Premise-based versus outcome-based information aggregation ☆

Geoffroy de Clippel [a], Kfir Eliaz [b]

[a] Brown University, Department of Economics, Providence, RI, USA
[b] Eitan Berglas School of Economics, Tel Aviv University and Economics Department, University of Michigan, USA

A B S T R A C T

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.

© 2014 Elsevier Inc. All rights reserved.

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 pieced 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 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 super-majority 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 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).

---

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

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

[3] We 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.

[4] Our 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, 2012 and the references therein).