Decision Framing in Judgment Aggregation

Fabrizio Cariani   Marc Pauly   Josh Snyder

Draft – August 4, 2006

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 its normative desirability. We show that in the canonical judgment aggregation model, no reasonable judgment aggregation functions are translation invariant. These results motivate a more general model of judgment aggregation, which gives up the requirement that judgment sets be complete.

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 (section 3). The idea that a judgment aggregation procedure should not depend on the decision frame is defended and formalized in section 4 using the notion of translation invariance. Unfortunately, section 5 shows that only very undesirable aggregation functions satisfy translation invariance. These aggregation functions will generally fail to be anonymous; furthermore, if they are non-dictatorial they will be manipulable.

In our view, these negative results stem from 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 section 6, we show that by giving up this requirement, translation invariance can be achieved by reasonable aggregation procedures (though translation invariance remains a significant constraint).

2 Example

Consider a committee of three people that is faced with the following situation. A married couple Anne (A) and Bill (B) applied for a job. The committee needs to decide who to hire, and it can decide to hire both applicants. It is common knowledge that the applicants will accept the job only if it is offered to both.

Decision Frame 1: The committee wants to decide whether or not to hire both applicants, and the individual committee members $\{1, 2, 3\}$ have the views 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 Anne should be hired but not Bill, and hence believes that not both should be hired. How should the committee decide whether or not to hire both candidates? The *conclusion-based procedure* would apply majority voting to $A \wedge B$, and since only a minority accepts the conjunction, the original question (whether to hire both candidates) 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: The committee still wants to decide whether to hire both applicants, but instead of considering whether to hire Anne and whether to hire Bill, they consider whether to hire Anne (A), and whether Anne and Bill are equally qualified (B).[1] These new atoms A and B can be expressed in the language of the previous frame as A and $A \equiv B$, respectively. So changing the language 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

Note that in both frames, $A \wedge B$ express the same fact (that both should be hired), so 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. We will model the reframing of a

[1] It is a simple sort of job, for which you are either qualified or not; there are no situations in which two candidates can be qualified, but one more qualified than the other.

decision as a translation from one language to another–and in this paper, for simplicity, we will actually treat it as a translation from a language to itself (as in the example above). 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 propositional variables L_0 using the standard Boolean connectives of conjunction ($\wedge$), negation ($\neg$), disjunction ($\vee$), etc. A *valuation* or *truth assignment* $v : L_0 \to \{0,1\}$ assigns a truth value to each atom of the language, and we let V_L be the set of valuations for L.

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.

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. 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 section 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*) $\mathcal{A} : (V_L)^n \to V_L$ 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 $f : N \to N$ 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 $f_\varphi : \{0, 1\}^n \to \{0, 1\}$ such that for all $v_1, \ldots, v_n$, $\mathcal{A}(v_1, ..., v_n) = f_\varphi(v_1(\varphi), ..., v_n(\varphi))$.*

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 $\mathcal{A}^s : \mathcal{C}^n \to \mathcal{C}$ 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.[2] 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 (2004, 2005), Dietrich (2004), Nehring and Puppe (2005) and Pauly and van Hees (2003). 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 conflict 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 section 6.

[2] 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.

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 (2001).

A main goal of scientific theories, e.g. in physics, is to be empirically adequate. 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 (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 importance

Here are two reasons why translation invariance is an important feature of a judgment aggregation function:

(1) A translation variant aggregation function seems to fail to latch onto a sufficiently objective notion of aggregation. In principle, we should get the same aggregated judgments *regardless* of the expressive means we happen to adopt. That certain sentences count as atoms in one language while they can only be expressed by means of complexes in another ought to be of no consequence. What normatively significant considerations could adjudicate between different expressively equivalent frames?

(2) There is also a concrete problem with using translation variant aggregation functions. If the judgment aggregation procedure in use yields different results in different decision frames, then even before voting begins, there must be a preliminary debate over how the decision is to be framed. This raises two concerns. First, it is not clear what the substance of such a debate would be. Second, this situation creates the possibility for strategic manipulation of results by careful selection of the framework. Such substanceless strategic battles are avoided when a translation invariant aggregation procedure is in use.

Reason (1) above is also the reason that translation invariance is desirable in an account of truthlikeness: translation variance is an indicator that an account has failed to latch on to objective, language-independent aspects of the world. There have been two responses to this argument in the philosophy of science literature:

(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 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.

In light of the arguments outlined above, we feel that translation invariance is a normatively desirable property that is even more defensible in judgment aggregation than in the philosophy of science, where it originated.

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* translation *is any map $\tau : L \to L$ that preserves the logical operations, i.e. $\tau(\alpha \wedge \beta) = \tau(\alpha) \wedge \tau(\beta)$, $\tau(\neg\varphi) = \neg\tau(\varphi)$, etc. Consequently, a translation is fully specified by its behavior on the atoms L_0.*

Note that we have defined translations as functions from L into L, not as a function from L to some L'. The reason for this move is mainly notational simplicity at no conceptual expense.[3] One consequence of this definition is that if two formulas φ and ψ are logically equivalent, i.e. $V_L(\varphi) = V_L(\psi)$, then their translations must also be logically equivalent, i.e. for any translation τ, $V_L(\tau(\varphi)) = V_L(\tau(\psi))$. The converse does not necessarily hold.

Definition 6 (Semantic Translation) *A* translation *is any map $\hat{\tau} : V_L \to 2^{V_L}$ mapping valuations to sets of valuations.*

For any function f, let $f[X] = \{y \mid \exists x \in X : f(x) = y\}$ 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

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

diagram commutes:

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

It is easy to see that for every syntactic translation relation there is a corresponding semantic translation relation and vice versa. Given a syntactic translation relation τ, define $\hat{\tau}(v) = V_L(\tau(\varphi_v))$. Conversely, given a semantic translation relation $\hat{\tau}$, define $\tau(\varphi) = \varphi_{\hat{\tau}[V_L(\varphi)]}$. Given this correspondence, when no confusion can arise, we will often identify semantic and syntactic translation, denoting both by τ.

Note that in general, $\tau(v)$ might not be a singleton set. If $|\tau(v)| > 1$, this means that our target language is able to make distinctions our source language does not. On the other hand, if τ is not one-to-one, e.g. $\tau(v_1) = \tau(v_2) = \{v'\}$, the source language makes distinctions the target language does not. We are interested in investigating translations between expressively equivalent decision frames. For this reason, we will postulate here that translation functions are all injective functions that map a valuation to a singleton set of valuations. As a consequence, such a translation function can be represented as a bijection or permutation $\tau : V_L \to V_L$. Hence, from now on, a translation function shall denote such a bijection on V_L.[4]

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.

[4]Note that this means that not every syntactically specified translation function will correspond to such a bijection. This is yet another reason for working with the semantic notion of translation.

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 $\vec{v} = \langle v_1, ..., v_n \rangle$, 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 $\mathcal{O} : \mathbb{P} \to N$ 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 $\mathcal{A}(\vec{w}) = w_i$. $\square$

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. $\square$

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 section 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 and List (2004) and on a new theorem.

Theorem 2 (Dietrich and List (2004)) *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 Appendix 1. □

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. □

We face a trilemma: one 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 $\mathcal{M} : \mathbb{N} \times \mathbb{N} \to \mathbb{A}$ 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.[5] Any judgment aggregation method must select a judgment

[5] That is to say by appropriately quantifying over the agents and atoms parameters.

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 Modified framework

We obtained a number of negative results regarding the possibility of translation invariant judgment aggregation within the canonical model. The plausibility and simplicity of the assumptions that went into the results provides a reason for re-examining the framework. The coherence of the canonical model has been under considerable pressure from impossibility results published in recent years, and we feel it is time to look to alternatives.

In search of a model that can accommodate translation invariance, it is attractive to give up the requirement that the collective judgment set be complete; the constraint of consistency is non-negotiable.[6] 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 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 $X_1, \dots, X_n \subseteq V_L$ we have $\tau[\mathcal{A}(X_1, \dots, X_n)] = \mathcal{A}(\tau[X_1], \dots, \tau[X_n])$.

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 $\mathcal{A}_\cup(X_1, \dots, X_n) = \bigcup_{i \leq n} X_i$. 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.

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.

[6] 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.

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

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

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.

6.2 Translation invariance and distance measures

Even though translation invariance is now achievable in this generalized aggregation framework, it remains a non-trivial constraint. For example, it cuts against any aggregation function that depends on a non-trivial distance measure between judgments, such as the fusion procedure proposed by Pigozzi (2005). This is a consequence of Theorem 4 below.[7] Recall that a distance measure d on valuations is a function $d : V_L \times V_L \to \mathbb{R}$ satisfying the 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 11 (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 4 *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 Appendix 2. □

In short, only trivial distance measures are translation invariant. And using trivial distance measures means that voters can only indicate which

[7]This result relating distance measures, valuations, and translation invariance is also relevant to the truthlikeness debate. It strikes us as likely that it is not new with us, but we have been unable to find a reference.

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, translation invariance is an impossible ideal. In reaction to that result, we considered a broader framework that allowed incomplete judgments. Though translation invariance remained a non-trivial constraint, we found some 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.

References

BOVENS, Luc and RABINOWICZ, Wlodek (2004), 'Democratic Answers to Complex Questions – An Epistemic Perspective', in Sintonen M. (ed.), *The Socratic Tradiction – Questioning as Philosophy and as Method* (Kluwer Academic Publishing).

BRAMS, Steven and FISHBURN, Peter (1983), *Approval Voting* (Birkhauser).

DIETRICH, Franz (2004), 'Judgment Aggregation in General Logics', *(unpublished)*.

DIETRICH, Franz and LIST, Christian (2004), 'Strategy Proof Judgment Aggregation', *(unpublished)*.

——— (2005), 'Judgment Aggregation by Quota Rules', *unpublished*.

LIST, Christian (2005), 'Group Knowledge and Group Rationality: a Judgment Aggregation Perspective', *Episteme*, 2(1), 25–38.

List, Christian and Pettit, Philip (2002), ‘Aggregating Sets of Judgments: An Impossibility Result’, *Economics and Philosophy*, 18, 89–110.

——— (2004), ‘Aggregating Sets of Judgments: Two Impossibility Results Compared.’, *Synthese*, 140(1-2), 207–235.

Miller, David (1974), ‘Popper’s Qualitative Theory of Verisimilitude’, *The British Journal of the Philosophy of Science*, 25, 166–177.

Nehring, Klaus and Puppe, Clemens (2005), ‘Consistent Judgment Aggregation: A Characterization’, *(unpublished)*.

Oddie, Graham (2001), ‘Truthlikeness’, in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. URL `<http://plato.stanford.edu/archives/fall2001/entries/truthlikeness/>`.

Pauly, Marc and van Hees, Martin (2003), ‘Logical Constraints on Judgment Aggregation’, *Journal of Philosophical Logic*.

Pigozzi, Gabriella (2005), ‘Judgment Aggregation without Paradox’, *(forthcoming)*.

Popper, Karl (1963), *Conjectures and Refutations* (Routledge (London)).

Tichy, Pavel (1974), ‘On Poppers’s Definitions of verisimillitude’, *The British Journal of the Philosophy of Science*, 25, 155–160.

Zwart, Sjoerd (1998), *Approaches to Truth*, Ph.D. thesis, ILLC Amsterdam.

Appendix 1: 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 12 ($P_{\vec{v},\Gamma}$) *Given a profile* $\vec{v} = \langle v_1, ..., v_n \rangle$ *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 13 (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 14 (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. 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 use the 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 the previous application. We invariably reach a point where there is a winning singleton sect, thus yielding that $\mathcal{A}$ is a dictatorship, which is a contradiction. □

Appendix 2: Proof of Theorem 4

($\Leftarrow$) 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 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.