Synthese (2008) 163:1–24
Knowledge, Rationality & Action 1–24
DOI 10.1007/s11229-008-9306-x

Decision framing in judgment aggregation

Fabrizio Cariani · Marc Pauly · Josh Snyder

Received: 28 December 2006 / Accepted: 18 January 2008 / Published online: 9 February 2008

Abstract Judgment aggregation problems are language dependent in that they may be framed in different yet equivalent ways. We formalize this dependence via the notion of translation invariance, adopted from the philosophy of science, and we argue for the normative desirability of translation invariance. We characterize the class of translation invariant aggregation functions in the canonical judgment aggregation model, which requires collective judgments to be complete. Since there are reasonable translation invariant aggregation functions, our result can be viewed as a possibility theorem. At the same time, we show that translation invariance does have certain normatively undesirable consequences (e.g. failure of anonymity). We present a way of circumventing them by moving to a more general model of judgment aggregation, one that allows for incomplete collective judgments.

Keywords Social choice theory · Judgment aggregation · Translation · Language dependence

F. Cariani (✉)
Department of Philosophy, University of California, Berkeley,
314 Moses Hall, Berkeley, CA 94270-2390, USA
e-mail: fcariani@berkeley.edu

M. Pauly · J. Snyder
Stanford University, Stanford, CA, USA

M. Pauly
e-mail: marc.pauly@stanford.edu

J. Snyder
e-mail: joshsnyder@earthlink.net

1 Introduction

The goal of judgment aggregation is to form a number of individual judgments into a collective judgment. These judgments are usually expressed in a formal language, yet little thought has been given to date about whether choosing to formalize judgments in one way rather than another has any impact. In the philosophy of science literature, however, the ability to select different facts about the world as the atomic ones has played an important role. In this paper, we formalize and extend this insight from philosophy of science, bringing it to bear on the realm of judgment aggregation.

After illustrating the problem of different but expressively equivalent decision frames for a particular judgment aggregation problem, we introduce a relatively standard model of judgment aggregation (Sect. 3). The idea that a judgment aggregation procedure should not depend on the decision frame is defended and formalized in Sect. 4 using the notion of translation invariance. In Sect. 5, we characterize the class of aggregation functions that satisfy translation invariance (Theorem 1). As it turns out, these aggregation functions will generally fail to be anonymous (Corollary 1); furthermore, if they are non-dictatorial they will be manipulable (Corollary 2).

In spite of these negative corollaries, we view our characterization result in a positive light. As we illustrate in Sect. 6, there are reasonable aggregation functions that are translation invariant, even though they fail to satisfy anonymity. We will suggest that these negative results are not caused by translation invariance but rather by a weakness in the underlying judgment aggregation model. This model, which we call the canonical model due to its prevalence in the literature, requires collective judgments to be complete—i.e. the collective can never be undecided on any issue under consideration. In Sect. 6, we show that by giving up this requirement, translation invariance can be achieved together with anonymity, though translation invariance remains a significant constraint.

2 Example

We shall start by describing an example—a small tweak on a real-life situation. In 1991, the German parliament staged a debate on whether the parliament should move from Bonn to Berlin (Leininger 1992). Among the motions considered were A (the parliament should move to Berlin), and B (the seat of government should move to Berlin). Disregarding the actual voting procedure adopted by the council of the elders of the German parliament, let's imagine a three people committee faced with this decision:

Decision Frame 1: The committee wants to decide whether or not to move the parliament, the government or both. The views of committee members (we use the set $\{1, 2, 3\}$ to label them) are represented in the table below.

	A	B	$A \wedge B$
1	1	1	1
2	1	0	0
3	0	1	0

So for instance, the second committee member thinks the parliament but not the government should move. Based on this, we can assign to him an opinion on the conjunction $A \wedge B$. How should the committee decide on $A \wedge B$? The *conclusion-based procedure* would apply majority voting to $A \wedge B$, and since only a minority accepts the conjunction, the original question would be answered NO. Alternatively, the *premise-based procedure* would apply majority voting to the premises A and B individually. As it turns out, there is a majority for both A and B, and hence the conjunction should be accepted as well.

The divergence of these two procedures, known as the *doctrinal paradox* or the *discursive dilemma*, has led to an investigation of judgment aggregation more generally, starting with (List and Pettit 2002). Here, the idea is that we want to aggregate individual judgments into a group judgment. In the example above, each committee member endorses a set of judgments, represented by formulas of some logical language such as propositional logic. Thus, the first committee member endorses judgment set $\{A, B, A \wedge B\}$, the second member endorses $\{A, \neg B, \neg(A \wedge B)\}$ and the third member $\{\neg A, B, \neg(A \wedge B)\}$. The aim is to find acceptable aggregation procedures that produce consistent group judgments from consistent individual judgments. The discursive dilemma illustrates that proposition-wise majority voting fails to yield consistent group judgments in general, for the judgment set of the group would be $\{A, B, \neg(A \wedge B)\}$.

However, occurrence of the discursive dilemma here is dependent on the decision frame used. The reason for this is that the premise-based procedure is sensitive to the decision frame (Bovens and Rabinowicz 2004). To illustrate this, consider the following reformalization of the earlier decision problem.

Decision Frame 2: As a matter of fact, the German parliament investigated a further motion, namely that the *parliament and the government should not be geographically separated*. In the first decision frame, this would have been represented by $A \equiv B$.[1]

But another decision frame might have looked equally good. In the new frame, the basic motions they consider are whether to move the parliament (A), and whether parliament and government should be in the same city (B).[2] Of course, this new frame is intertranslatable with the previous one. So changing the frame but leaving the committee's views about the world unaltered, we have:

	A	B	$A \wedge B$
1	1	1	1
2	1	0	0
3	0	0	0

[1] This might seem inaccurate. After all, the truth of $\neg A \wedge \neg B$ does not guarantee that parliament and government will be in the same city. For example, the parliament might move to Dresden, while the government stays in Bonn. Our response: the assignment of sentences of the modeling language to claims in the informal case that we're trying to model depends partly on background facts that we can simply stipulate. In this case, we stipulate that each of the institution will have to be either in Berlin or in Bonn. It is *relative* to these additional facts that $A \equiv B$ correctly represents the proposition that parliament and government shall not be separated.

[2] Note that the extra-systematic interpretation of the symbol 'B' in the formal model is now different.

Note that, in both frames, $A \wedge B$ expresses the same proposition (namely, that both should be moved): this second decision frame is indeed equivalent to the first one.

The conclusion-based procedure still leads to a NO answer, just like in the first decision frame. However, there is no longer a divergence between the conclusion- and the premise-based procedures! Applying majority voting to the premises, we obtain a majority for A (as before), but no majority for B. Hence, under the premise-based procedure, the committee decision would again be NO. The voters' judgments are unchanged, but the paradox has vanished, since the formalization has changed.

In short, the decision frame matters. The general problem of sensitivity to the different ways of framing an issue has been studied under the rubric of *agenda manipulability*.[3] We will model the reframing of a decision here as a translation from one language to another—and in this paper, we will actually treat it as a translation from a language to itself, as in the example above (See Sect. 4.3 for details). In these translations, the crucial difference is simply that certain facts are expressible as atoms in one language, whereas those same facts are expressible only as complexes after translation.

We are concerned here to learn which judgment aggregation procedures are not affected by the choice of decision frame. Thus we ask: which aggregation functions are invariant under translation? And what (normative and formal) properties do these aggregation functions have?

To answer these questions, we set up a formal model for judgment aggregation, which we mostly derive from the literature on the subject, as well as a formal model for translations that can interact with it.

3 Introduction to judgment aggregation

3.1 Language and semantics

A *propositional language* L is generated from a finite set of atoms or loal variables L_0 using the standard Boolean connectives of conjunction ($\wedge$), negation ($\neg$), disjunction ($\vee$), etc. A *valuation* or *truth assignment* v: $L_0 \rightarrow \{0, 1\}$ assigns a truth value to each atomic sentence of the language, and we let V_L be the set of valuations for L. By the usual recursive definition, each formula of L is assigned a truth-value in a valuation.

For a formula $\varphi \in L$, we let $V_L(\varphi) = \{v \in V_L | v(\varphi) = 1\}$ denote the valuations satisfying φ. Similarly, for a set of formulas $\Phi \subseteq L$, we let $V_L(\Phi) = \{v \in V_L | \forall \varphi \in \Phi$: $v(\varphi) = 1\}$ denote the set of valuations satisfying all formulas in Φ. For a valuation v, we let $L(v) = \{\varphi \in L | v(\varphi) = 1\}$ denote the formulas satisfied by v. Similarly, for a set of valuations $X \subseteq V_L$, we let $L(X) = \{\varphi \in L | \forall v \in X$: $v(\varphi) = 1\}$ denote the set of formulas satisfied by all valuations $v \in X$. We call a set of valuations X *consistent* iff $|X| \geq 1$ and *complete* iff $|X| \leq 1$. A set of formulas Φ is *consistent/complete* iff $V_L(\Phi)$ is consistent/complete. Let Φ be a set of formulas of L, and θ a single formula: we write $\Phi \models \theta$ iff $V_L(\Phi) \subseteq V_L(\theta)$.

In contrast with most work on judgment aggregation, our formal framework will not include an explicit agenda set $A \subseteq L$ of formulas that form the topic of discussion.

[3] Dietrich (2006) provides a systematic study of agenda manipulability.

Thus, we take the agenda to consist of the whole language L. This will simplify the following definitions and results, though it is often an unrealistic assumption. However, most agendas provide enough constraints to extend judgments to the entire language given consistency and completeness (Why choose a formal language whose expressiveness exceeds what is required?). Also, we want to find aggregation procedures with good general properties, and so in particular any proposed aggregation procedure should function properly in cases where the agenda is the whole language.

3.2 Judgment aggregation procedures

An aggregation procedure takes individual judgments as inputs and returns as output a collective judgment. These judgments (both individual and collective) are usually treated syntactically, as sets of sentences. For this paper, however, it will be far more convenient to treat them semantically, as sets of valuations. In Sect. 3.3, we will situate our model in the judgment aggregation literature.

Let $N = \{1, 2, \ldots, n\}$ be the finite set of agents or voters, where $n \geq 2$. A *judgment aggregation procedure* (or *function*) $\mathcal{A}$: $(V_L)^n \rightarrow V_L$ maps n individual judgments (valuations) to a group judgment. An aggregation function input $\vec{v} = \langle v_1, \ldots, v_n \rangle$ is called a *profile*.

Some commonly discussed properties of aggregation procedures will play a role:

Definition 1 (*Anonymity*) Suppose f: $N \rightarrow N$ is a permutation. We use $\vec{f}(\langle v_1, \ldots, v_n \rangle)$ as shorthand for $\langle v_{f(1)}, \ldots, v_{f(n)} \rangle$. An aggregation function $\mathcal{A}$ is *anonymous* iff for any permutation f and for any profile $\vec{v}$, $\mathcal{A}(\vec{v}) = \mathcal{A}(\vec{f}(\vec{v}))$.

Definition 2 (*Dictatorship*) An aggregation function $\mathcal{A}$ is a dictatorship iff there is some $i \in N$ such that for all $v_1, \ldots, v_n$ we have $\mathcal{A}(v_1, \ldots, v_n) = v_i$.

Definition 3 (*Manipulability*) An aggregation function $\mathcal{A}$ is manipulable iff there exist some voter i, proposition φ, and profile $\langle v_1, \ldots, v_n \rangle$ such that $v_i(\varphi) \neq \mathcal{A}(v_1, \ldots, v_i, \ldots, v_n)(\varphi)$, but $v_i(\varphi) = \mathcal{A}(v_1, \ldots, v_i^*, \ldots, v_n)(\varphi)$ for some alternative valuation v_i^*.

Definition 4 (*Independence*) An aggregation function $\mathcal{A}$ is independent iff for all propositions φ there is a function f_φ: $\{0, 1\}^n \rightarrow \{0, 1\}$ such that for all $v_1, \ldots, v_n$, $\mathcal{A}(v_1, \ldots, v_n)(\varphi) = f_\varphi(v_1(\varphi), \ldots, v_n(\varphi))$.

3.3 Discussion of the canonical model

In our formal model, we have taken judgment aggregation procedures to map individual valuations to a group valuation. This model is usually presented in a syntactic version. Recall that a set of formulas is Φ consistent (complete) provided that $V_L(\Phi)$ is consistent (complete). Furthermore, call Φ *deductively closed* iff for every formula δ, $\delta \in \Phi$ whenever $V_L(\Phi) \subseteq V_L(\delta)$. Let $\mathcal{C} \subseteq 2^L$ be the collection of sets of propositions that are consistent, complete, and deductively closed. Then the syntactic analogue of a judgment aggregation function (as defined above semantically) is a function $\mathcal{A}^s$: $\mathcal{C}^n \rightarrow \mathcal{C}$ on complete, consistent, and deductively closed sets of formulas.

Any valuation v can then be treated as equivalent to the set of all formulas $L(v)$ that it satisfies, and any complete, consistent, and deductively closed set of formulas Φ can then be treated as equivalent to the valuation $V_L(\Phi)$ that satisfies it.[4] Note that without deductive closure, a syntactic aggregation function can distinguish, e.g. $\{p, q\}$ from $\{p, q, p \wedge q\}$, a distinction we cannot make semantically.

The syntactic definition of aggregation function further illustrates the assumptions that are built into our aggregation model: both the individuals and the group must be logically consistent and must have an opinion on every proposition. This idealized model has had a privileged, canonical status in the literature on judgment aggregation: it is adopted as a basic framework by List and Pettit (2002, 2004), Dietrich and List (unpublished, 2007), Dietrich (unpublished), Nehring (unpublished). However, this almost universal adoption need not be understood as reflecting widespread belief in the canonical model. On the contrary, many aggregation procedures studied in the literature fail to satisfy the assumptions built into the canonical model. List (2005) argues that the point of impossibility theorems in judgment aggregation is not to conclude that judgment aggregation is impossible, but to partition the logical space of aggregation procedures. It is in this connection that the canonical model has proven indispensable. Many of the impossibility results published so far depend on its adoption.

We will demonstrate that there is a tension between requiring robustness under switches of decision frame and some of the assumptions of the canonical model. Hence, if such robustness is desirable, we should move towards a more general model of judgment aggregation; we do just that in Sect. 6.

4 Introduction to translation invariance

4.1 History

The notion of translation invariance has its origins in the philosophy of science. We will sketch these origins before discussing the normative status of translation invariance and our formal model. For a thorough introduction and detailed references, the reader is referred to Oddie (2007).

A main goal of scientific theories, e.g. in physics, is empirical adequacy. Thus, so the standard story goes, Einstein's theory of relativity improves on Newtonian physics since it correctly predicts more observations. According to one philosophical stance, realism, this is because Einstein's theory is *closer to the truth* than Newton's theory.

At the same time, both theories are (likely to be) false, strictly speaking. Hence, one of the tasks of philosophy of science is to be able to compare different scientific theories according to *truthlikeness*. Given two scientific theories A and B, both potentially false, we would like criteria that tell us when A is closer to the truth than B. Note that we are concerned here with a semantic rather than an epistemic problem

[4] Later, when we relax completeness, we will still be able to give a semantic representation of syntactic aggregation functions, by representing a consistent and deductively closed set of formulas by the set of valuations that satisfy all those formulas.

(Zwart 1998): we are not interested in *how we can tell* that A is closer to the truth than B. We only want a definition of *what it means* to be closer to the truth.

Popper (1963) was the first to give a definition of truthlikeness, following what is now known as a *content approach*. For Popper, a theory A is at least as close to the truth as theory B iff the true consequences of B are included in those of A, and the false consequences of A are included in those of B. Hence, the truth content of B does not exceed that of A, and the falsity content of A does not exceed that of B. But while intuitively appealing, it was soon realized that this definition collapses since it entails that false theories are all incomparable (Tichy 1974; Miller 1974).

In reaction to this failure, some people have proposed a *likeness approach* which compares theories with the true theory according to an explicit similarity relation. This relation can be defined in many ways, but one measure that has been proposed is to average the number of atomic formulas on which models of a theory differ from the true theory. Against this kind of proposal, Miller (1974) argued that it violated translation invariance, since formalizing theories differently led to different truthlikeness values. It has become a standard objection in the philosophy of science.

4.2 Normative interest

So, what is the normative status of translation invariance as a condition in the judgment aggregation model? We see the normative status of translation invariance as similar to that of anonymity and neutrality, conditions which are usually desirable, but certainly not always. Similarly, translation invariance will often be one desirable property among many, and we have little to say about the difficult trade-offs one might be forced to make.

Is translation invariance normatively desirable in this sense? If the judgment aggregation procedure in use yields different results in different decision frames, then even before voting begins, there must be a way of determining how the decision is to be framed. In particular circumstances, the group may *agree* to pick a class of logically independent propositions to treat as premises in the judgment. However, first, there is no general guarantee that this should be the possible. If the group does not agree, and no one has the authority to fix a frame, then presumably there must be a debate: it is not clear, in abstract, what the substance of such a debate would be.[5] However, again, there is no reason to always expect this debate to settle on a fixed way of framing the problem at hand: distinct frames may be just as good by all criteria. Translation invariance applies, in our view, when a group lacks the ability or the opportunity to carry this debate through.

A second point, translation variance creates the possibility for strategic manipulation of results by careful selection of the framework. Such strategic battles are avoided when a translation invariant aggregation procedure is in use. In this sense, its normative

[5] It is clearly outside of the scope of our work to make sense of this question. One can imagine *pragmatic* sorts of argument here: perhaps, *sometimes* judgments about the relative priority of various issues can be obtained by getting clear about which issues are the most *urgent*.

desirability, is similar to that of other requirements of robustness under certain types of manipulation.

Third, there *are* real-life examples of collective deliberation in which alternative frames for a given problem sound just as good—witness our initial Bonn/Berlin example. In it, we can distinguish three decision frames that seem just as good as each other: $\{p, g\}$, $\{p, p \equiv g\}$, $\{g, p \equiv g\}$. In this example, there is no apparent reason to treat the 'complex' proposition such as $p \equiv g$ differently from the 'basic' propositions. It is in these cases that the choice of decision frame seems completely open, yet possibly decisive.

In the literature on truth-likeness, there have been two responses to the argument from translation variance:

(1) Translation variance is not a problem because there is a privileged language. This is a form of semantic realism that holds that certain properties are natural kinds, and hence these must be taken as the atoms in any theory. While this may be a viable position in metaphysics and philosophy of science, we think that it is not a strong response in the domain of judgment aggregation. We see no reason why there should be conceptually primary formulations of all decision problems.

(2) Translation variance is not a problem because the truthlikeness measure (or aggregation procedure) should change with the language. When the measure/procedure is appropriately indexed to the language, the problem can be avoided. Whether this is a plausible response depends on whether any specific indexed proposals can be motivated well. Note also that the concrete challenges of implementing such an aggregation system are formidable.

Finally, we want to address a specific concern raised by an anonymous referee: Every translation invariant aggregation function (in the formal sense that we shall define in the next section) satisfies a rather strong version of the unanimity principle, namely for all $p \in \mathcal{L}$ and $\vec{v}$, if [for all $v_i \in \vec{v}$ we have $v_i(p) = 1$] then $\mathcal{A}(\vec{v})(p) = 1$. The problem is that a strong unanimity requirement is often considered problematic (see for example Nehring unpublished). There are two kinds of objections here: (i) the unanimity principle is a factor in very strong impossibility results (Mongin unpublished; Nehring unpublished) (ii) the unanimity principle is unacceptable on intuitive grounds: a suitably gerrymandered disjunction may be supported by unanimous consensus, but the unanimity principle should intuitively not apply to that.

We will start from (ii). Nehring (unpublished) (Sect. 7) provides two arguments to the effect that violations of unanimity as applied to arbitrary propositions should look not only excusable, but actually desirable. Nehring's conclusions are not incompatible with our view. Each of Nehring's arguments is premised on the acceptability of representation of aggregation problems with logically independent propositions. Indeed, one of them is premised on the stronger condition that the propositions be, as Nehring says, *epistemically independent* (meaning that the reasons that affect the judgment on each premise have no bearing on the judgment on the other premises). We refer back to our comments at the beginning of this subsection: translation invariance applies when an aggregation problem should be represented by giving equal weight to alternatives that are mutually exclusive and exhaustive of logical space. This is a special case of judgment aggregation, but so is the case of epistemically independent premises. As far as we are aware, no one has given compelling reasons to suppose that every judgment aggregation problem will be representable by means of logically independent premises.

As for (i): Mongin (unpublished) shows that (given conditions on the agenda) there are no strongly unanimous, non-dictatorial aggregation functions that satisfy a weakening of independence (independence restricted to propositional variables: he calls this 'IIPA'). We think of this type of result as continuous in spirit to ours: unlike Mongin, we think that even the extremely thin IIPA, may, on occasion, be foregone. But, like Mongin (p. 16), we think that the proper response to increasingly simpler impossibility results is to weaken general framework assumptions. We make a step in this direction in Sect. 6, and defer the investigation of alternatives to future research.

4.3 Formal framework

Translations can be approached via a semantic and via a syntactic route. We shall develop both routes here, although the semantic one will be central to our results.

Definition 5 (*Syntactic Translation*) A syntactic translation is any map $\tau\colon L \to L$ that

(i) preserves the logical operations ($\tau(\alpha \wedge \beta) = \tau(\alpha) \wedge \tau(\beta)$, etc.), and
(ii) preserves entailment and satisfiability, i.e. for all φ and ψ, $\varphi \models \psi \Leftrightarrow \tau(\varphi) \models \tau(\psi)$[6]

By clause (i) syntactic translations are fully specified by their behavior on the atomic sentences L_0.

Note that we have defined translations as functions from L into L, not as functions from L to some L'. The reason for this move is mainly notational simplicity at no conceptual expense.[7]

Definition 6 (*Semantic Translation*) A quasi-translation is any map $\hat{\tau}\colon V_L \to 2^{V_L}$ mapping valuations to sets of valuations. $\hat{\tau}$ is a *semantic translations* iff $\hat{\tau}$ is a permutation of V_L.

For any function f, let $f[X] = \{y \mid \exists x \in X\colon f(x) = y\}$ be the image of X under f. Given that we have finitely many atoms, note that any set of valuations $X \subseteq V_L$ can be characterized completely by some formula φ_X, i.e. there is some formula φ_X such that $V_L(\varphi_X) = X$. We say that a syntactic translation relation τ and a semantic translation $\hat{\tau}$ *correspond* iff for all formulas $\varphi \in L$, we have $V_L(\tau(\varphi)) = \hat{\tau}[V_L(\varphi)]$. That is, the following diagram commutes:

6 We don't call these maps "syntactic" because they are defined using syntactic concepts only (they aren't, as clause (ii) shows). Rather the label refers to the fact that these maps are defined on syntactic objects (formulas of the language), in contrast with the associated semantic maps defined below.

7 Indeed, talking about translations across distinct formal languages only requires appending the appropriate subscripts.

$$\begin{array}{ccc} L & \xrightarrow{\tau(\cdot)} & L \\ \Big\downarrow {\scriptstyle V_L(\cdot)} & & \Big\downarrow {\scriptstyle V_L(\cdot)} \\ V_L & \xrightarrow{\hat{\tau}[\cdot]} & V_L \end{array}$$

If $\hat{\tau}$ is a quasi-translation, $\hat{\tau}(v)$ might fail to be a singleton set. If $|\hat{\tau}(v)| > 1$, this means that our target language is able to make distinctions our source language does not. On the other hand, if $\hat{\tau}$ is not one-to-one, e.g. $\hat{\tau}(v_1) = \hat{\tau}(v_2) = \{v'\}$, the source language makes distinctions the target language does not. However, neither situation is possible if $\hat{\tau}$ is a semantic translation.

It is easy to see that for every syntactic translation relation τ there is a corresponding quasi-translation $\hat{\tau} = V_L(\tau(\varphi_v))$. In fact, something stronger is true.

Proposition 1 (Representation) *For any syntactic translation τ, there is a unique corresponding semantic translation $\hat{\tau}$. Furthermore, for any semantic translation $\hat{\tau}$, there is a corresponding syntactic translation σ (and σ is unique up to logical equivalence).*

Proof See the Appendix. □

Given this correspondence, we will often identify semantic translations and syntactic translation, denoting both by τ (as in the following definition).

Now we are ready to define our central notion, that of a translation invariant judgment aggregation function. Let $\vec{\tau}(\vec{v})$ signify the profile resulting from distributing τ over $\vec{v}$. That is, if $\vec{v} = \langle v_1, \ldots, v_n \rangle$, then $\vec{\tau}(\vec{v}) = \langle \tau(v_1), \ldots, \tau(v_n) \rangle$.

Definition 7 (*Translation Invariance*) An aggregation function $\mathcal{A}$ is translation invariant iff for all translations τ and profiles $\vec{v}$, $\tau(\mathcal{A}(\vec{v})) = \mathcal{A}(\vec{\tau}(\vec{v}))$.

5 Results

We now turn to an investigation of the interaction between judgment aggregation and translation. We begin with a characterization theorem for translation invariant aggregation functions.

5.1 Characterization theorem

To present the characterization theorem, we introduce a number of additional notions. The first is that of a rolling dictatorship; an aggregation function is a rolling dictatorship when the collective judgment always equals one of the voters' judgments (though not necessarily always the same voter).

Definition 8 (*Rolling dictatorship*) An aggregation function $\mathcal{A}$ is a rolling dictatorship iff for all profiles $\langle v_1, \ldots, v_n \rangle$, there is an $i \in N$ such that $\mathcal{A}(v_1, \ldots, v_n) = v_i$.

Being a rolling dictatorship is a necessary condition for an aggregation function to be translation invariant, given a low bound on the number of atoms.

Lemma 1 *Given* $|L_0| \geq \log_2(n+2)$*, all translation invariant aggregation functions are rolling dictatorships.*

Proof Suppose $|L_0| \geq \log_2(n+2)$. We now prove the contrapositive. Suppose $\mathcal{A}$ is not a rolling dictatorship. Then there is some input profile $\vec{v}$ and some valuation $w \neq v_1, \ldots, v_n$ such that $\mathcal{A}(\vec{v}) = w$.

Define a translation τ as follows. Find an x such that $x \neq v_1, \ldots, v_n, w$. We are guaranteed that this is possible, since with $\log_2(n+2)$ atoms, we have $n+2$ distinct valuations. Let $\tau(y) = y$ for all $y \neq w, x$, and let $\tau(w) = x$ and let $\tau(x) = w$.

Now suppose, towards a contradiction, that $\mathcal{A}(\vec{\tau}(\vec{v})) = \tau(\mathcal{A}(\vec{v})) = \tau(w)$. But $\tau(v_i) = v_i$ for $i \in N$, so $\mathcal{A}(\tau(\vec{v})) = \mathcal{A}(\vec{v}) = w$. It follows that $\tau(w) = w$. But $\tau(w) = x$, and $x \neq w$ by assumption, a contradiction. □

The bound on the number of atoms is tight.

Lemma 2 *If* $|L_0| < \log_2(n+2)$*, then there exist translation invariant aggregation functions that are not rolling dictatorships.*

Proof Define an aggregation function as follows. Whenever there exists a unique $x \in V_L$ such that for all i, $x \neq v_i$, let $\mathcal{A}(\vec{v}) = x$. Otherwise, let $\mathcal{A}(\vec{v}) = v_1$. This function is translation invariant but not a rolling dictatorship. □

Central to the characterization theorem is the notion of a sectarian aggregation function. The idea is that a sectarian aggregation function pays attention only to how voters are distributed in clusters, or sects—determined by which voters have identical judgments—and then selects a winning sect. The collective judgment is then equal to the judgment of the winning sect. Observe that all sectarian aggregation functions are thus rolling dictatorships; this fact is used later in the proof of the characterization theorem.

We first formalize the notion of a sect in terms of partitions of N.

Definition 9 ($P_{\vec{v}}$) Given a profile $\vec{v} = \langle v_1, \ldots, v_n \rangle$, let $P_{\vec{v}}$ be the partition of N where i and j are in the same block iff $v_i = v_j$.

Lemma 3 $P_{\vec{v}} = P_{\vec{w}}$ *iff there is a translation* τ *such that for all i,* $\tau(v_i) = w_i$.

Proof ($\Rightarrow$) Suppose $P_{\vec{v}} = P_{\vec{w}}$. We can simply define the translation τ to be such that $\tau(v_i) = w_i$ for all i, and $\tau(w) = w$ for any $w \neq v_1, \ldots, v_n$. Such a τ is well defined because the fact that $P_{\vec{v}} = P_{\vec{w}}$ ensures us that for any i, j, $v_i = v_j$ iff $w_i = w_j$.

($\Leftarrow$) Suppose τ is a translation with the stated property. Then, i and j are in the same class in P_v iff $v_i = v_j$ iff $w_i = \tau(v_i) = \tau(v_j) = w_j$. □

We now define the notion of a sectarian aggregation function.

Definition 10 (*Sectarian*) Let $\mathbb{P}$ be the set of all partitions of N. An aggregation function $\mathcal{A}$ is *sectarian* iff there is a function $\mathcal{O}\colon \mathbb{P} \to N$ such that for all $\vec{v} = \langle v_1, \ldots, v_n \rangle$, we have $\mathcal{A}(\vec{v}) = v_{\mathcal{O}(P_{\vec{v}})}$.

We can now state the characterization theorem.

Theorem 1 *Given* $|L_0| \geq \log_2(n+2)$*, a judgment aggregation function is translation invariant iff it is sectarian.*

Proof Suppose that $|L_0| \geq \log_2(n+2)$.

($\Leftarrow$) Suppose there is some $\mathcal{O}$ such that for all input profiles $\vec{v}$, we have $\mathcal{A}(\vec{v}) = v_{\mathcal{O}(P_{\vec{v}})}$. Consider an arbitrary permutation τ and an arbitrary input $\vec{v}$. We want to show that $\tau(\mathcal{A}(\vec{v})) = \mathcal{A}(\vec{\tau}(\vec{v}))$.

But τ is one to one, so $\tau(v_i) = \tau(v_j)$ iff $v_i = v_j$. Thus $P_{\vec{v}} = P_{\tau(\vec{v})}$, and thus $\mathcal{O}(P_{\vec{v}}) = \mathcal{O}(P_{\tau(\vec{v})})$. Thus $\mathcal{A}(\vec{\tau}(\vec{v})) = \tau(v_{\mathcal{O}(P_{\tau(\vec{v})})}) = \tau(v_{\mathcal{O}(P_{\vec{v}})}) = \tau(\mathcal{A}(\vec{v}))$. Thus $\mathcal{A}$ is translation invariant.

($\Rightarrow$) Assume $\mathcal{A}$ is translation invariant. Consider any two profiles $\vec{v}$ and $\vec{w}$ such that $P_{\vec{v}} = P_{\vec{w}}$. We aim to show that there is some $i \in N$ such that $\mathcal{A}(\vec{v}) = v_i$ and $\mathcal{A}(\vec{w}) = w_i$. Furthermore, this i will depend only on the choice of $\vec{v}$. This will suffice to show that the function $\mathcal{O}$ can be found (in this case, $\mathcal{O}(P_{\vec{v}}) = i$).

Since $P_{\vec{v}} = P_{\vec{w}}$, by Lemma 3, there is some τ such that $\tau(\vec{v}) = \vec{w}$. Since $\mathcal{A}$ is a rolling dictatorship (Lemma 1), there is some $i \in N$ such that $\mathcal{A}(\vec{v}) = v_i$. Since $\mathcal{A}$ is translation invariant, $\tau(v_i) = \tau(\mathcal{A}(\vec{v})) = \mathcal{A}(\vec{\tau}(\vec{v})) = \mathcal{A}(\vec{w})$. But $\tau(v_i) = w_i$, so $\mathcal{A}(\vec{w}) = w_i$. □

Again, the bound on the number of atoms is tight; the proof uses the same construction as in Lemma 2.

5.2 Consequences of the characterization theorem

This characterization theorem provides for two interesting consequences. The first is that anonymity and translation invariance are at odds.

Corollary 1 *Given* $|L_0| \geq \log_2(n+2)$*, there are no anonymous translation invariant judgment aggregation functions.*

Proof Suppose $|L_0| \geq \log_2(n+2)$, and for a contradiction, suppose that $\mathcal{A}$ is translation invariant and anonymous. By Theorem 1, $\mathcal{A}$ is sectarian. Let $\vec{v}$ be a profile such that for all $i \neq j \in N$, $v_i \neq v_j$. There must be some $i \in N$ such that $\mathcal{A}(\vec{v}) = v_i$, and thus $\mathcal{O}(P_{\vec{v}}) = i$. Let f be a permutation of N such that $f(i) = j \neq i$.

Observe that $P_{\vec{v}} = P_{f(\vec{v})}$, so $\mathcal{A}(\vec{f}(\vec{v})) = v_{f(i)}$. However, $f(i) = j \neq i$, so $v_{f(i)} \neq v_i$. This contradicts $v_i = \mathcal{A}(\vec{v}) = \mathcal{A}(\vec{f}(\vec{v})) = v_{f(i)}$, which is given by anonymity. □

This is unfortunate, since both translation invariance and anonymity are in general very desirable. If we want to keep both anonymity and translation invariance, we must reject some other assumption that went into the framework, a route we shall explore in Sect. 6.

The second interesting consequence concerns manipulability.

Corollary 2 *Given* $|L_0| \geq \log_2(n+2)$*, translation invariant judgment aggregation functions are either dictatorships or manipulable.*

The proof of this relies on a weakening of a result of Dietrich (unpublished) and on a new theorem.

Theorem 2 (Dietrich unpublished) *If $\mathcal{A}$ is an aggregation function that fails independence, then $\mathcal{A}$ is manipulable.*

Theorem 3 *There are no independent, non-dictatorial, sectarian aggregation functions.*

Proof See the Appendix.

With these, we can prove Corollary 2.

Proof Suppose $\mathcal{A}$ is translation invariant; by Theorem 1, it is sectarian. Now suppose $\mathcal{A}$ is not a dictatorship. Then by Theorem 3, $\mathcal{A}$ is not independent. But then by Theorem 2, $\mathcal{A}$ is manipulable. □

Theorem 3 is essential in proving Corollary 2. However, it is also significant in its own right. Most of the impossibility results in the judgment aggregation literature depend on some kind of independence assumption. As Dietrich (unpublished) show, in the presence of independence assumptions, impossibility results cannot even be avoided by retreating to a model that allows incomplete outcomes. It is natural then to suggest that, if our goal is to avoid impossibility results, we should drop independence altogether. However, Theorem 3 shows that translation invariance and independence occupy (nearly) disjoint regions of logical space. As a consequence, insofar as Corollary 1 is read as an impossibility result, it is of an altogether different kind from traditional impossibility theorems.

Looking at Corollary 2, we face a trilemma: we must accept either a dictatorship, manipulability, or translation variance. Dictatorships are even worse than failures of anonymity and are in some sense at odds with the very notion of judgment *aggregation*. This leaves us with manipulability or translation variance, neither of which are attractive. Left to choose between them, they suggest an interesting general problem with strategic voting. We argued already that, when using a translation variant aggregation function, you open up the possibility of manipulability in a pre-debate about framing. On the other hand, we now see that if you prevent that debate by using a translation invariant aggregation function, you open up the possibility of manipulability in the voting stage. Either way, you end up with manipulability.

5.3 Aside: the bound

We saw that the characterization theorem and its consequences depended on a tight lower bound on the number of atoms. Does this make our results less interesting? No, for two reasons.

(1) The bound is *really low*, logarithmic in the number of voters, so low in fact that most examples found in the literature satisfy it.

(2) In the grand scheme of things, we think that we should concern ourselves with judgment aggregation *methods*—ways of aggregating judgments that are applicable in a variety of circumstances, and thus are independent of the numbers of agents and atoms. We can define a *judgment aggregation method* to be a function $\mathcal{M}\colon \mathbb{N} \times \mathbb{N} \to \mathbb{A}$ that maps the number of voters and the number of atoms

to an appropriate aggregation function. The notions applicable to aggregation functions can be lifted to judgment aggregation methods in a component-wise manner.[8] Any judgment aggregation method must select a judgment aggregation function for cases that satisfy the bound requirement, and our impossibility results will apply in these cases.

Thus the bound, while worth noting from a technical point of view, has little impact on the significance of our results.

6 Judgment aggregation without completeness

6.1 Possibility versus impossibility: a modified framework

While Theorem 1, our characterization theorem, may not be immediately interpreted as either a possibility or an impossibility theorem, Corollaries 1 and 2 seem to suggest the latter interpretation. However, we do consider our characterization to be a possibility theorem, since there are reasonable aggregation functions that are sectarian. The most natural example would be the plurality rule with tie breaker $\mathcal{A}_p$ defined as follows: Given input valuations $v_1, \ldots, v_n$, let $\mathcal{A}_p(v_1, \ldots, v_n)$ be that valuation which occurs most often among the input valuations, and if there is a tie, among the tied valuations, choose the valuation of the voter with the lowest index. For example, we will have $\mathcal{A}_p(v, v', v', v'') = v'$, and $\mathcal{A}_p(v, v', v'', v'', v') = v'$. This aggregation function is sectarian, but fails anonymity. As this example demonstrates, the failure of anonymity is not so much due to translation invariance, but rather due to the need to break ties. This need to break ties is built into the canonical model since it requires collective judgments to be complete.

The coherence of the canonical model has been under considerable pressure from impossibility results published in recent years, and we feel that the results obtained concerning translation invariance lend further support to move to a more general model of judgment aggregation. In search of a model that can accommodate translation invariance, two options suggest themselves naturally; we could allow judgment aggregation functions to be undefined on some profiles, or we could give up the requirement that the collective judgment set be complete[9] (the constraint of consistency is non-negotiable). The second option (giving up completeness) seems more intuitive, and has already been explored in other contexts (e.g. Pigozzi 2006) This broader framework already allows for reasonable translation invariant aggregation functions. We take one further step; once we allow the collective judgment to be incomplete, it is natural to also allow individuals this freedom.

We end up with a (*generalized*) *aggregation function* $\mathcal{A}$: $(2^{V_L})^n \to 2^{V_L}$ that maps individual sets of valuations to a collective set of valuations. Our earlier notion of

[8] That is to say by appropriately quantifying over the agents and atoms parameters. For example we could say that a method $\mathcal{M}$ is anonymous iff for all n and k, the aggregation procedure $\mathcal{M}(n, k)$ is anonymous. For manipulability, on the other hand, existential quantification over agents and atoms is more appropriate.

[9] List (2005) points out that groups sometimes do not have the option of not returning a complete collective judgment. However, the solution to their predicaments might well lie beyond the scope of judgment aggregation theory.

translation invariance can easily be lifted to this more general type of aggregation function. We say that $\mathcal{A}$ is *translation invariant* iff for all translations τ and all $X_1, \ldots, X_n \subseteq V_L$ we have $\tau[\mathcal{A}(X_1, \ldots, X_n)] = \mathcal{A}(\tau[X_1], \ldots, \tau[X_n])$.

With complete judgments, translation invariance was an almost impossible ideal; in this generalized aggregation framework, translation invariance is easier to come by. Consider for instance the aggregation function $\mathcal{A}_\cup$ defined as $\mathcal{A}_\cup(X_1, \ldots, X_n) = \bigcup_{i \leq n} X_i$. This function is easily seen to be translation invariant. In the case of the discursive dilemma, it simply fails to make a decision by returning the set of all the individuals' valuations.

A more interesting aggregation function is the following function $\mathcal{A}_{AV}$, which we may view as approval voting (Brams and Fishburn 1983) in the context of judgment aggregation: it returns as the group judgment the most popular of the voters' judgments. Formally, we have

$$\mathcal{A}_{AV}(X_1, \ldots, X_n) = \{v \in V_L \mid \forall v' \in V_L\colon |\{i|v \in X_i\}| \geq |\{i|v' \in X_i\}|\}$$

In the situation of the discursive dilemma, $\mathcal{A}_{AV}$ fails to make a decision and simply returns all individual valuations. Like in the case of $\mathcal{A}_\cup$, $\mathcal{A}_{AV}$ is easily seen to be translation invariant.

Note that these positive results concerning translation invariance also imply translation invariance for the special case where individuals are required to have complete judgment sets. Hence, if all we are interested in is to restore obtain translation invariance, it suffices to generalize the canonical model on the collective side.

In this context (completeness for inputs, but not for outputs), we can prove, given appropriate extensions of the definitions, an analogue of Theorem 1.

Definition 11 Let $\mathbb{P}$ be the set of all partitions of N. Say that $\mathcal{A}$ is generalized sectarian iff there is a function $\mathcal{O}\colon \mathbb{P} \mapsto \mathcal{P}(N)$ such that for all $\vec{v}$, $\mathcal{A}(\vec{v}) = \{w \in V_L \mid \exists i \in N, v_i = w$ and $i \in \mathcal{O}(P_{\vec{v}})\}$.

The analogue of Theorem 1 requires a condition that plays the role of the 'bound' in Theorem 1 (it turns out that in this context, the bound-assumption is not sufficient to carry through the proof of the analogue of Lemma 1).

Definition 12 A generalized aggregation function $\mathcal{A}$ is non-covering iff for every profile $\vec{v}$, there is a valuation z such that $z \neq v_i$, for all $i \in N$ and $z \notin \mathcal{A}(\vec{v})$.

In the original setting, every aggregation function above the bound is non-covering. We can now state the extended characterization result:

Theorem 4 *If the generalized aggregation function* $\mathcal{A}$ *is non-covering, then*: $\mathcal{A}$ *is* translation invariant *iff* $\mathcal{A}$ *is generalized sectarian.*

A proof is sketched in the appendix. Similar results also hold for the fully generalized context, in which completeness is relaxed for both inputs and outputs, although each such result requires its own set of definitions.

Theorem 4 reinforces the observation that translation invariance is a demanding constraint, even at the generalized level. However, it would be hasty to conclude that translation invariant generalized functions are never desirable. To repeat, verdicts of normative desirability typically depend on more information than is represented in the abstract judgment aggregation model. For example, enriching the informational environment of voting may affect the desirability of various procedures. These issues are discussed in Goodin and List (2006), in relation to May's theorem.[10]

There is another dimension of significance for this possibility result, towards which is important to gesture (if only in a rather programmatic way). The plurality rules that we have found to satisfy translation invariance might look implausible, because they are excessively 'holistic'. For one thing, they are insensitive to the similarities among valuations. However, 'holistic' aggregation rules, can be used as 'components' of less holistic ones. Suppose we start with an aggregation problem that involves a large boolean algebra. Given a proposition A, we might try to identify a sub-algebra containing those propositions *relevant* to A. The difficulty, of course, is in understanding 'relevance' so that given a proposition A, the propositions relevant to A do indeed form a boolean algebra. It might, then, be suggested that we apply a plurality rule to the relevant sub-algebra. This would not guarantee *full translation-invariance*, but a more local invariance (i.e. invariance relative to the relevant sub-algebra. More vividly, this is translation-invariance, subject to the constraint that a fragment of the original language is held fixed). Our result would be of significance–since it would show that, one way of ensuring a limited version of translation-invariance, is adopting a plurality rule to the judgments on the sub-algebra. This is clearly a topic for further research, and for a higher interaction of judgment aggregation with central areas of philosophical logic.

6.2 Translation invariance and distance measures

There is yet another respect in which translation-invariance remains a non-trivial constraint at the generalized level. It cuts against any aggregation function that depends on a non-trivial distance measure between judgments, such as the fusion procedure proposed by Pigozzi (2006). This is a consequence of Theorem 5 below.[11] Recall that a distance measure d on valuations is a function $d\colon V_L \times V_L \rightarrow \mathbb{R}$ satisfying the

[10] It is natural to ask: do the results concerning independence and manipulability (Corollary 2 and Theorem 3) hold in the general setting? In fact, the natural analogue of Corollary 2 fails in the generalized setting: there are translation invariant functions that are independent (in some appropriately generalized sense) and non-dictatorial. The union rule discussed in the main text is one such rule. In fact, the union rule is *systematic* in the sense that a function f determines whether $\mathcal{A}(\vec{X})(\varphi) = 1$ regardless of which proposition φ we is in question. The relevant function f maps each unanimous profile $\langle 1 \times \cdots \times 1 \rangle$ to 1, the unanimous profile $\langle 0 \times \cdots \times 0 \rangle$ to 0, and is undefined everywhere else (meaning that φ is neither accepted nor rejected).
On the other hand, by adding further conditions, it is possible to *characterize* what we called approval voting—an aggregation rule that violates independence. Goodin and List (2006) discuss a characterization theorem which can be reinterpreted in our notation as a characterization theorem for approval voting in the generalized setting.

[11] This result relating distance measures, valuations, and translation invariance is also relevant to the truthlikeness debate. Miller (1975) proves a somewhat related point.

following properties for all $x, y, z \in V_L$: (1) $d(x, y) \geq 0$, (2) $d(x, y) = 0$ iff $x = y$, (3) $d(x, y) = d(y, x)$, and (4) $d(x, y) + d(y, z) \geq d(x, z)$.

Definition 13 (*Distance function translation invariance*) A distance function d is invariant under translation iff for any τ, v and w, $d(v, w) = d(\tau(v), \tau(w))$.

Theorem 5 *A distance function d is invariant under translation iff there exists some* $c \geq 0$ *such that* $d(v, w) = c \cdot \bar{\delta}(v, w)$, *where* $\bar{\delta}(v, w) = \begin{cases} 0 \text{ if } v = w \\ 1 \text{ if } v \neq w \end{cases}$

Proof See the Appendix. □

In short, only trivial distance measures are translation invariant. And using trivial distance measures means that voters can only indicate which valuations they approve of and which they do not. Thus any judgment aggregation procedure that depends on a non-trivial distance measure will fail translation invariance.

7 Conclusion

Sensitivity to decision framing has been noted but not systematically studied in the literature on judgment aggregation. Representing switches of decision frame as switches of language, we were able to give precise content to the notion of an aggregation procedure being robust under changes of frame.

The immediate challenge then was to characterize which aggregation functions are translation invariant. Our exploration revealed that in the standard framework for judgment aggregation, while there are reasonable translation-invariant aggregation functions, translation invariance implies manipulability and violations of anonymity. In reaction to these results, we considered a broader framework that allowed incomplete judgments. Though translation invariance remained a non-trivial constraint, we found more breathing room there.

This is in many ways only an initial foray into this territory; there is considerable work still to be done. But we hope to have shown that the issue of language variance in judgment aggregation is significant, is subject to formal treatment, and should be taken seriously.

Acknowledgements We would like to thank Alexei Angelides, Johan van Benthem, Kenny Easwaran, Branden Fitelson, Itzhak Gilboa, Alistair Isaac, Martin van Hees, Christian List, Klaus Nehring, Gabriella Pigozzi, Mike Titelbaum, audiences at the Berkeley-Stanford logic workshop, LOFT 2006, the 2006 Cowles Workshop, the 2007 Pacific APA and two anonymous referees for their valuable comments and suggestions regarding this paper.

Appendix

Proof of Proposition 1

Recall that, for any valuation w, φ_w is a sentence of the language that is true exactly in w. Recall also that the core properties of syntactic translations are component-wise definition for complex statements and preservation of both deducibility and non-deducibility.

Proof of Part (i) Let τ be a syntactic translation. We want to show that $\hat{\tau}$ (defined canonically as: $\hat{\tau}(v) = V_L(\tau(\varphi_v))$) is a semantic translation. As pointed out in the text, it is easy to see that $\hat{\tau}$ must be a quasi-translation. We want to verify that it is *semantic translation*—i.e. that $\hat{\tau}$ (a) outputs only singletons (b) is 1–1 and (c) onto the set of singletons in 2^{V_L}.

A few preliminaries. Group the sentences of L in equivalence classes under the relation of logical equivalence. Write $[\varphi]$ for the equivalence class that φ belongs to, and let $\mathcal{E}$ be the set of equivalence classes for L. Note that since L_0 is finite, $\mathcal{E}$ is finite. The core properties of translations imply:

(1) $[\varphi] = [\psi]$ iff $[\tau(\varphi)] = [\tau(\psi)]$

Let $<_\mathcal{E}$ be the partial order on $\mathcal{E}$ defined by $[\varphi] <_\mathcal{E} [\psi]$ iff $\varphi \models \psi$. (1) allows us to define the function $h([\varphi]) = [\tau(\varphi)]$, and also to claim that h is a permutation of $\mathcal{E}$. (The proof below shows that h is an automorphism of the structure $\langle \mathcal{E}, <_\mathcal{E}\rangle$, but we do not assume this in the proof.)

Ad (a): Suppose for some v, $\hat{\tau}(v)$ is not a singleton. Then, since $\hat{\tau}(v) = V_L(\tau(\varphi_v))$:

Case (i) $V_L(\tau(\varphi_v))$ is empty, or
Case (ii) $|V_L(\tau(\varphi_v))| > 1$.

In case (i), φ_v is a consistent sentence of L that is mapped onto an inconsistent one, thus violating the core properties. In case (ii), let w be a valuation in $V_L(\tau(\varphi_v))$. Consider φ_w and its equivalence class $[\varphi_w]$. Since h is onto $\mathcal{E}$, $[\varphi_w]$ is in the range of h. Hence some sentence θ is both logically equivalent to φ_w and in the range of τ. Because $w \in V_L(\tau(\varphi_v))$, we now have $\varphi_w \models \tau(\varphi_v)$. So $\theta \models \tau(\varphi_v)$, because θ and φ_w are equivalent. However, $\tau(\varphi_v) \not\models \varphi_w$ and so $\tau(\varphi_v) \not\models \theta$. Since θ is in the range of τ, for some ψ, $\tau(\psi) = \theta$ (fix this ψ). Then the following hold:

(2) $\tau(\psi) \models \tau(\varphi_v)$
(3) $\psi \models \varphi_v$

(2) holds because $\tau(\psi) = \theta$, (3) by the properties of τ. Now ψ can only imply φ_v if it is a contradiction or if $[\psi] = [\varphi_v]$. However, ψ can't be a contradiction, because $\tau(\psi)$ is consistent. But neither can $[\psi] = [\varphi_v]$: if this were the case, $[\tau(\psi)] = [\tau(\varphi_v)]$, i.e. θ would be logically equivalent to $\tau(\varphi_v)$ which it is not, because θ is logically equivalent to φ_w and $\tau(\varphi_v) \not\models \varphi_w$.

Ad (b): Suppose $v \neq w$. Then (*) φ_v and φ_w are mutually inconsistent. Suppose (by reductio) that $\hat{\tau}(v) = \hat{\tau}(w)$. Then $[\varphi_{\hat{\tau}(v)}] = [\varphi_{\hat{\tau}(w)}]$. By the definition of $\hat{\tau}$ this means: $[\varphi_{V_L(\tau(\varphi_v))}] = [\varphi_{V_L(\tau(\varphi_v))}]$, which implies $[\tau(\varphi_v)] = [\tau(\varphi_w)]$ but this with (*) contradicts (1).

Ad (c): Suppose that there is a valuation v such that $v \notin Ran(\hat{\tau})$. Then for no valuation w, $v \in V[\tau(\varphi_v)]$. Let $\top$ be a tautology and index each valuation with a label from the set $\{1, \ldots, 2^n\}$. $\top \models (\varphi_{v_1} \vee \cdots \vee \ \varphi_{v_{2n}})$, but $\tau(\top) \not\models \tau(\varphi_{v_1}) \vee \cdots \vee \ \tau(\varphi_{v_{2n}})$, but this must be in violation of one of the two core properties of τ.

Uniqueness of $\hat{\tau}$: Suppose σ and $\hat{\tau}$ are semantic translations that correspond to τ (i.e. for all φ, $V_L(\tau(\varphi)) = \sigma[V_L(\varphi)]$). Let v be a valuation:

$$\hat{\tau}(v) = \hat{\tau}[V_L(\varphi_v)] = V_L(\tau(\varphi_v)) = \sigma[V_L(\varphi_v)] = \sigma(v)$$

This completes the proof of part (i). □

Proof of Part (ii) Let $\hat{\tau}$ be a semantic translation. We want to specify a canonical syntactic translation (which we call $\bar{\tau}$) associated with $\hat{\tau}$. Let $f: L_0 \mapsto L$ be a function that maps any atomic statement A of L to $\varphi_{\hat{\tau}[V_L(A)]}$. Now extend f component-wise to a function $\bar{\tau}: L \mapsto L$. Clearly, $\bar{\tau}$ obeys the first property of syntactic translations by construction. It is easy to see, by induction, that $\bar{\tau}$ has the "correspondence" property, viz.

$$(**) V_L(\bar{\tau}(\alpha)) = \hat{\tau}[V_L(\alpha)]$$

Now, suppose that $\hat{\tau}$ is a permutation of valuations. We want to show $\alpha \models \beta \Leftrightarrow \bar{\tau}(\alpha) \models \bar{\tau}(\beta)$. Recall that, $\alpha \models \beta$ iff $V_L(\alpha) \subseteq V_L(\beta)$. Now we have:

$$V_L(\alpha) \subseteq V_L(\beta) \Leftrightarrow \hat{\tau}[V_L(\alpha)] \subseteq \hat{\tau}[V_L(\beta)] \Leftrightarrow V_L(\bar{\tau}(\alpha)) \subseteq V_L(\bar{\tau}(\beta))$$

The first biconditional is justified by the fact that $\hat{\tau}$ is a permutation of valuations. The rightmost statement, together with (**), implies $\hat{\tau}(\alpha) \models \hat{\tau}(\beta)$. So $\bar{\tau}$ is a syntactic translation corresponding to $\hat{\tau}$.

Now $\bar{\tau}$ was defined as an extension of the initial f, and there are infinitely many different statements that would do the job of $\varphi_{\hat{\tau}(V_L[A])}$ in the definition of f. Call any such statement *admissible for A*. Notice that admissible statements must all belong to the same equivalence class (under logical equivalence) and must exhaust the class in question (anything in the equivalence class is admissible). It is easy to see using the definition of correspondence that any translation that corresponds to a semantic translation must make an admissible choice for each of the atomic statements. So $\bar{\tau}$ is determined up to logical equivalence and no finer than that. □

Proof of Theorem 3

To prove this, new definitions and terminology will be helpful. We first extend our partition notation, so that partitions can be induced by valuations v_i restricted to sets of propositions Γ.

Definition 14 ($P_{\vec{v},\Gamma}$) Given a profile $\vec{v} = \langle v_1, \ldots, v_n \rangle$ and a set of propositions Γ, let $P_{\vec{v},\Gamma}$ be the partition of N where i and j are in the same block iff $v_i(\varphi) = v_j(\varphi)$ for all $\varphi \in \Gamma$.

When Γ is a singleton set, we omit the braces. Observe that when $L_0 \subseteq \Gamma$, the two definitions coincide: $P_{\vec{v}} = P_{\vec{v},\Gamma}$.

Definition 15 (*Agreement*) Two profiles $\vec{v}$ and $\vec{w}$ agree on proposition φ iff for all $i \in N$, $v_i(\varphi) = w_i(\varphi)$.

Observe that if two profiles agree on a proposition, then an independent aggregation function will return the same value for the proposition for both profiles.

Definition 16 (*Winning sect*) $S \subseteq N$ is the winning sect on atomic proposition p in profile $\vec{v}$ iff

(1) $S \in P_{\vec{v},p}$ and
(2) for $i \in S$, $\mathcal{A}(\vec{v})(p) = v_i(p)$.

Condition (1) ensures that S is actually a sect at all (when attention is restricted to p). Condition (2) ensures that S is the winning sect—that the collective judgment on p is the same as the sect's judgment on p.

Lemma 4 *If*

(a) *a judgment aggregation function $\mathcal{A}$ is independent and sectarian, and*
(b) *there exists a profile $\vec{v}$, an atom p, and a voter i such that $\{i\}$ is the winning sect on p in $\vec{v}$*

then for any profile $\vec{w}$ and for any atom $q \neq p$, $\mathcal{A}(\vec{w})(q) = w_i(q)$.

Proof Suppose that $\mathcal{A}$ is sectarian and independent and that there exist $\vec{v}$, p, and i satisfying (b). Consider an arbitrary profile $\vec{w}$ and atom $q \neq p$.

Let $\vec{x}$ be a profile such that $\vec{x}$ agrees with $\vec{v}$ on p and $\vec{x}$ agrees with $\vec{w}$ on q. Since $\mathcal{A}$ is independent and $\vec{v}$ and $\vec{x}$ agree on p, $\mathcal{A}(\vec{x})(p) = \mathcal{A}(\vec{v})(p)$ and $v_i(p) = x_i(p)$. In conjunction with condition (b), this implies that $\mathcal{A}(\vec{x})(p) = x_i(p)$. Furthermore, the agreement of $\vec{v}$ and $\vec{x}$ on p ensures that for all $j \neq i$, $x_i(p) \neq x_j(p)$.

But since $\mathcal{A}$ is sectarian, it is a rolling dictatorship. In the case of $\vec{x}$, voter i must be the (rolling) dictator, since all other agents disagree with the collective on p. But since i is the (rolling) dictator, $\mathcal{A}(\vec{x})(q) = x_i(q)$. And since $\mathcal{A}$ is independent and $\vec{x}$ and $\vec{w}$ agree on q, $\mathcal{A}(\vec{x})(q) = \mathcal{A}(\vec{w})(q)$ and $x_i(q) = w_i(q)$. Thus $\mathcal{A}(\vec{w})(q) = w_i(q)$. □

Lemma 5 *If*

(a) *a judgment aggregation function $\mathcal{A}$ is independent and sectarian, and*
(b) *there exists a profile $\vec{v}$, an atom p, and a voter i such that $\{i\}$ is the winning sect on p in $\vec{v}$*

then $\mathcal{A}$ is a dictatorship.

Proof Suppose that $\mathcal{A}$ is sectarian and independent and that there exist $\vec{v}$, p, and i satisfying (b). Then by Lemma 4, for any profile $\vec{w}$ and for any atom $q \neq p$, $\mathcal{A}(\vec{w})(q) = w_i(q)$. Consider a profile $\vec{w}$ such that $\{i\} \in P_{\vec{w},q}$. Then $\{i\}$ is the winning sect on q in $\vec{w}$; applying Lemma 4 again, we obtain that for any profile $\vec{y}$ and for any atom $r \neq q$, $\mathcal{A}(\vec{y})(r) = y_i(r)$. But all atoms are different from at least one of p and q, since $p \neq q$. Therefore, for any profile $\vec{z}$ and any atom s, $\mathcal{A}(\vec{z})(s) = z_i(s)$. Thus $\mathcal{A}$ is a dictatorship. □

Lemma 6 *If*

(a) *a judgment aggregation function $\mathcal{A}$ is independent, sectarian, and non-dictatorial, and* (*b*) *there exists a profile* $\vec{v}$*, an atom* p*, and a sect* S *such that* $|S| > 1$ *and* S *is the winning sect on* p *in* $\vec{v}$

then there exists a profile $\vec{w}$*, an atom* q*, and a sect* T *such that* $|T| = |S| - 1$ *and* T *is the winning sect on* q *in* $\vec{w}$*.*

Proof Suppose that $\mathcal{A}$ is sectarian, independent, and non-dictatorial and there there exist $\vec{v}$, p, and S such that $|S| > 1$ and S is the winning sect on p in $\vec{v}$.

In what follows, we will refer extensively to three subsets of N, so we define them here. Let $M = \{min(S)\}$. Let $S^- = S - M$. Let $N^- = N - S$.

Construct a profile $\vec{w}$ as follows. Let $\vec{w}$ agree with $\vec{v}$ on p. Let $\vec{w}_i(q) = \vec{w}_i(p)$ for all $i \in S^-$ and let $\vec{w}_i(q) = 1 - \vec{w}_i(p)$ for all $i \notin S^-$. Let $\vec{w}_i(r) = 0$ for all i and for all $r \neq p, q$.

Construct also a profile $\vec{x}$ as follows. Let $\vec{x}$ agree with $\vec{v}$ on p. Let $\vec{x}_i(q) = \vec{x}_i(p)$ for all $i \in M$ and let $\vec{x}_i(q) = 1 - \vec{x}_i(p)$ for all $i \notin M$. Let $\vec{x}_i(r) = 0$ for all i and for all $r \neq p, q$.

Note that $P_{\vec{w}} = P_{\vec{x}} = \{M, S^-, N^-\}$. Since $\mathcal{A}$ is sectarian, there must be a rolling dictator for each of $\vec{x}$ and $\vec{w}$, drawn from one of these three sects. Furthermore, since $\mathcal{A}$ is sectarian, the rolling dictator must be the same for $\vec{x}$ and $\vec{w}$.

The rolling dictator cannot come from N^-, because (by independence) S is the winning sect on p in $\vec{x}$ and $\vec{w}$.

The rolling dictator also cannot come from M. If it did, then M would be the winning sect on q in $\vec{x}$. But then by Lemma 5, $\mathcal{A}$ would be a dictatorship, and this violates one of our assumptions.

Thus the rolling dictatorship must come from S^-, and thus S^- is the winning sect on q in $\vec{w}$. Since $|S^-| = |S| - 1$, we are done. □

We now turn to the proof of Proposition 3.

Proof (by contradiction) Suppose there were an independent, non-dictatorial, sectarian aggregation function $\mathcal{A}$, and consider an arbitrary profile $\vec{v}$ and atom p. There is some winning sect S on p in $\vec{v}$, of size $|S|$.

If $|S| = 1$, then by Lemma 5, $\mathcal{A}$ is a dictatorship, a contradiction.

If $|S| > 1$, we appeal to Lemma 6 $|S| - 1$ times. Each time we conclude that exists a profile $\vec{w}$ and an atom q such that the winning sect has size one smaller than on the previous appeal. We invariably reach a point where there is a winning singleton sect, thus yielding that $\mathcal{A}$ is a dictatorship, which is a contradiction. □

Proof of Theorem 4

The proof of this theorem mimics the proof of Theorem 1. Given a profile $\vec{v}$, we identify the set $S_{\vec{v}}$ of the valuations in that profile. We also let $C_{\vec{v}}$ be the set $V_L - S_{\vec{v}}$ of valuations in the language that do not occur in $\vec{v}$. Lemma 1, involved the notion of a *rolling dictatorship*. Lifted to the new level, this notion becomes:

Definition 17 A generalized aggregation function $\mathcal{A}$ is a rolling dictatorship* iff for all profiles $\vec{v}$, $\mathcal{A}(\vec{v}) \subseteq S_{\vec{v}}$.

Note again, in the special case in which $\mathcal{A}$ only outputs singletons, the notion of rolling dictatorship* is equivalent to the old notion of rolling dictatorship.

The equivalent of Lemma 1 says that if $\mathcal{A}$ is non-covering, and $\mathcal{A}$ is not a rolling dictatorship*, then it is not translation invariant. The proof of this lemma goes through just as the proof of Lemma 1. Suppose by reductio that $\mathcal{A}$ is translation invariant but not rolling*. Since $\mathcal{A}$ is not a rolling dictatorship* some w is both in $\mathcal{A}(\vec{v})$ and $C_{\vec{v}}$. Since $\mathcal{A}$ is *non-covering*, there is a valuation $z \notin \mathcal{A}(\vec{v})$ but $z \in C_{\vec{v}}$, hence $z \neq w$. Let τ be a translation that fixes every element of $S_{\vec{v}}$, but swaps z and w. Then we get:

1. $w \in \mathcal{A}(\vec{v})$.
2. $z \notin \mathcal{A}(\vec{v})$.
3. $\tau(w) \in \tau(\mathcal{A}(\vec{v}))$.
4. $\tau(w) = z$.
5. $\mathcal{A}(\vec{v}) = \mathcal{A}(\tau(\vec{v})) = \tau(\mathcal{A}(\vec{v}))$.
6. $z = \tau(w) \in \mathcal{A}(\vec{v})$.

(3) follows from (1). The first equality in (5) turns on the fact that τ fixes $\vec{v}$. The second depends on the translation invariance of $\mathcal{A}$. (6) follows from (3), (4) and (5). However, (6) contradicts (2).

We can now prove the theorem. The proof of the ($\Rightarrow$) direction is *mutatis mutandis* the same proof as in Theorem 1. For the ($\Leftarrow$) direction, assume $\mathcal{A}$ is translation invariant. Consider any two profiles $\vec{v}$ and $\vec{w}$ such that $P_{\vec{v}} = P_{\vec{w}}$. We aim to show that for any $i \in N$,

$$(*) v_i \in \mathcal{A}(\vec{v}) \Leftrightarrow w_i \in \mathcal{A}(\vec{w})$$

Let $D_{\vec{v}} = \{i \in N | v_i \in \mathcal{A}(\vec{v})\}$. Note that $D_{\vec{v}}$ depends only on the choice of $\vec{v}$. Thus we will be able to set for all $\vec{w} \in P_{\vec{v}}$, $\mathcal{A}(\vec{w}) = \{w_i | i \in D_{\vec{v}}\}$. This will suffice to show that the function $\mathcal{O}$ can be defined (in particular $\mathcal{O}(P_{\vec{v}}) = D_{\vec{v}}$).

Since $P_{\vec{v}} = P_{\vec{w}}$ (by Lemma 3) there is some τ such that $\vec{\tau}(\vec{v}) = \vec{w}$. By the analogue of Lemma 1, $\mathcal{A}$ is a rolling dictatorship*; hence there is some $D \subseteq N$ such that for all valuations x,

$$x \in \mathcal{A}(\vec{v}) \text{ iff there is an } i \in D \text{ such that } v_i = x$$

Clearly,

$$v_i \in \mathcal{A}(\vec{v}) \iff \tau(v_i) \in \tau(\mathcal{A}(\vec{v}))$$

Since $\mathcal{A}$ is translation invariant,

$$\tau(v_i) \in \tau(\mathcal{A}(\vec{v})) \iff \tau(v_i) \in \mathcal{A}(\vec{\tau}(\vec{v}))$$

But

$$\mathcal{A}(\vec{\tau}(\vec{v})) = \mathcal{A}(\vec{w})$$

and $\tau(v_i) = w_i$. The last identity, together with the previous two equivalences, establishes (*).

Proof of Theorem 5

($\Leftarrow$) Suppose there exists some $c \in [0, \infty)$ such that $d(v, w) = c \cdot \bar{\delta}(v, w)$, and suppose τ is an arbitrary permutation. Since τ is one to one, $\bar{\delta}(v, w) = \bar{\delta}(\tau(v), \tau(w))$. Therefore, $d(v, w) = c \cdot \bar{\delta}(v, w) = c \cdot \bar{\delta}(\tau(v), \tau(w)) = d(\tau(v), \tau(w))$.

($\Rightarrow$) We prove the contrapositive. Suppose there is no $c \in [0, \infty)$ such that $d(v, w) = c \cdot \bar{\delta}(v, w)$. We show that d is not invariant under translation.

First observe that d is not a constant function. If it were, then since $d(v, v) = 0$, we would have in general $d(v, w) = 0 = 0 \cdot \bar{\delta}(v, w)$.

Since $d(v, v) = d(w, w) = 0$, and since d is not constant, there must be $v \neq w$ such that $d(v, w) = e > 0$. Observe that $\bar{\delta}(v, w) = 1$, so in this case, $d(v, w) = e \cdot \bar{\delta}(v, w)$. But since there is no c such that in general $d(v, w) = c \cdot \bar{\delta}(v, w)$, there must be $x \neq y$ such that $d(x, y) = f$ where $f \neq e$. (Note that if $x = y$ then $d(x, y) = 0 = e \cdot \bar{\delta}(x, y)$, which is what x and y were supposed to rule out.)

We consider four cases.

(1) $v = x$ and $w = y$. This case does not arise. Observe that we could then obtain $e = d(v, w) = d(x, y) = f$, which contradicts $e \neq f$.
(2) $v = x$ and $w \neq y$. For all $z \neq w, y$, let $\tau(z) = z$. Let $\tau(w) = y$, $\tau(y) = w$. τ is well-defined, since $w \neq v, x$ and $y \neq x, v$. (In each of these, the first inequality is given by construction and the second is yielded by $v = x$.) Suppose, towards a contradiction, that d is invariant under translation. But then $e = d(v, w) = d(\tau(v), \tau(w)) = d(v, y) = d(x, y) = f$, which contradicts $e \neq f$. So d is not invariant under translation.
(3) $w = y$ and $v \neq x$. Similar to case (2).
(4) $v \neq x$ and $w \neq y$. For all $z \neq v, w, x, y$, let $\tau(z) = z$. Let $\tau(v) = x$, $\tau(x) = v$, $\tau(w) = y$, and $\tau(y) = w$. τ is well-defined, since v, w, x, and y are all different. Suppose, towards a contradiction, that d is invariant under translation. But then $e = d(v, w) = d(\tau(v), \tau(w)) = d(x, y) = f$, which contradicts $e \neq f$. So d is not invariant under translation.

Since the cases are exhaustive, we conclude d is not invariant under translation.

References

Bovens, L., & Rabinowicz, W. (2004). Democratic answers to complex questions—an epistemic perspective. In M. Sintonen (Ed.), *The socratic tradition—questioning as philosophy and as method*. Kluwer Academic Publishing.

Brams, S., & Fishburn, P. (1983). *Approval voting*. Birkhauser.

Dietrich, F. (2006). Judgment aggregation: (Im)possibility theorems. *Journal of Economic Theory, 126*(1), 286–298.

Dietrich, F., & List, C. (2007). Judgment aggregation by quota rules. *Journal of Theoretical Politics*, *19*(4), (in press).

Goodin, R., & List, C. (2006). A Conditional defense of plurality rule: Generalizing May's theorem in a restricted informational environment. *American Journal of Political Science*, *50*(4), 940–949.

Leininger, W. (1992). The fatal vote: Berlin versus Bonn, *Finanzarchiv, 50*, 1–20.

List, C. (2005). Group knowledge and group rationality: A judgment aggregation perspective. *Episteme, 2*(1), 25–38.

List, C., & Pettit, P. (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy, 18*, 89–110.

List, C., & Pettit, P. (2004). Aggregating sets of judgments: Two impossibility results compared. *Synthese, 140*(1–2), 207–235.
Miller, D. (1974). Popper's qualitative theory of Verisimilitude. *The British Journal of the Philosophy of Science, 25*, 166–177.
Miller, D. (1975). The accuracy of predictions. *Synthese, 30*, 159–191.
Oddie, G. (2007). Truthlikeness. Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2007 Edition), http://plato.stanford.edu/archives/sum2007/entries/truthlikeness/.
Pigozzi, G. (2006). Belief merging and the discursive dilemma: An argument-based account to paradoxes of judgment aggregation. *Synthese*, *152*(2), 285–298.
Popper, K. (1963). *Conjectures and refutations*. London: Routledge.
Tichy, P. (1974). On Poppers's definitions of verisimillitude. *The British Journal of the Philosophy of Science, 25*, 155–160.
Zwart, S. (1998). Approaches to truth. Ph.D. thesis, ILLC Amsterdam.