and plays the action $(0, \dots, 0)$ with probability $1 - \hat{\beta}_i(s_i)$. It follows that for a given vector of signal realizations, the probability that the decision is passed under β is equal to the probability that at least a fraction θ of the voters chose $\bar{\omega}$, which in turn, equals to the probability that at least a fraction θ of the voters chose the action 1 under $\hat{\beta}$.

Assume β is not a BNE. Then there is a player i , a signal s_i and a *pure* action a_i , such that, given s_i and β_{-i} , player i strictly prefers to choose $a_i \in \{0, 1\}^K$ over $\beta_i(s_i)$. Voter i 's expected utility when playing a_i after receiving a signal s_i is equal to

$$\sum_{\omega \in \Omega} \text{Prob}(\omega | s_i) \sum_{s_{-i}} \text{Prob}(s_{-i} | \omega, s_i) \sum_{a_{-i} \in \{0, 1\}^{K \times (I \setminus \{i\})}} \text{Prob}(a_{-i} | s_{-i}, \beta_{-i}) u_i(a, \omega),$$

where $u_i(a, \omega)$ denotes i 's utility when each player j has submitted the ballot a_j (which determines the collective decision to be implemented) and ω is the true state.

By definition of β , $\text{Prob}(a_{-i} | s_{-i}, \beta_{-i})$ is positive only if $a_j = (0, \dots, 0)$ or $\bar{\omega}$, for all $j \in I \setminus \{i\}$. Let X_i be the subset of such ballot profiles with k^* ballots of the form $(0, \dots, 0)$ and $n - 1 - k^*$ ballots of the form $\bar{\omega}$. Voter i 's action has no impact on the final outcome when $a_{-i} \notin X_i$, and we can thus ignore those terms when computing i 's optimal action. In other words, it must be that

$$\sum_{\omega \in \Omega} \text{Prob}(\omega | s_i) \sum_{s_{-i}} \text{Prob}(s_{-i} | \omega, s_i) \sum_{a_{-i} \in X_i} \text{Prob}(a_{-i} | s_{-i}, \beta_{-i}) u_i(a, \omega)$$

is strictly larger than the same expression where a_i is replaced by $\beta_i(s_i)$. Given the definition of X_i , either the collective decision is one (“pass”) when playing a_i against *any* $a_{-i} \in X_i$ or the collective decision is zero (“fail”) when playing a_i against *any* $a_{-i} \in X_i$.

Consider now the outcome-based game, and let \hat{X}_i be the set of ballot profiles for voters different from i such that k^* of them cast a 0-ballot and $n - k^* - 1$ cast a 1-ballot. A similar reasoning to the one developed for the premise-based game implies that voter i of type s_i picks an action $\hat{a}_i \in \{0, 1\}$ so as to maximize

$$\sum_{\omega \in \Omega} \text{Prob}(\omega | s_i) \sum_{s_{-i}} \text{Prob}(s_{-i} | \omega, s_i) \sum_{\hat{a}_{-i} \in \hat{X}_i} \text{Prob}(\hat{a}_{-i} | s_{-i}, \hat{\beta}_{-i}) u_i(\hat{a}, \omega),$$

where \hat{X}_i is the set of vectors $\hat{a}_{-i} \in \{0, 1\}^{I \setminus \{i\}}$ with exactly k^* null components.

By construction of β , it must be that $Prob(a_{-i}|s_{-i}, \beta_{-i}) = Prob(\hat{a}_{-i}|s_{-i}, \hat{\beta}_{-i})$, for each $a_{-i} \in X_i$, with $\hat{a}_{-i} \in \hat{X}_i$ given by $\hat{a}_j = 1$ when $a_j = \bar{\omega}$ and $\hat{a}_j = 0$ when $a_j = (0, \dots, 0)$. So, if a_i induces a positive collective decision (“pass”) whenever i is pivotal in the premise-based game, then voter i of type s_i would be better off in the outcome-based game by playing $\hat{a}_i = 1$ instead of $\hat{\beta}_i(s_i)$ when others play $\hat{\beta}_{-i}$. If a_i induces the zero collective decision (“fail”) whenever i is pivotal in the premise-based game, then voter i of type s_i would be better off in the outcome-based game by playing $\hat{a}_i = 0$ instead of $\hat{\beta}_i(s_i)$ when others play $\hat{\beta}_{-i}$. This contradicts the fact that $\hat{\beta}$ forms a BNE of the outcome-based game, and we have thus proved that β is a BNE of the premise-based game. ■

We emphasize that Proposition 1 is *not* a corollary of McLennan (1998). That paper showed that ex-ante efficient symmetric strategy profiles must be BNE in games with common interest. Since the strategy space is larger in the premise-based approach, it follows as a corollary that the most efficient BNE of the premise-based game is at least as efficient than the most efficient BNE of the outcome-based game. Yet common-interest games often admit BNEs other than the efficient symmetric one. Proposition 1 establishes that *any* symmetric BNE of the outcome-based game can be replicated by a symmetric BNE of the premise-based game.

The converse does not hold. There are robust situations where symmetric BNE outcomes of the premise-based game cannot be achieved by any symmetric BNE of the outcome-based game. The next example illustrates an even stronger point, namely that the maximal ex-ante welfare achievable by a symmetric BNE can be strictly larger in the premise-based than in the outcome-based game.

Example 1 Consider a problem with three premises, and three individuals. The society’s standard is to take a positive decision if and only if two of three premises are positive. Suppose that $q = 1/2$, meaning that the disutility from a false positive is equal to the disutility from a false negative. Let π be any prior such that $\pi_{101} = 0.4$ and $\pi_{100} = \pi_{001} = 0.15$. Suppose that conditional probabilities are such that $p_{101}(101) = 1/2$, $p_{101}(100) = 1/4$, and $p_{101}(001) =$

1/4, while $p_\omega(\omega) = 1$ for all $\omega \neq (1, 0, 1)$. We start by computing the best ex-ante welfare that can be achieved by a symmetric BNE in the outcome-based game. Given McLennan's (1998) result, it amounts to computing the largest ex-ante welfare that can be achieved via any symmetric strategy profile. Given the signal structure, it is easy to check that an optimal strategy profile will have $\hat{\beta}(s) = 0$ for all s such that $f(s) = 0$ and s is different from $(1, 0, 0)$ and $(0, 0, 1)$, and $\hat{\beta}(s) = 1$ for all s such that $f(s) = 1$. Symmetry then allows us to restrict attention without loss of generality to the case $\hat{\beta}(1, 0, 0) = \hat{\beta}(0, 0, 1)$. If α denotes this common number, then rather straightforward computations allow to show that the ex-ante welfare of an individual (same for all, given that they share a common interest) is equal to:

$$\frac{1}{4}\alpha^3 - \frac{9}{20}\alpha^2 + \frac{3}{20}\alpha - \frac{1}{10}.$$

This expression is maximized at $\alpha = 1/5$. One can check that the ex-ante welfare from truth-telling in the premise-based game is already strictly larger than the maximal welfare achievable in the outcome-based game (the exact difference is equal to 0.0235).

Observe that this example is robust in the sense that there exists an open set of conditional probabilities in $\Delta(\Omega)^\Omega$ around p such that the maximal ex-ante welfare associated to a symmetric BNE is strictly larger in the premise-based game than in the outcome-based game, for any of profile r of conditional probabilities in that set. For each strategy profile, the ex-ante welfare (in both games) is a linear function of the probabilities associated with each profile of votes conditional on the various states. These probabilities vary continuously with r , and hence the sequence of functions that determines the ex-ante welfare as a function of the strategy profile converges uniformly to the ex-ante welfare function associated with p as r approaches p . The maximal value of the ex-ante welfare (derived from choosing the optimal strategy profile) will thus converge to the maximal value at p when r approaches p , and it must thus be that the premise-based becomes strictly superior compared to the outcome-based whenever q is close enough to p .

4. LARGE GROUPS AND EFFICIENCY AT THE LIMIT

We start by proving a simple lemma that will allow us to derive a useful sufficient condition for asymptotic efficiency in the premise-based game.

Lemma 1 *Let $\theta \in [1/2, 1)$, and $(\alpha_l)_{l=1}^L, (\beta_l)_{l=1}^L$ be two vectors with components in $[0, 1]$. If $\min_l \alpha_l > \max_l \beta_l$, then there exists $(x, y) \in [0, 1] \times [0, 1]$ such that*

$$(\forall l = 1, \dots, L) : \alpha_l x + (1 - \alpha_l)y > \theta$$

$$(\forall l = 1, \dots, L) : \beta_l x + (1 - \beta_l)y < \theta$$

Proof: Let $\alpha = \min_l \alpha_l$ and $\beta = \max_l \beta_l$. Consider the vector

$$(x, y) = (\theta, \theta) + (\epsilon, -\frac{\epsilon}{2}[\frac{\alpha}{1-\alpha} + \frac{\beta}{1-\beta}]).$$

If ϵ is small enough, then $(x, y) \in [0, 1]^2$. Observe that $\alpha_l x + (1 - \alpha_l)y \geq \alpha x + (1 - \alpha)y$, for all l , because $x > y$ and $\alpha_l \geq \alpha$. In addition,

$$\alpha x + (1 - \alpha)y = \theta + \alpha\epsilon - (1 - \alpha)\frac{\epsilon}{2}[\frac{\alpha}{1-\alpha} + \frac{\beta}{1-\beta}] > \theta + \alpha\epsilon - (1 - \alpha)\frac{\epsilon}{2}[\frac{\alpha}{1-\alpha} + \frac{\alpha}{1-\alpha}] = \theta,$$

where the inequality follows from the fact that $\alpha > \beta$. Hence $\alpha_l x + (1 - \alpha_l)y > \theta$, as desired.

Similarly, observe that $\beta_l x + (1 - \beta_l)y \leq \beta x + (1 - \beta)y$, for all l , because $x > y$ and $\beta_l \leq \beta$. In addition,

$$\beta x + (1 - \beta)y = \theta + \beta\epsilon - (1 - \beta)\frac{\epsilon}{2}[\frac{\alpha}{1-\alpha} + \frac{\beta}{1-\beta}] < \theta + \beta\epsilon - (1 - \beta)\frac{\epsilon}{2}[\frac{\beta}{1-\beta} + \frac{\beta}{1-\beta}] = \theta,$$

where the inequality follows from the fact that $\alpha > \beta$. Hence $\beta_l x + (1 - \beta_l)y < \theta$, as desired. ■

The lemma is useful in the following way. Suppose that the objective is to implement some positive decision when the state falls in some subset A of Ω , and to implement the other decision when the state falls in a disjoint subset B of Ω . Suppose also that there is a set S of signals such that the probability

of getting a signal in S given any true state in A (cf. the α_i 's) is strictly larger than the probability of getting a signal in S given any true state in B (cf. the β_i 's). Then it is possible to find a “separating mixed strategy” - voting for the positive decision with probability x when the signal falls in S , and with probability y when the signals falls outside of S – such that the desired decisions are implemented for sure in a large population of independent voters who follow this strategy in a super-majority game with threshold θ . The next proposition applies this idea issue by issue.

Proposition 2 *Suppose that, for all $k \in \{1, \dots, K\}$, there exists a set $S_k \subseteq \{0, 1\}^K$ of signals such that the probability of receiving a signal in S_k at any state ω such that $\omega_k = 1$ is strictly larger than the probability of receiving a signal in S_k at any state ω such that $\omega_k = 0$:*

$$\min_{\omega \text{ s.t. } \omega_k=1} \sum_{s \in S_k} p_\omega(s) > \max_{\omega \text{ s.t. } \omega_k=0} \sum_{s \in S_k} p_\omega(s).$$

Then there exists a sequence of symmetric BNEs in the premise-based game such that the probability of error in each state goes down to zero as n grows.

Proof: Observe that McLennan’s (1998) result allows to limit ourselves to prove a weaker statement, namely that there exists a sequence of symmetric strategies (not necessarily in equilibrium) in the premise-based game such that the probability of error in each state decreases to zero as n grows. Indeed, any symmetric strategy profile will be dominated by another strategy profile that is ex-ante efficient, and which will thus be a symmetric BNE, as well as achieving efficiency at the limit.

Fix a premise k , and let (x, y) be the two numbers derived in the previous Lemma, with the various α ’s being equal to $\sum_{s \in S_k} p_\omega(s)$, with ω varying in $\{0, 1\}^K$ such that $\omega_k = 1$, and the β ’s being equal to $\sum_{s \in S_k} p_\omega(s)$, with ω varying in $\{0, 1\}^K$ such that $\omega_k = 0$. Let then $\hat{\beta}$ be a strategy such that each individual votes positively on premise k with probability x whenever the signal falls in S_k , and with probability y otherwise. The previous Lemma and the law of large numbers imply that the probability of having a number of

positive votes on premise k larger than θn converges to 1 at any state ω such that $\omega_k = 1$, and that the probability of having a number of positive votes on premise k smaller than θn converges to 1 at any state ω such that $\omega_k = 0$. In other words, the probability of error on any premise goes down to zero as n grows if these strategies are followed, and the probability of making the wrong decision also goes down to zero, a fortiori. ■

A similar technique would allow us to prove the following result for asymptotic efficiency in the outcome-based game: if there exists a set $S \subseteq \{0, 1\}^K$ of signals such that the probability of receiving a signal in S at any state ω with $f(\omega) = 1$ is strictly larger than the probability of receiving a signal in S at any state ω with $f(\omega) = 0$, then there exists a sequence of symmetric BNEs in the outcome-based game such that the probability of error in each state goes down to zero as n grows. Instead of proving this claim, we establish an even weaker sufficient condition for the outcome-based game. This new condition will allow us to derive several key results in this section and the next. From now on, let $\Omega := \{0, 1\}^K$ denote the set of states (or signals), let \mathcal{I} be the set of states ω such that $f(\omega) = 0$, and let $\mathcal{G} = \Omega \setminus \mathcal{I}$.

Proposition 3 *Fix $\theta \in [1/2, 1)$. Suppose that the collection of vectors $(p_\omega)_{\omega \in \Omega}$ describing the probability distribution of signals as a function of the true state satisfy the following condition:*

$$(\exists x \in \mathbb{R}^\Omega) : p_\omega \cdot x > 0, \text{ for all } \omega \in \mathcal{G} \text{ and } p_\omega \cdot x < 0, \text{ for all } \omega \in \mathcal{I}.$$

Then there exists a sequence of symmetric BNEs in the outcome-based game such that the probability of error in each state goes down to zero as n grows.

Proof: As in the proof of the previous Proposition, McLennan's (1998) result allows to limit ourselves to prove a weaker statement, namely that there exists a sequence of symmetric strategies (not necessarily in equilibrium) in the outcome-based game such that the probability of error in each state decreases to zero as n grows.

Notice that x can be taken as close as desired to the null vector in \mathbb{R}^Ω , as λx satisfies the same set of inequalities as x , for each $\lambda > 0$. Hence consider an x

small enough so that $(\theta, \dots, \theta) + x \in [0, 1]^\Omega$. Consider then the mixed strategy β where an individual votes positively on the final decision with probability $\theta + x(s)$ if he receives the signal s . Trivially, we have that $p_\omega \cdot \beta > \theta$, for all $\omega \in \mathcal{G}$, and $p_\omega \cdot \beta < \theta$, for all $\omega \in \mathcal{I}$. The law of large numbers implies that the probability of having a number of positive votes on the final decision larger than θn converges to 1 at any state $\omega \in \mathcal{G}$, and that the probability of having a number of positive votes on the final decision smaller than θn converges to 1 at any state $\omega \in \mathcal{I}$. In other words, the probability of making the wrong decision goes down to zero when n grows, as desired. ■

The stronger sufficient condition, stated right before Proposition 3, can now be derived as a corollary of that proposition and Lemma 1.

Corollary 2 *Suppose that there exists a set $S \subseteq \Omega$ of signals such that:*

$$\min_{\omega \in \mathcal{G}} \sum_{s \in S} p_\omega(s) > \max_{\omega \in \mathcal{I}} \sum_{s \in S} p_\omega(s).$$

Then there exists a sequence of symmetric BNEs in the outcome-based game such that the probability of error in each state goes down to zero as n grows.

Proof: By Lemma 1, there exists $(y, z) \in [0, 1]^2$ such that $y \sum_{s \in \mathcal{G}} p_\omega(s) + z \sum_{s \in \mathcal{I}} p_\omega(s)$ is strictly larger than θ , for all $\omega \in \mathcal{G}$, and strictly smaller than θ for all $\omega \in \mathcal{I}$. Hence the assumption of Proposition 3 is satisfied for x defined as follows: $x(s) = y - \theta$ if $s \in \mathcal{G}$ and $= z - \theta$ if $s \in \mathcal{I}$. ■

Another immediate corollary of Proposition 3 establishes that the outcome-based game is generically asymptotically efficient.

Corollary 3 *Fix $\theta \in [1/2, 1)$. Suppose that the collection of vectors $(p_\omega)_{\omega \in \{0,1\}^\kappa}$ describing the probability distribution of signals as a function of the true state are linearly independent. Then there exists a sequence of symmetric BNEs in the outcome-based game such that the probability of error in each state goes down to zero as n grows. In particular, asymptotic efficiency is guaranteed at almost all $p \in \Delta(\Omega)^\Omega$.*

Proof: Fix $\epsilon > 0$. Consider now the system of equations $p_\omega \cdot x = \epsilon$, for all $\omega \in \mathcal{G}$ and $p_\omega \cdot x = -\epsilon$, for all $\omega \in \mathcal{I}$. This system admits a solution x since the matrix $(p_\omega)_{\omega \in \{0,1\}^K}$ is invertible, and the previous proposition thus applies. ■

By Proposition 1, asymptotic efficiency is guaranteed in the premise-based game whenever it happens in its outcome-based variant. We conclude this section by providing an example showing that the converse does not hold. It also shows that the sufficient condition stated in Proposition 3 does not imply the one stated in Proposition 2. This comes in sharp contrast with Ahn and Oliveros (2011b) where asymptotic efficiency is shown to be achievable in a trial-by-trial mechanism if and only if it is achievable via combined voting. The difference in results stems from the fact that our outcome-based game is different from their split trial scenario, as we explained in the Introduction.

Example 2 Consider a problem with three premises – $K = 3$ – and the society’s standard being the simple majority – $f(\omega) = 1$ if and only if at least two premises are positive. Suppose that conditional probabilities are defined as follows: p_{010} picks $(1, 1, 1)$ with probability α and $(0, 0, 0)$ with probability $1 - \alpha$, p_{110} picks $(1, 1, 1)$ with probability β and $(0, 0, 0)$ with probability $1 - \beta$, p_{011} picks $(1, 1, 1)$ with probability γ and $(0, 0, 0)$ with probability $1 - \gamma$, p_{111} picks $(1, 1, 0)$, $(1, 0, 1)$ and $(0, 1, 1)$ with equal probability, while p_ω picks ω for sure for any other state ω . We claim that, whenever $\gamma > 0$ and $\max\{\gamma, 1/3\} < \alpha < \min\{\beta, 2/3\}$, there exists a sequence of symmetric strategy profiles in the premise-based game such that the probability of error converges to zero in every state, but not in the outcome-based game.

We begin by showing that Proposition 2 applies. Take S_1 as the set of all signals for which the first component is positive. The probability of getting a signal in S_1 is $2/3$ in $(1, 1, 1)$, β in $(1, 1, 0)$, and 1 in any other ω such that $\omega_1 = 1$. The probability of getting a signal in S_1 is α in $(0, 1, 0)$, γ in $(0, 1, 1)$, and 0 in any other ω such that $\omega_1 = 0$. It is thus true that the probability of getting a signal in S_1 is larger at any state ω for which $\omega_1 = 1$ than at any state ω' for which $\omega'_1 = 0$. Take S_2 as the set of all signals for which the second component is positive. It is easy to check that the minimal probability of getting a signal in S_2 , conditional on the second component of the state

being positive, is equal to $\min\{2/3, \alpha, \beta, \gamma\}$, while the probability of getting a signal in S_2 , conditional on the second component of the state being zero, is zero. Finally, take S_3 as the set of signals for which the third component is positive, minus $(1, 1, 1)$, plus $(0, 0, 0)$. The probability of getting a signal in S_3 is $2/3$ in $(1, 1, 1)$, $1 - \gamma$ in $(0, 1, 1)$, and 1 in any other ω such that $\omega_3 = 1$. The probability of getting a signal in S_3 is $1 - \alpha$ in $(0, 1, 0)$, $1 - \beta$ in $(1, 1, 0)$, and 0 in any other ω such that $\omega_1 = 0$. It is thus true that the probability of getting a signal in S_3 is larger at any state ω for which $\omega_1 = 1$ than at any state ω' for which $\omega'_1 = 0$.

We now show that it is impossible to reach asymptotic efficiency in the outcome-based game. Consider a group of size n . The distribution of the number of positive votes when the state is $(0, 1, 0)$ is the sum of n independent draws of a variable that picks 1 with probability $\alpha x_n(111) + (1 - \alpha)x_n(000)$ and 0 with the complementary probability, where x_n denotes the probability ascribed by an optimal strategy for a group of size n to send a positive message as a function of the signal. Similarly, the distribution of the number of positive votes when the state is $(1, 1, 0)$ is the sum of n independent draws of a variable that picks 1 with probability $\beta x_n(111) + (1 - \beta)x_n(000)$ and 0 with the complementary probability. Also, the distribution of the number of positive votes when the state is $(0, 1, 1)$ is the sum of n independent draws of a variable that picks 1 with probability $\gamma x_n(111) + (1 - \gamma)x_n(000)$ and 0 with the complementary probability. Given that α falls strictly in between γ and β , it must be that $\alpha x_n(111) + (1 - \alpha)x_n(000)$ is larger or equal to the minimum of $\beta x_n(111) + (1 - \beta)x_n(000)$ and $\gamma x_n(111) + (1 - \gamma)x_n(000)$. Hence, if the probability of having a super-majority of positive votes increases to 1 in both $(1, 1, 0)$ and $(0, 1, 1)$, then it also increases to 1 in $(0, 1, 0)$. This shows that asymptotic efficiency cannot be reached in the outcome-based game.

5. FURTHER COMPARISON WITH AHN AND OLIVEROS (2011b)

In this last section, we show that a natural analogue of Ahn and Oliveros' (2011b) sufficient condition for asymptotic efficiency to our outcome-based game is significantly stronger than the one derived in Proposition 3. Here is

the natural analogue of their condition in our framework:

$$\text{There exists } S_I \text{ and } S_G \subseteq \Omega \text{ such that } S_I \cap S_G = \emptyset \text{ and} \quad (1)$$

1. $(\forall \omega \in \mathcal{I}) : \sum_{s \in S_I} p_\omega(s) > \sum_{s \in S_G} p_\omega(s)$, and
2. $(\forall \omega \in \mathcal{G}) : \sum_{s \in S_G} p_\omega(s) > \sum_{s \in S_I} p_\omega(s)$.

In other words, it is strictly more likely to get a signal in S_G rather than a signal in S_I when the true state is associated with a positive decision, and vice versa when the optimal decision at the true state is zero.

First observe that condition (1) trivially implies the condition stated in Proposition 3, simply by taking $x(s) = 1$ if $s \in S_G$, -1 if $s \in S_I$, and 0 otherwise. On the other hand, condition (1) is significantly stronger than the condition appearing in Proposition 3. For instance, it does not allow to establish the generic asymptotic efficiency of the outcome-based game, as we did in Corollary 3. To see this, consider for instance the case of two premises, with a society standard where the positive decision is optimal if and only if both premises are positive. It is easy to check that condition (1) does not apply for any collection of conditional probabilities such that $p_{11}(11) < 1/2$ and $p_{00}(11) < 1/2$. Such a set clearly contains an open subset of $\Delta(\Omega)^\Omega$. Notice that such distributions are not unreasonable at all. It is quite possible, for instance, that signals on the two premises are drawn independently, and that the probability of getting a positive signal is systematically rather low (e.g. it is overall difficult to establish the evidence of a crime even if they are factually true). Asymptotic efficiency at the limit can be guaranteed, though, if the probability of getting a positive signal conditional on the premise being true is larger than the probability of receiving a positive signal conditional on the premise being false. Indeed, let α_k be the probability of receiving a positive signal when premise k is true, and β_k be the probability of receiving a positive signal when premise k is false. Corollary 2 implies that asymptotic efficiency is guaranteed in the outcome-based game whenever $\alpha_k > \beta_k$ (simply take $S = \{11\}$, since $\alpha_1\alpha_2 > \alpha_1\beta_2, \beta_1\alpha_2 > \beta_1\beta_2$). Yet both $p_{11}(11)$ and

$p_{11}(00)$ will be smaller than $1/2$ (and hence condition (1) does not apply) if $\alpha_1\alpha_2 < 1/2$.

References

- Ahn, D. S., and S. Oliveros**, 2011a. Combinatorial Voting. *Econometrica*, forthcoming.
- Ahn, D. S., and S. Oliveros**, 2011b. The Condorcet Jur(ies) Theorem. Mimeo, Berkeley University.
- Austen-Smith, D. and J. S. Banks**, 1996. Information Aggregation, Rationality, and the Condorcet Jury Theorem. *American Political Science Review* **90**, 34-45.
- Bovens, L., and W. Rabinowicz**, 2006. Democratic Answers to Complex Questions - an Epistemic Perspective. *Synthese* **150**, 131-153.
- Bozbay, I., F. Dietrich, and H. Peters**, 2011. Judgment Aggregation in Search for the Truth. Mimeo, LSE and Maastricht University.
- Brennan, G.**, 2001. Collective Coherence? *International Review of Law and Economics* **21**, 197-211.
- Dietrich, F., and C. List**, 2007. Strategy-Proof Judgment Aggregation. *Economics and Philosophy* **23**, 269-300.
- Feddersen, T. and W. Pesendorfer**, 1997. Voting Behavior and Information Aggregation in Elections With Private Information. *Econometrica* **65**, 1029-1058.
- Feddersen, T. and W. Pesendorfer**, 1998. Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts Under Strategic Voting. *American Political Science Review* **92**, 23-35.
- Kornhauser, L. A., and L. G. Sager**, 1986. Unpacking the Court. *Yale Law Journal* **96**, 82-117.
- List, C.** 2005. The Probability of Inconsistencies in Complex Collective Decisions. *Social Choice and Welfare* **24**, 3-32.

- List, C.**, 2006. The Discursive Dilemma and Public Reason. *Ethics* **116**, 362-402.
- List, C., and C. Puppe**, 2009. Judgement Aggregation. *Handbook of Rational and Social Choice*, Chapter 19, 457-483.
- McLennan, A.**, 1998. Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents. *American Political Science Review* **92**, 413-418.
- Persico, N.**, 2004. Committee Design with Endogenous Information. *Review of Economic Studies* **71**, 165–191.
- Pettit, P.**, 2001. Deliberative Democracy and the Discursive Dilemma. *Philosophical Issues* **11**, 268-299.