INDIVIDUAL COHERENCE AND GROUP COHERENCE

RACHAEL BRIGGS, FABRIZIO CARIANI, KENNY EASWARAN, BRANDEN FITELSON

ABSTRACT. Paradoxes of individual coherence (*e.g.*, the preface paradox for individual judgment) and group coherence (*e.g.*, the doctrinal paradox for judgment aggregation) typically presuppose that deductive consistency is a coherence requirement for both individual and group judgment. In this paper, we introduce a new coherence requirement for (individual) full belief, and we explain how this new approach to individual coherence leads to an amelioration of the traditional paradoxes. In particular, we explain why our new coherence requirement gets around the standard doctrinal paradox. However, we also prove a new impossibility result, which reveals that (more complex) varieties of the doctrinal paradox can arise even for our new notion of coherence.

1. INDIVIDUAL COHERENCE[1]

1.1. **Deductive Consistency: The Recent Dialectic.** It is often assumed that an epistemically rational agent's (full) beliefs ought to be deductively consistent. That is, the following is often taken to be a (synchronic) epistemic coherence requirement for individual agents.[2]

(CB) **Consistency Norm for Belief**. Epistemically rational agents should (at any given time) have logically consistent belief sets.

One popular motivation for imposing such a requirement is the presupposition that epistemically rational agents should, in fact, obey the following norm:

Date: Septmeber 16, 2013. Draft: final version to appear in *Essays in Collective Epistemology*, J. Lackey (*ed.*), Oxford University Press.

Key words and phrases. judgment aggregation, coherence, consistency, doctrinal paradox, preface paradox.

We have discussed various aspects of this project with a large number of individuals and groups over the last few years. While we cannot list all of these helpful individuals and groups here, we would like to single a few of them out. Thanks to an invitation by Jason Konek, we presented this material at the Philosophy Department at the University of Michigan, where we received very valuable feedback on the project. Nick Leonard, Gabriella Pigozzi, and an anonymous referee read various versions of this paper and gave us helpful notes. Finally, we'd like to thank Jennifer Lackey for the invitation which prompted this particular collaboration.

[1]This section is an abridged version of a much longer story we have written about individual coherence (Easwaran & Fitelson 2013). In that longer paper, we layout the philosophical and formal framework in greater detail. Specifically, see that longer paper for a detailed discussion of the various idealizations that are involved in the present framework.

[2]Notable advocates of (CB) include Pollock (1990), Ryan (1991, 1996), and Kaplan (2013). Christensen (2004), Kolodny (2007), Foley (1992), Klein (1985), and Kyburg (1970) all reject (CB). See (Easwaran & Fitelson 2013) for an extended discussion of the recent dialectic concerning (CB), as well as a more in-depth discussion of the present alternative(s) to (CB).

(TB) **Truth Norm for Belief**. Epistemically rational agents should (at any given time) believe propositions that are true.

These two norms differ in one fundamental respect: (TB) is *local* in the sense that an agent complies with it only if each particular belief the agent holds (at a given time) has some property (in this case: truth). On the other hand, (CB) is a *global* norm: whether or not an agent's doxastic state (at a given time) is in accordance with (CB) is a more holistic matter, which trades essentially on properties of their entire belief set. While these two epistemic norms differ in this respect, they are also intimately related, logically. We may say that one norm n *entails* another norm n' just in case everything that is permissible according to n is permissible according to n'. In this sense, (TB) *asymmetrically entails* (CB). That is, if an agent is in accordance with (TB), then they must also be in accordance with (CB), but not conversely.

Although (CB) accords well with (TB) there is a strong case to be made that (CB) conflicts with other plausible local norms, in particular:

(EB) **Evidential Norm for Belief**. Epistemically rational agents should (at any given time) believe propositions that are supported by their evidence.

It is plausible to interpret preface cases as revealing a tension between (EB) and (CB). Here is a rendition of the preface that we find particularly compelling:

> **Preface Paradox**. Let **B** be the set containing all of S's justified first-order beliefs. Assuming S is a suitably interesting inquirer, this set **B** will be a very rich and complex set of judgments. And, because S is fallible, it is reasonable to expect that some of S's first-order evidence is misleading. As a result, it seems reasonable to believe that some beliefs in **B** are false. Indeed, we think S herself would be justified in believing this very second-order claim. But, of course, adding this second-order belief to **B** renders S's total belief set inconsistent.

We take it that, in suitably constructed preface cases (such as this one), it would be epistemically permissible for S to satisfy (EB) but violate (CB). That is, we think that some preface cases are counterexamples to (CB). It is not our aim here to investigate whether this is the correct response to the preface paradox.[3] Presently, we simply take this claim as a *datum*. In this sense, our response to the preface is similar to the recent responses of Christensen (2004) and Kolodny (2007).

However, our approach to individual coherence diverges from Christensen's and Kolodny's. Christensen and Kolodny (and almost everyone else in this literature) would be inclined to accept the following conditional:

(†) If there are any (synchronic, epistemic) coherence requirements for full belief, then (CB) is among them.

[3]We think Christensen (2004) has given compelling arguments for the epistemic rationality of certain preface cases (*i.e.*, for the rationality of some inconsistent belief sets).

Christensen (2004) urges his readers to focus on partial belief (*viz., credence*). He suggests that all epistemological explanations (worth having) can be couched solely in terms of credences. In other words, Christensen seems to think that epistemology can (in some sense) do without full belief. As such, Christensen would be inclined to deny the antecedent of (†), which is a (trivial) way of accepting (†).[4]

Unlike Christensen, Kolodny does not think we can do without full belief (in epistemology). On the contrary, Kolodny thinks full belief is indispensable in epistemology (for proper accounts of practical and theoretical reasoning). However, Kolodny thinks that the only (synchronic) epistemic requirement on full belief is (EB). That is, Kolodny argues that there are no coherence requirements for full belief *per se*, and he offers a sophisticated error theory to explain away our intuitions to the contrary. As a result, Kolodny would also deny the antecedent of (†), but for different reasons than Christensen.

Our response to preface cases (and other epistemic paradoxes involving deductive consistency) differs from both Christensen's and Kolodny's. Whereas they would both abandon the idea that we should bother trying to articulate (synchronic, epistemic) coherence requirements for full belief, we would be inclined to say that (†) is false. Having said that, we do think there is a kernel of truth in each of Christensen's and Kolodny's responses. Unlike Kolodny, Christensen thinks there are coherence requirements for credences (*i.e.*, requirements of probabilistic coherence). We agree. In fact, our approach to grounding new coherence requirements for full belief was inspired by an existing approach to grounding probabilism as a coherence requirement for credences. Indeed, one of the main virtues of our approach is that it gives a unified framework for grounding both quantitative and qualitative coherence requirements. Unlike Christensen, Kolodny thinks that full belief is indispensable (in epistemology), and that (EB) is a *bona fide* epistemic requirement for full belief. We agree. And, this is why we think it's important to try to articulate and defend an alternative to (CB), which is consonant with (EB).

1.2. **A Principled Alternative to Deductive Consistency.** Our alternative to (CB) was not motivated by thinking about paradoxes of deductive consistency (like the preface). It was inspired by some recent arguments for probabilism as a (synchronic, epistemic) coherence requirement for credences. James Joyce (1998, 2009) has offered arguments for probabilism that are rooted in considerations of accuracy. We won't get into the details of Joyce's arguments here.[5] Instead, we present a general framework for grounding coherence requirements for sets of judgments of various types, including both credences and full beliefs. Our unified framework constitutes a generalization of Joyce's argument for probabilism. Moreover, when our approach is applied to full belief, it yields coherence requirements that are

[4]Here, we are going beyond what Christensen explicitly says in his book. He doesn't explicitly endorse an eliminativist stance regarding full belief (in epistemology). But, he does seem to imply that (*qua* epistemologists) we don't need to invoke coherence requirements for full belief *per se*. That is, he seems to think epistemology only requires coherence requirements for credences.

[5]There are some important disanalogies between Joyce's argument for probabilism and our analogous arguments regarding coherence requirements for full belief. See (Easwaran & Fitelson 2012) for discussion.

superior to (CB), in light of preface cases (and other similar paradoxes of consistency).

Applying our framework to judgment sets $\mathbf{J}$ of type $\mathfrak{J}$ only requires completing three steps. The three steps are as follows:

> **Step 1**. Say what it means for a set $\mathbf{J}$ to be *perfectly accurate* (at a possible world w). We use the term "vindicated" to describe the perfectly accurate set of judgments of type $\mathfrak{J}$, at w, and we use the abbreviation $\mathbf{J}_w$ to denote the vindicated set of judgments of type $\mathfrak{J}$, at w.[6]
>
> **Step 2**. Define a *measure of distance between judgment sets*, $d(\mathbf{J}, \mathbf{J}')$. We apply this measure to gauge the distance between a given set of judgments $\mathbf{J}$ of type $\mathfrak{J}$ and the vindicated set $\mathbf{J}_w$.
>
> **Step 3**. Adopt a *fundamental epistemic principle*, which uses $d(\mathbf{J}, \mathbf{J}_w)$ to ground a (synchronic, epistemic) coherence requirement for judgment sets $\mathbf{J}$ of type $\mathfrak{J}$.

This is all very abstract. To make things more concrete, let's look at the simplest application of our framework — to the case of *opinionated full belief*. Let:

$B(p) =_{df} S$ believes that p.

$D(p) =_{df} S$ disbelieves that p.

For simplicity, we suppose that S is opinionated, and that S forms judgments involving propositions drawn from a finite Boolean algebra of propositions. More precisely, let $\mathcal{A}$ be an agenda, which is a (possibly proper) subset of some finite boolean algebra of propositions. For each $p \in \mathcal{A}$, S either believes p or S disbelieves p, and not both.[7] In this way, an agent can be represented by her "belief set" $\mathbf{B}$, which is just the set of her beliefs (B) and disbeliefs (D) over some salient agenda $\mathcal{A}$. More precisely, $\mathbf{B}$ is a set of proposition-attitude pairs, with propositions drawn from $\mathcal{A}$ and attitudes taken by S toward those propositions (at a given time). Similarly, we think of propositions as sets of possible worlds, so that a proposition is true at any world that it contains, and false at any world it doesn't contain.[8]

[6]As a heuristic, you can think of $\mathbf{J}_w$ as the set of judgments of type $\mathfrak{J}$ that an omniscient agent (*i.e.*, an agent who is omniscient about the facts at world w) would have.

[7]The assumption of opinionation results in no loss of generality for present purposes. This is for two reasons. First, as Christensen (2004) convincingly argues, suspension of belief is not a plausible way out of the preface paradox (or other similar paradoxes of consistency). Second, in the context of judgment aggregation, it is typically assumed that judges are opinionated (at least, with regard to the agendas on which they are jointly making judgments). For exceptions, see, *e.g.*, Dietrich and List (2008, 2007b). In general, we would want to be able to model suspension of judgment in our framework (Friedman 2013). See (Easwaran 2012) for just such a generalization of the present framework.

[8]It is implicit in this formalism that agents satisfy a weak sort of logical omniscience, in the sense that if two propositions are logically equivalent, then they are in fact the same proposition, and so the agent can't have distinct attitudes toward them. However, it is not assumed that agents satisfy a stronger sort of logical omniscience — an agent may believe some propositions while disbelieving some other proposition that is entailed by them (*i.e.*, our logical omniscience assumption does not imply closure).

Step 1 is straightforward. It is clear what it means for a set **B** of this type to be perfectly accurate/vindicated. The vindicated set $\mathbf{B}_w$ is given by the following definition:

> $\mathbf{B}_w$ contains $B(p)$ $[D(p)]$ just in case p is true [false] at w.

This is clearly the best explication of $\mathbf{B}_w$, since $B(p)$ $[D(p)]$ is accurate just in case p is true [false]. So, in this context, Step 1 is uncontroversial.

Step 2 is less straightforward, because there are many ways one could measure "distance between judgment sets". For simplicity, we adopt perhaps the most naïve distance measure, which is given by:

> $d(\mathbf{B}, \mathbf{B}') =_{df}$ the number of judgments on which **B** and **B**′ disagree.[9]

In particular, if you want to know how far your judgment set **B** is from vindication (at w) just count the number of mistakes you have made (at w). To be sure, this is a very naïve measure of distance from vindication. In this paper, we will not delve into the various worries one might have about $d(\mathbf{B}, \mathbf{B}_w)$, or the plethora of alternative distance measures one could adopt. Here, our aim is primarily to explain the ramifications of our new approach to individual coherence for the existing dialectic concerning group coherence and judgment aggregation.

Step 3 is the philosophically crucial step. Given our setup, there is a choice of fundamental epistemic principle that yields (CB) as a coherence requirement for full belief. Specifically, consider the following principle:

> **Possible Vindication** (PV). There exists *some* possible world w at which *all* of the judgments in **B** are accurate. Or, to put this more formally, in terms of our distance measure d: $(\exists w)[d(\mathbf{B}, \mathbf{B}_w) = 0]$.

Given our setup, it is easy to see that (PV) is equivalent to (CB). As such, a defender of (TB) would presumably find (PV) attractive as a fundamental epistemic principle. However, in light of preface cases (and other paradoxes of consistency), many philosophers would be inclined to say that (PV) is too strong to yield a (plausible, binding) coherence requirement for full belief. Indeed, we ultimately opt for fundamental principles that are strictly weaker than (PV). But, as we mentioned above, our rejection of (PV) was not (initially) motivated by prefaces and the like. Rather, our adoption of fundamental principles that are weaker than (PV) was motivated (initially) by analogy with Joyce's arguments for probabilism as a coherence requirement for credences.

In the case of credences, the analogue of (PV) is clearly inappropriate. The vindicated set of credences (*i.e.*, the credences an omniscient agent would have) are such that they assign maximal credence to all truths and minimal credence to all falsehoods (Joyce, 1998). As a result, in the credal case, (PV) would require that all of one's credences be extremal. One doesn't need preface-like cases (or any other

[9] This is called the *Hamming distance* between the binary vectors **B** and **B**′ (Deza and Deza 2009). On distance measures between judgment sets, see, *e.g.*, Pigozzi (2006); Miller and Osherson (2009); Duddy and Piggins (2012).

subtle or paradoxical cases) to see that this would be an unreasonably strong requirement. It is for this reason that Joyce (and all others who argue in this way for probabilism) back away from the analogue of (PV) to strictly weaker epistemic principles — specifically, to accuracy-dominance avoidance principles, which are credal analogues of the following fundamental epistemic principle.

> **Weak Accuracy-Dominance Avoidance** (WADA). **B** is *not weakly*[10] *dominated* in distance from vindication. Or, to put this more formally (in terms of d), there does *not* exist an alternative belief set **B**′ such that:
>
> (i) $(\forall w)[d(\mathbf{B}', \mathbf{B}_w) \le d(\mathbf{B}, \mathbf{B}_w)]$, and
>
> (ii) $(\exists w)[d(\mathbf{B}', \mathbf{B}_w) < d(\mathbf{B}, \mathbf{B}_w)]$.

(WADA) is a very natural principle to adopt, if one is not going to require that it be possible to achieve perfect accuracy. Backing off (PV) to (WADA) is analogous to what one does in decision theory, when one adopts a weak dominance principle rather than a principle of *maximizing (actual) utility*.

Initially, it may seem undesirable for an account of epistemic rationality to allow for doxastic states that cannot be perfectly accurate. But, as Richard Foley (1992) explains, an epistemic strategy that is guaranteed to be imperfect is sometimes preferable to one that leaves open the possibility of vindication.

> ...if the avoidance of recognizable inconsistency were an absolute prerequisite of rational belief, we could not rationally believe each member of a set of propositions and also rationally believe of this set that at least one of its members is false. But this in turn pressures us to be unduly cautious. It pressures us to believe only those propositions that are certain or at least close to certain for us, since otherwise we are likely to have reasons to believe that at least one of these propositions is false. At first glance, the requirement that we avoid recognizable inconsistency seems little enough to ask in the name of rationality. It asks only that we avoid certain error. It turns out, however, that this is far too much to ask.

We agree with Foley's assessment that (PV) is too demanding. (WADA), however, seems to be a better candidate fundamental epistemic principle. As we will explain below, if S violates (WADA), then S's doxastic state *must* be defective — from both alethic and evidential points of view.

If an agent S satisfies (WADA) — *i.e.*, if S's belief set is *non-dominated* in distance from vindication — then we say S is **coherent** (we'll also apply the term 'coherent' to belief sets). To wit, our new coherence (*viz.*, non-dominance) requirement is

> (NDB) Epistemically rational agents should (at any given time) be coherent.

Interestingly, (NDB) is strictly weaker than (CB). Moreover, (NDB) is weaker than (CB) in the right way, in light of the preface case (and other similar paradoxes of consistency). Our first two theorems help to explain why.

[10] Strictly speaking, Joyce *et al.* opt for *strict* dominance-avoidance principles. However, in the credal case (assuming continuous, strictly proper scoring rules), there is no difference between weak and strict dominance (Schervish *et al.* 2009). So, there is no serious disanalogy here.

The first theorem states a necessary and sufficient condition for (*i.e.*, a characterization of) coherence: we call it *Negative* because it identifies certain objects, the *non*-existence of which is necessary and sufficient for coherence. The second theorem states a sufficient condition for coherence: we call it *Positive* because it states that in order to show that a certain belief set **B** is coherent, it's enough to construct a certain type of object.

Definition 1 (Witnessing Sets)**.** **S** *is a* ***witnessing set*** *iff (a) at every world, at least half of the judgments*[11] *in* **S** *are inaccurate; and, (b) at some world, more than half of the judgments in* **S** *are inaccurate.*

If **S** is a witnessing set and no proper subset of it is a witnessing set, then **S** is a **minimal witnessing set**. Notice that if **S** is a witnessing set, then it must contain a minimal witnessing set. Theorem 1 shows that the name "witnessing set" is apt, since these entities provide a witness to incoherence.

Theorem 1 (Negative)**.** **B** *is coherent if and only if no subset of* **B** *is a witnessing set.*

It is an immediate corollary of this first theorem that if **B** is logically consistent [*i.e*, if **B** satisfies (PV)], then **B** is coherent. After all, if **B** is logically consistent, then there is a world w such that no judgments in **B** are inaccurate at w. However, while consistency guarantees coherence, the converse is not the case. That is, coherence does not guarantee consistency. This will be most perspicuous as a consequence of our second central theorem:

Definition 2. *A probability function* Pr **represents** *a belief set* **B** *iff for every* $p \in \mathcal{A}$*:*

(i) **B** *contains* $B(p)$ *iff* $\Pr(p) > 1/2$.

(ii) **B** *contains* $D(p)$ *iff* $\Pr(p) < 1/2$.

Theorem 2 (Positive)**.** **B** *is coherent if*[12] *there is a probability function* Pr *that represents* **B**.

To appreciate the significance of Theorem 2, it helps to think about a standard lottery case.[13] Consider a fair lottery with n tickets, exactly one of which is the winner. For each $j \leqslant n$ (for $n \geqslant 3$), let p_j be the proposition that the j^{th} ticket is not the winning ticket. And, let q be the proposition that some ticket is the winner. Finally, let Lottery be the following belief set:

$$\left\{B(p_j) \mid 1 \leqslant j \leqslant n\right\} \cup \{B(q)\}.$$

Lottery is clearly coherent (just consider the probability function that assigns each ticket equal probability of winning), but it is not logically consistent. This

[11]Throughout the paper, we rely on naïve counting. This is unproblematic since all of our algebras are finite.

[12]For counterexamples to the converse of Theorem 2, see (Easwaran & Fitelson 2013).

[13] We are *not endorsing* the belief set Lottery in this example as *epistemically rational.* Indeed, we think that the lottery paradox is not as compelling — as a counterexample to (CB) — as the preface paradox is. On this score, we agree with Pollock (1990) and Nelkin (2000). We are just using this lottery example to make a formal point about the logical relationship between (CB) and (NDB).

explains why (NDB) is strictly weaker than (CB). Moreover, this example is a nice illustration of the fact that (NDB) is weaker than (CB) in a desirable way. More precisely, we can now show that (NDB) is entailed by both alethic considerations [(TB)/(CB)] and evidential considerations [(EB)].

While there is much disagreement about the precise content of (EB), there is widespread agreement that the following is a necessary condition for (EB).

> **Necessary Condition for Satisfying (EB).** S satisfies (EB), *i.e.*, all of S's judgments are *justified*, ***only*** *if*:
>
> ($\mathcal{R}$) There exists *some* probability function that probabilifies (*i.e.*, assigns probability greater than 1/2 to) each of S's beliefs and dis-probabilifies (*i.e.*, assigns probability less than 1/2 to) each of S's disbeliefs.

Many evidentialists agree that probabilification — relative to some probability function — is a necessary condition for justification. Admittedly, there is a lot of disagreement about which probability function is implicated in ($\mathcal{R}$).[14] But, because our Theorem 2 only requires the existence of some probability function that probabilifies S's beliefs and dis-probabilifies S's disbeliefs, it is sufficient to ensure (on most evidentialist views) that (EB) entails (NDB). And, given our assumptions about prefaces (and perhaps even lotteries), this is precisely the entailment that fails for (CB). Thus, by grounding coherence for full beliefs in the same way Joyce grounds probabilism for credences, we are naturally led to a coherence requirement for full belief that is a plausible alternative to (CB). This gives us a principled way to reject (†), and to offer a new type of response to preface cases (and other similar paradoxes of consistency). Figure 1 depicts the logical relations between the norms discussed in this section.

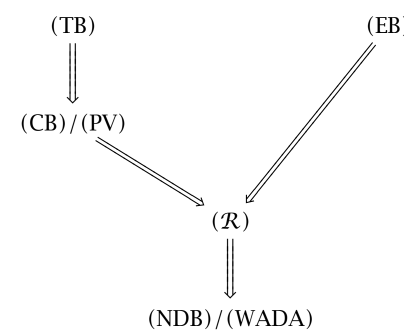

Figure 1. The logical relations between epistemic norms

[14]Internalists like Fumerton (1995) require that the function $\Pr(\cdot)$ which undergirds (EB) should be "internally accessible" to the agent (in various ways). Externalists like Williamson (2000) allow for "inaccessible" evidential probabilities. And, subjective Bayesians like Joyce (2005) say that $\Pr(\cdot)$ should reflect the agent's subjective degrees of belief (*viz.*, credences). Despite this disagreement, most evidentialists agree that (EB) entails ($\mathcal{R}$), which is all we need for present purposes.

In the remainder of the paper, we will explain how our new approach to individual coherence can undergird an interesting new conception of group coherence. This has important ramifications for judgment aggregation.

2. Consistency Preservation in Judgment Aggregation

2.1. **The Standard Doctrinal Paradox.** Recent interest in Judgment Aggregation has been partly fueled by interest in a paradox concerning the aggregation of individual judgment into collective judgment (Kornhauser and Sager, 1986). The basic idea is quite simple: when groups aggregate their opinions on logically connected propositions there may be cases in which the majority rule (and indeed any supermajority short of unanimity), may fail to preserve consistency.

As an example, consider a case in which a group of three judges $\{j_1, j_2, j_3\}$ makes the following judgments regarding three propositions $\{p, q, p \,\&\, q\}$.

	p	q	$p \,\&\, q$
j_1	D	B	D
j_2	B	D	D
j_3	B	B	B
majority	B	B	D

Despite the consistent judgments of the individual group members, the majority opinion is inconsistent. To see that the problem applies to any supermajority, one must simply generalize this pattern to larger sets of premises and larger group sizes.[15] In fact, this observation has been quite significantly generalized by thinking axiomatically about aggregation rules. A battery of impossibility results (List and Pettit, 2002; Pauly and van Hees, 2006; Dietrich, 2006; Dietrich and List, 2007a) has shown that many combinations of attractive properties are incompatible.

There is a fairly deep analogy between this paradox and the lottery paradox.[16] What is important for our current purposes is that analyzing the doctrinal paradox in terms of coherence turns out to be a fruitful endeavor (just as it is fruitful to analyze the lottery paradox in terms of coherence). If *coherence*, rather than *consistency* is our central normative concept, then we should investigate the possibility (or impossibility, as the case might be) of coherence-preservation in groups.

2.2. **Aggregation Framework.** Our aggregation framework is a slight generalization of the standard framework from List and Pettit (2002).[17]

Let $\mathcal{G}$ be a set of individuals (named $1, 2, \ldots, n$, with $n \geqslant 3$). Let $\mathcal{L}$ be a propositional language generated by a finite set of atomic sentences. An **agenda** $\mathcal{A}$ is a subset of $\mathcal{L}$ that is closed under negation. Let $\mathbf{B}_i$ be the **belief set** of individual i (on the propositions in $\mathcal{A}$): $\mathbf{B}_i$ is an assignment of exactly one of belief (B) or disbelief (D) to every proposition in $\mathcal{A}$. A **profile** is a sequence $\vec{\mathbf{B}}$ of belief sets (one judgment set for each judge in $\mathcal{G}$). Let $\mathcal{B}$ be the set of all possible belief sets, and let $\vec{\mathcal{B}}$ be the set of all possible profiles. Any function $f : \vec{\mathcal{B}} \to \mathcal{B}$ can be called an **aggregation function**. It is interpreted as assigning a "group" belief set to any given profile.

The space of aggregation functions can be effectively investigated by laying down various properties. If f is defined on every possible profile, f is said to have **universal domain**.[18] Our definition of an aggregation function implicitly requires universal domain, and our claims below implicitly presuppose it. f is **unanimous** iff, whenever all judges agree on their judgment about a proposition, the aggregated judgment set agrees with them as well. f is **dictatorial** iff there is a judge i such that for all profiles the collective belief set always coincides with i's belief set. f is **independent** iff for each proposition p, the collective judgment depends only on the individual judgments on p. An independent f is **inversive** if a profile with the opposite pattern of judgments on p gives the opposite collective judgment. f is **systematic** if the pattern of dependence does not vary across propositions.[19]

Formally, f is **dictatorial** iff there is an $i \in \mathcal{G}$, such that for all $\vec{\mathbf{B}}$, $f(\vec{\mathbf{B}}) = \mathbf{B}_i$. Let $\vec{\mathbf{B}}|p$ the sequence of beliefs on p alone (as determined by $\vec{\mathbf{B}}$). f is **independent** iff for each $p \in \mathcal{A}$, there is a function $h_p : \{B, D\}^{\mathcal{G}} \to \{B, D\}$, s.t. $f(\vec{\mathbf{B}})(p) = B \Leftrightarrow h_p(\vec{\mathbf{B}}|p) = B$. Let $\vec{b}$ be a sequence of judgments from $\{B, D\}^{\mathcal{G}}$ and let $\vec{b}^{-1}$ be the sequence of judgments with B and D reversed. f is **inversive** iff f is independent and for each $p \in \mathcal{A}$, $h_p(\vec{b}) \neq h_p(\vec{b}^{-1})$. f is **systematic** iff f is independent and for all $p, q \in \mathcal{A}$, $h_p = h_q$.

3. A Coherence-Based Perspective on the Doctrinal Paradox

Consider the Majority rule defined on odd-sized groups: for every proposition p in $\mathcal{A}$, $MAJ(\vec{\mathbf{B}})$ assigns whichever attitude to p more members of the group hold in $\vec{\mathbf{B}}$. (This rule is evidently not defined for even-sized groups, since the two attitudes could possibly be held by the same number of members.) The standard doctrinal paradox stems from the observation that the Majority rule does not preserve consistency. That is to say, assuming that the individual judges submit consistent opinions, there is no guarantee that the group majority is consistent. In this section, we investigate how this result extends to our more permissive epistemic norm. We start by providing two reasons for optimism: in many ordinary contexts of aggregation, an analogue of the doctrinal paradox for coherence simply

[15] Suppose for example that the acceptance threshold was 99%. Consider a case in which there are 101 relevant propositions, namely: $p_1, \ldots, p_{100}, (p_1 \& \ldots \& p_{100})$. Suppose that there are 101 judges: for all i between 1 and 100, j_i rejects p_i and accepts p_k for each $i \neq k$. Suppose finally that j_{101} accepts all of $p_1, \ldots, p_{100}$. Then the resulting 99%-supermajority opinion is: $\{p_1, \ldots, p_{100}, \sim(p_1 \,\&\, \ldots \,\&\, p_{100})\}$.

[16] As was first remarked by Levi (2004). See also Douven and Romeijn (2007); Chandler (2013).

[17] See also List and Puppe (2009), Grossi and Pigozzi (2012) and Mongin (2012) for recent surveys.

[18] Our usage of "universal domain" differs from standard usage, which presupposes that the possible profiles are restricted to profiles of belief sets that are themselves consistent. See §5.

[19] The inversive and systematic properties are logically independent. To see that inversive does not imply systematic, consider an aggregation rule that has one player as the dictator for some propositions and another player as the dictator for others. Although this is a strange rule for determining group beliefs, it is inversive, but not systematic. (It is inversive because reversing everyone's judgments includes reversing both dictators. It is not systematic because different dictatorships mean there is different dependence for different propositions.) To see that systematic does not imply inversive, consider an aggregation rule that says that the group believes every proposition in the agenda. This rule is systematic (since every proposition depends in the same, trivial, way on group beliefs) but not inversive (since no change to the group's beliefs reverses the group judgments).

doesn't arise, because in such contexts, aggregating by Majority produces coherent outcomes. Specifically, we argue that if either the judges are all consistent or the agenda has a certain kind of standard form, then Majority preserves coherence. We complement this argument with the observation that in the general case, there are some, rather complicated, failures of coherence preservation.

First, a general word about the logic of the argument is in order. Suppose that we have two normative constraints C_1 and C_2 such that C_1 is stronger than C_2. For example, these might be consistency and coherence, respectively. Suppose that we know that C_1 is not preserved by Majority. When investigating whether C_2 is preserved by Majority, we must account for two distinct effects. First, by weakening the normative constraint, we make it easier for the collective judgment to comply with the normative constraint itself. In the present setting, those cases in which Majority fails to preserve consistency turn out to be cases in which the majority opinion is coherent. The second effect is that by weakening the constraints to C_2, we must countenance new possible inputs to the aggregation process. Specifically, in the present setting, we must also consider what happens when we allow individuals to be coherent, but not consistent.

These two effects can be studied separately. Our first theorem shows that that, as long as the individuals that compose the group submit logically consistent attitudes, the group opinion is guaranteed to be coherent. (The proof in the Appendix shows that this follows immediately from Theorem 2.)

Theorem 3. *For every agenda $\mathcal{A}$, odd-sized group $\mathcal{G}$, and profile $\vec{\mathbf{B}}$, if all $\mathbf{B}_i$'s are consistent, then $MAJ(\vec{\mathbf{B}})$ is coherent.*

No incoherence can arise as long as the individuals submit perfectly consistent and complete judgments.

What about the general case? That is, what about the question whether our notion of rationality can be preserved by Majority for any *coherent* input profile? Here we distinguish two special cases.

We say that an agenda $\mathcal{A}$ is **simple** iff every minimal witnessing set of attitudes over the propositions in $\mathcal{A}$ assigns attitudes to exactly two propositions.[20]

Theorem 4. *Majority preserves coherence for simple agendas and odd-sized groups.*

To illustrate this theorem, note the following definition:

Definition 3. *An agenda $\mathcal{A}$ is truth-functional iff $\mathcal{A}$ can be partitioned in two subsets $\mathcal{A}^P$ and $\mathcal{A}^C$, such that:*
(i) no member of $\mathcal{A}$ is a tautology or a contradiction
(ii) any inconsistent subset of $\mathcal{A}^P$ contains a proposition-negation pair (in other words, there are no logical dependencies among the sentences in $\mathcal{A}^P$ except for those involving proposition-negation pairs).
(iii) $\mathcal{A}^P$ is closed under negation.

[20]Our usage of "simple" deviates from traditional usage, where "simple" is often interpreted as "contains no minimally inconsistent subset of size greater than two" (Dietrich & List 2007b).

(iv) $\mathcal{A}^C$ consists of a single proposition/negation pair.
(v) any maximal consistent subset of $\mathcal{A}^P$ entails a member of $\mathcal{A}^C$.

To illustrate, this definition characterizes agendas of the form:

$$\textbf{Conjunctive: } \{p_1, ..., p_n, (p_1 \,\&\, ... \,\&\, p_n), negations\}.$$

$$\textbf{Disjunctive: } \{p_1, ..., p_n, (p_1 \vee ... \vee p_n), negations\}.$$

In general, these are agendas that contain a set of "premises" and a "conclusion" that is some boolean compound of those premises, *e.g.*, the conjunctive agenda $\{p_1, ..., p_n, (p_1 \& ... \& p_n), negations\}$ is the union of the set $\{p_1, ..., p_n, negations\}$ (the *premises*) and the set $\{(p_1 \,\&\, ... \,\&\, p_n), \sim(p_1 \,\&\, ... \,\&\, p_n)\}$ (the *conclusion*).

Truth-functional agendas are an important and much-discussed class of agendas in Judgment Aggregation.[21]

Theorem 5. *Every truth-functional agenda is simple.*

Putting together theorems 4 and 5, we get:

Corollary 1. *Restricted to truth-functional agendas and odd-sized groups, Majority preserves coherence.*

It is a remarkable fact that such a large class of instances of the doctrinal paradox is eliminated when we move from consistency preservation to coherence preservation.[22]

However, the theory of judgment aggregation for coherent judgment sets is not as simple as these results might suggest. When we consider sufficiently elaborate agendas, it turns out that Majority does not preserve coherence. In fact, something a bit stronger is true: we can prove an impossibility result for coherence analogous to standard results that apply to consistency. That is to say, it turns out that a set of properties that are all satisfied by Majority (as well as by many other rules) are incompatible with (universal) coherence-preservation. The following definition characterizes one way in which an agenda can be "sufficiently elaborate".

Definition 4. *A belief set $\mathbf{B}$ is α-**almost-coherent** iff it is incoherent, but there are two distinct proposition-negation pairs $\{p, \sim p\}$ and $\{q, \sim q\}$ such that the belief set resulting from reversing judgments on either pair or both is coherent.*

*A belief set $\mathbf{B}$ is β-**almost-coherent** iff it is incoherent, but there are three distinct proposition-negation pairs $\{p, \sim p\}$, $\{q, \sim q\}$, and $\{r, \sim r\}$ such that reversing its judgments on any one of these pairs results in a coherent belief set.*

[21]In the early literature, almost all discussions of judgment aggregation were restricted to truth-functional agendas. More recently, broader classes of agendas have received some attention, *e.g.*, the atomically closed agendas of Pauly and van Hees (2006), and the agenda used by Dietrich and List (2007a) to derive Arrow's theorem as a corollary of a result in judgment aggregation.

[22]Note that Corollary 1 can be combined with other preservation results to yield slightly stronger results. For example, List (2013) characterizes a weakening of consistency that he calls 2-consistency: a belief set $\mathbf{B}$ is 2-consistent iff it does not include a proposition-negation pair. It is evident that Majority preserves 2-consistency, so it is a simple consequence of Corollary 1 that Majority preserves *the combination* of coherence and 2-consistency.

*An agenda $\mathcal{A}$ is **complex** iff there is a belief set* **B** *over $\mathcal{A}$ that is α-almost-coherent and there is a belief set* **B**′ *over $\mathcal{A}$ that is β-almost-coherent.*

Theorem 6. *If the agenda $\mathcal{A}$ is complex, then every inversive, systematic, unanimous rule that preserves coherence on $\mathcal{A}$ is dictatorial.*

In the Appendix, we give an example of a belief set that is both α and β-almost-coherent. Our example involves an agenda with five proposition-negation pairs. Unsurprisingly (since it follows from Theorem 6) we show that on the agenda for this belief set there are coherent profiles such that their majoritarian aggregate is not coherent.

4. Related Work.

We have investigated the preservation properties of a rationality constraint that is logically weaker than coherence. Naturally, much further work needs to be done to extend our approach. In this section we want to locate our analysis with respect to other studies that have also considered weaker rationality constraints.

One such study, and in fact one that is motivated in ways that are somewhat similar to ours, is List (2013). List advances a general distinction between *blatant* and *non-blatant* inconsistency. This distinction is formalized *via* the notion of k-consistency: a belief set **B** is k-consistent iff its smallest minimally inconsistent subset has size k. For example, compare two belief sets:

$\mathbf{B}_1$ assigns belief to every member of $\{p, q, \sim p\}$.
$\mathbf{B}_2$ assigns belief to every member of $\{p, q, \sim(p \,\&\, q)\}$

$\mathbf{B}_1$ is 2-inconsistent: its smallest minimally inconsistent subset has size 2 (it is the belief set that assigns belief to $\{p, \sim p\}$. By contrast, $\mathbf{B}_2$ is 3-inconsistent, but not 2-inconsistent. The result in List (2013) is that k-consistency is preserved by a supermajority with threshold $k - 1/k$.

The important point, for our purposes, is that our notion of coherence is neither stronger nor weaker than any of List's notions of k-consistency.

Theorem 7. *For every k, there is a set of propositions that is k-consistent but not $k+1$-consistent, such that an agent who believes all of them is coherent, and there is also a set of propositions that is k-consistent but not $k+1$-consistent, such that an agent who believes all of them is incoherent.*

Thus, these notions cross-cut each other.

It is however, elementarily possible to combine these notions of rationality to obtain stricter global norms (all of them weaker than consistency). For example, List's 2-consistency can be combined with our notion of *coherence* to produce a global rule that is evidence-sensitive but also rules out proposition-negation pairs.

In addition to List's work, there are studies, such as Dietrich's 2007 that explore generalized notions of consistency. Dietrich's aim is to extend the core results in the theory of judgment aggregation to a variety of non-classical logics. We stress

a technical and a philosophical point: technically, our notion of *coherence* does not meet all the requirements that Dietrich imposes on his generalized notion of consistency. Philosophically, Dietrich's work allows us to point out that our notion is not based on a non-classical understanding of the semantic apparatus. It is based, rather, on a non-classical understanding of what it is for a belief set to be in compliance with epistemic norms.

5. Conclusion

We have defended a coherence norm for belief, according to which epistemically rational agents should adopt beliefs that avoid weak accuracy-domination. This coherence norm is strictly weaker than the norm of deductive consistency, but possesses a number of advantages over the consistency norm.

First, our coherence norm can be justified by alethic considerations, via a simple dominance argument. An analogous justification of the consistency norm would require an implausibly strong 'possible vindication' premise, whose partial-belief analogue meets with general skepticism.

Second, our coherence norm is compatible with the evidential norm for belief, which states that epistemically rational agents should believe the propositions supported by their evidence. In the preface paradox, the evidential norm conflicts with deductive consistency. One common response is to give up on deductive consistency; this raises the question of what (if anything) to put in its place. Christensen suggests replacing it with an epistemology of partial belief, where partial beliefs are governed by the Kolmogorov axioms. If the Kolmogorov axioms are compatible with the evidential norm, then so is our coherence norm, since (as Theorem 2 states) every probability function can be represented by a coherent set of full beliefs.

Third, our coherence norm, unlike the deductive consistency norm, sets an attainable standard for collective rationality. It is well known that for a large class of agendas, Majority voting fails to preserve deductive consistency: even in a group whose members all have consistent opinions, there is no guarantee that the majority opinion will be consistent.

The move from consistency to coherence helps in two ways. First, if we keep the consistency requirement for individuals, but adopt a weaker coherence requirement for group beliefs, then beginning with permissible individual beliefs and taking a majority vote will always yield permissible group beliefs. In a group whose members all have consistent opinions, the majority opinion is guaranteed to be coherent (Theorem 3). Second, if we adopt the weaker coherence requirement for both individuals and groups, then beginning with permissible individual beliefs and taking a majority vote will still yield permissible group beliefs in an important class of cases. In particular, in a group whose members all have coherent opinions, the majority opinion is guaranteed to be coherent provided the group is odd-sized and the agenda is simple (Theorem 4). Results about simple agendas are of interest because all truth-functional agendas are simple (Theorem 5). Though the move from consistency to coherence is not a panacea for all voting paradoxes (see Theorem 6), it nonetheless represents significant progress.

We have shown that the coherence norm has useful implications for both individual and social epistemology. In both areas, we expect further investigation of coherence to yield further insight. In individual epistemology, there are unanswered questions about the relationship between partial beliefs and coherent full beliefs. What additional constraints must a coherent set of full beliefs satisfy in order to be represented by a probability function? Can these constraints be independently justified? In social epistemology, there are unanswered questions about the behavior of majority on agendas that are neither simple nor complex, and about the behavior of other rules. Our early investigations into coherence are only the beginning of a potentially fruitful research program.

Appendix

Proof of Theorem 1.

[**B** is coherent iff (⇔) it contains no witnessing set.]

(⇒) We'll prove the contrapositive. Suppose that $\mathbf{S} \subseteq \mathbf{B}$ is a witnessing set. Let $\mathbf{B}'$ agree with **B** on all judgments outside **S** and disagree with **B** on all judgments in **S**. By the definition of a witnessing set, $\mathbf{B}'$ must weakly dominate **B** in distance from vindication $[d(\mathbf{B}, \mathbf{B}_w)]$. Thus, **B** is incoherent.

(⇐) Again, we prove the contrapositive. Suppose that **B** is incoherent, *i.e.*, that there is some $\mathbf{B}'$ that weakly dominates **B** in distance from vindication $[d(\mathbf{B}, \mathbf{B}_w)]$. Let **S** be the set of judgments on which **B** and $\mathbf{B}'$ disagree. Then, **S** will be a witnessing set.

Proof of Theorem 2.

[**B** is coherent if there is a probability function Pr that represents **B**.]

Let Pr be a probability function that represents **B** in sense of Definition 2. Consider the expected distance from vindication of a belief set — the sum of $Pr(w)d(\mathbf{B}, \mathbf{B}_w)$. Since $d(\mathbf{B}, \mathbf{B}_w)$ is a sum of components for each proposition (1 if **B** disagrees with w on the proposition and 0 if they agree), and since expectations are linear, the expected distance from vindication is the sum of the expectation of these components. The expectation of the component for disbelieving p is $\Pr(p)$ while the expectation of the component for believing p is $1 - \Pr(p)$. Thus, if $\Pr(p) > 1/2$ then believing p is the attitude that uniquely minimizes the expectation, while if $\Pr(p) < 1/2$ then disbelieving p is the attitude that uniquely minimizes the expectation. Thus, since Pr represents **B**, this means that **B** has strictly lower expected distance from vindication than any other belief set with respect to Pr. Suppose, for *reductio*, that some $\mathbf{B}'$ (weakly) dominates **B**. Then, $\mathbf{B}'$ must be no farther from vindication than **B** in any world, and thus $\mathbf{B}'$ must have expected distance from vindication no greater than that of **B**. But **B** has strictly lower expected distance from vindication than any other belief set. Contradiction. Therefore, no $\mathbf{B}'$ can dominate **B**, and so **B** must be coherent.

Proof of Theorem 3.

[For every agenda $\mathcal{A}$, odd-sized group $\mathcal{G}$, and profile $\vec{\mathbf{B}}$, if all $\mathbf{B}_i$'s are consistent, then $MAJ(\vec{\mathbf{B}})$ is coherent.]

Since each judge is consistent, there must be some world in which the judge is accurate on every proposition in the agenda. Take one such world for each judge, and assign equal probability to each of these worlds (counting multiplicity, if the same world is repeated). Then any proposition that is accepted by the majority has probability greater than $1/2$ and any proposition that is rejected by the majority has probability less than $1/2$. Thus, the majority aggregate is representable by a probability function, and thus by Theorem 2, the majority aggregate is coherent.

Proof of Theorem 4.

[Restricted to simple agendas and odd-sized groups, Majority preserves coherence.]

Let G be an odd-sized group and let $\mathcal{A}$ be a truth-functional agenda. Suppose (*by reductio*) that $\vec{\mathbf{B}}$ is a profile such that $MAJ(\vec{\mathbf{B}}) = \mathbf{E}$ where **E** is an incoherent set. This means that it must contain some minimal witnessing subset, which must have exactly two propositions. Each of these propositions must have the relevant attitude assigned by more than half of the judges. Thus, there must be some judge that assigns both attitudes. This judge therefore has a minimal witnessing subset of her attitudes, and is therefore incoherent.

Proof of Theorem 5.

[Every truth-functional agenda is simple]

Let $\mathcal{A}$ be a truth-functional agenda. Let **S** be a minimal witnessing subset in $\mathcal{A}$. We break down the proof into three observations.

(i) **S** does not assign the same attitude to a proposition and its negation.

Proof: If it did, then in every world, these two attitudes would contribute one accurate and one inaccurate judgment. Thus, removing these two judgments would result in a belief set that has exactly one fewer accurate judgment and one fewer inaccurate judgment in every world than **S**. Since **S** was a witnessing set, this proper subset would be too, and thus **S** would not be minimal.

(ii) **S** must assign attitudes to exactly the same number of members of $\mathcal{A}^C$ and $\mathcal{A}^P$.

Proof: Since **S** doesn't assign the same attitude to a proposition and its negation, the attitudes **S** assigns within $\mathcal{A}^C$ and within $\mathcal{A}^P$ must each be consistent. This is because the only inconsistent subsets of either are proposition-negation pairs. Thus, there must be a world in which all **S**'s judgments in $\mathcal{A}^C$ are accurate, and there must be a world in which all **S**'s judgments in $\mathcal{A}^P$ are accurate. If either set constituted a majority of **S**, then the relevant world would show that **S** was not a witnessing set, and therefore the two parts must be the same size.

Note further that since $\mathcal{A}^C$ has exactly two propositions, this means that **S** must consist of either exactly two members of both $\mathcal{A}^P$ and two members $\mathcal{A}^C$, or exactly one member of each.

(iii) **S** does not assign attitudes to exactly two members of $\mathcal{A}^C$ and exactly two members of $\mathcal{A}^P$.

Proof: The two members of $\mathcal{A}^C$ are a proposition-negation pair, and by the first observation, **S** has opposite judgments on them, so in every world they are either both accurate or both inaccurate. Neither is a tautology, so there must be some world w_1 in which these two judgments are both accurate. In every such world, both judgments **S** assigns to members of $\mathcal{A}^P$ must be inaccurate.

Because **S** is a witnessing set, there must also be some world w_2 in which a strict majority of **S**'s judgments are inaccurate. In w_2, both judgments **S** assigns in $\mathcal{A}^C$ must be inaccurate, and at least one other judgment, say, the one to p_1, must also be inaccurate. But if c_1 is any member of $\mathcal{A}^C$, then there is no world where the judgments on p_1 and c_1 are both accurate. Since w_2 is a world in which both are inaccurate, $\{c_1, p_1\}$ is a witnessing set, which means that **S** was not minimal.

Thus, any minimal witnessing subset of such an agenda must assign attitudes to exactly one member of $\mathcal{A}^P$ and exactly one member of $\mathcal{A}^C$, so the agenda is simple.

Proof of Theorem 6

[If an agenda has an α-almost-coherent belief set, and a β-almost-coherent belief set, then every inversive, systematic, unanimous rule that preserves coherence on this agenda is dictatorial.][23]

Fix an inversive, systematic, unanimous aggregation rule f that preserves coherence.

Let $\mathcal{G}$ be the set of judges. Say that $C \subseteq \mathcal{G}$ is a **winning coalition** for p, B iff for every $\vec{B} \in \{B, D\}^{\mathcal{G}}$ where C is exactly the set of judges that assign B to p, then $f(\vec{B})(p) = B$. Say that $C \subseteq \mathcal{G}$ is a winning coalition for p, D iff for every $\vec{B} \in \{B, D\}^{\mathcal{G}}$ where C is exactly the set of judges that assign D to p, then $f(\vec{B})(p) = D$.

By inversiveness, C is a winning coalition for p, B iff it is a winning coalition for p, D. By systematicity, C is a winning coalition for p, B iff it is a winning coalition for q, B. Thus, for an inversive, systematic aggregation rule, we can talk about C being a winning coalition *simpliciter*. By independence, if there is *some* $\vec{B} \in \{B, D\}^{\mathcal{G}}$ where C is exactly the set of judges that agree with $f(\vec{B})$ on p, then C is a winning coalition.

Lemma 1: If C is not a winning coalition then its complement $-C$ is. To see this, fix some $\vec{B}$ where all judges in C assign B to p and all judges in $-C$ assign D to p. If $f(\vec{B})(p) = B$ then C is a winning coalition, and otherwise $-C$ is.

Lemma 2: If C is a winning coalition, then so is any superset of C.

Assume that $C \subset C'$, and C is a winning coalition, but C' is not. Since the agenda is evenly negated, there is a belief assignment X that is α-almost-coherent. That is, X is incoherent, and there are two distinct proposition-negation pairs $\{p, \sim p\}$ and $\{q, \sim q\}$ such that switching its judgments on either or both both gives a coherent belief assignment. Let every judge in C assign X with a reversal on $\{p, \sim p\}$. Let every judge in $C' - C$ assign X with reversal on both pairs. Let every judge outside of C' assign X with reversal on $\{q, \sim q\}$.

We show that the aggregate is X, and is thus incoherent. Every judgment in X other than those on p, $\sim p$, q, and $\sim q$ is shared by every judge, and thus by unanimity, is accepted by the aggregate. p and $\sim p$ have the reverse of X on every judge in C', but have the value given by X outside of C'. Since C' is not a winning coalition, its complement is, and thus the aggregate must agree with X on p and $\sim p$. q and $\sim q$ have the value from X on every judge in C, but not on any judge outside of C. Since C is a winning coalition, the aggregate has the value from X on q and $\sim q$. Thus, the aggregate is an extension of X, and thus must be incoherent, so our assumption that there could be a superset of a winning coalition that is not winning was false.

Lemma 3: The intersection of any two winning coalitions is a winning coalition.

Assume C and C' are winning coalitions but their intersection $C \cap C'$ is not. Since the agenda is β-almost-coherent, there is some incoherent partial belief assignment X and three distinct proposition-negation pairs $\{p, \sim p\}$, $\{q, \sim q\}$, and $\{r, \sim r\}$ such that reversing X's judgment on any one of these pairs yields a coherent belief assignment. Let every judge in $C \cap C'$ assign X reversed on p and $\sim p$. Let every judge in $C' - C$ assign X reversed on q and $\sim q$. Let every judge outside C' assign X reversed on r and $\sim r$.

We show that the aggregate is X, and is thus incoherent. All judges agree with every judgment in X outside of p, q, r and their negations, and thus by unanimity, the aggregate does too. All judges outside of $C \cap C'$ agree with X's judgments on p and $\sim p$, while the judges in $C \cap C'$ disagree. Since $C \cap C'$ is not a winning coalition, its complement is, and thus X's judgments on p and $\sim p$ are shared by the aggregate. All judges outside of $C' - C$ agree with X's judgments on q and $\sim q$, while the judges in $C' - C$ disagree. But the complement of $C' - C$ is a superset of C, and is thus a winning coalition, by Lemma 2. Thus, X's judgments on q and $\sim q$ are shared by the aggregate as well. All judges in C' agree with X's judgments on r and $\sim r$. Since C' is a winning coalition, the aggregate does too. Thus, the aggregate includes all of X's judgments, so the aggregate must be incoherent, which contradicts the fact that f preserves coherence. Thus, our assumption that the intersection of two winning coalitions could fail to be a winning coalition was false.

Proof of theorem: By Lemma 1, for every singleton, either it or its complement is a winning coalition. If none of the singletons is a winning coalition, then every complement of a singleton is. But by Lemma 3, the intersection of any pair of winning coalitions is itself a winning coalition. But by repeated intersection of the complements of all singletons, we conclude that if no singleton is a winning coalition, then the empty set is, which is impossible. Thus, at least one singleton must be a winning coalition. So therefore, the aggregation rule must be a dictatorship, QED.

Example.

Here we want to provide an example of an agenda that satisfies the condition of Theorem 6. We also show that this agenda allows coherent input profiles with incoherent majoritarian aggregates.

[23] This proof replicates in a Coherence setting a standard proof of Impossibility from the Judgment Aggregation literature. We adapted our proof from the proof of (Grossi and Pigozzi 2012).

Consider the boolean algebra over (at least) 11 worlds, $w_1, w_2, \dots, w_{10}, w_{11}$.[24] Consider the agenda given by the following five propositions (and their negations — in what follows, we ignore mention of the negations, and assume that each judge makes consistent judgments on every proposition-negation pair):

$$\begin{aligned} A &= \{w_1, w_2, w_3, w_4\} \\ B &= \{w_1, w_5, w_6, w_7\} \\ C &= \{w_2, w_5, w_8, w_9\} \\ D &= \{w_3, w_6, w_8, w_{10}\} \\ E &= \{w_4, w_7, w_9, w_{10}\} \end{aligned}$$

The important thing is that for every pair of these propositions, there is exactly one world where both are true, and in any world where one of the propositions is true, exactly one other is as well.

On this agenda, note that the belief set $\langle B, B, B, B, B\rangle$ is incoherent — in every world it gets a majority of the propositions wrong (exactly three wrong in each of the worlds $w_1, \dots, w_{10}$, and all wrong in all other worlds), so it is dominated by $\langle D, D, D, D, D\rangle$.

The belief set $\langle B, B, B, B, D\rangle$ is both α and β-almost-coherent. To show this, it suffices to show three things: that $\langle B, B, B, B, D\rangle$ is incoherent; that $\langle B, B, B, D, D\rangle$ is coherent (symmetry considerations mean that this also shows that $\langle B, B, D, B, D\rangle$, $\langle B, D, B, B, D\rangle$, and $\langle D, B, B, B, D\rangle$ are coherent); and that $\langle B, B, D, D, D\rangle$ is coherent. Thus, the agenda satisfies the conditions of Theorem 6.

To see that $\langle B, B, B, B, D\rangle$ is incoherent, note that it is weakly dominated by $\langle D, D, D, D,$ On worlds $w_1, w_2, w_3, w_5, w_6, w_8$, both belief sets have distance 2 from vindication. On worlds w_4, w_7, w_9, w_{10}, the former has distance 4 from vindication while the latter has distance 2 from vindication. On world w_{11}, the former has distance 4 from vindication while the latter has distance 0 from vindication.

Now we show that $\langle B, B, B, D, D\rangle$ is coherent. Consider the probability distribution that assigns probability 0 to any world in D or E and any world not in any of the five propositions, and probability $1/3$ to each of the three remaining worlds. (In this case, w_1, w_2, w_3.) On this distribution, every proposition that is believed has probability strictly greater than $1/2$ (in fact, they all have probability $2/3$) and every proposition that is disbelieved has probability strictly less than $1/2$ (in fact, both have probability 0). Thus, by Theorem 2, the assignment is coherent.

Finally, we show that $\langle B, B, D, D, D\rangle$ is coherent. Consider the probability distribution that assigns probability $1/3$ to the unique world shared by the two propositions that are believed (in this case w_1), probability $1/9$ to each of the six worlds that are in exactly one of those two propositions (in this case $w_2, w_3, w_4, w_5, w_6, w_7$), and probability 0 to all other worlds. On such a distribution, the two believed propositions have probability $2/3$, and the three disbelieved propositions all have probability $2/9$, and thus by Theorem 2, the belief set is coherent.

[24] If you prefer to think of propositions sententially, then you can generate propositions with exactly this same logical structure by taking atoms p_1, p_2, p_3, p_4, and considering $w_1, \dots, w_{16}$ as the 16 state descriptions (conjunctions of these four atomic sentences or their negations), and the propositions as each being a disjunction of four state descriptions.

Therefore, this agenda satisfies the conditions of the theorem, and so the only aggregation function that is independent, inversive, systematic, and unanimous while preserving coherence is dictatorship.

In particular, we can see that Majority fails to preserve coherence. Just consider ten coherent judges who each have one permutation of $\langle B, B, B, D, D\rangle$. The majority judgment must be $\langle B, B, B, B, B\rangle$, which is incoherent.

Two final clarificatory notes on Theorem 6 are in order. First, “systematic” does not imply “inversive”. To see this, consider the rule that says the group believes every proposition on the agenda. This rule is systematic, but not inversive. Second, systematicity is essential to the theorem — independence alone does not suffice. To see this, consider the following example.[25] Let $\mathcal{A}$ be (any) complex agenda. Now consider the agenda $\mathcal{A}^+ = \mathcal{A} \cup \{q, \neg q\}$, where q is some (non-tautological and non-contradictory) proposition that is logically unrelated to any of the propositions in $\mathcal{A}$. Then, $\mathcal{A}^+$ will also be complex. Now consider the following aggregation procedure: On the subagenda $\mathcal{A}$, individual 1 determines the collective judgments, and on the subagenda $\{q, \neg q\}$, individual 2 determines the collective judgments. Although this aggregation procedure is dictatorial on each of these two subagendas, it is not dictatorial on the agenda $\mathcal{A}^+$ in its entirety, since there is no single dictator who determines the collective judgment for every proposition.

Proof of Theorem 7

[For every k, there is a set of propositions that is k-consistent but not $k+1$-consistent, such that an agent who believes all of them is coherent, and there is also a set of propositions that is k-consistent but not $k+1$-consistent, such that an agent who believes all of them is incoherent.]

First, we will construct a set of $2k-1$ propositions that is k-consistent but not $k+1$-consistent. Since there are worlds in which a majority of the propositions are true, the set will be coherent. Let S be the set of all subsets of $\{1, \dots, 2k-1\}$ of size exactly k. For each s in S, let there be a distinct world w_s, and let this be all the worlds the make up the boolean algebra. Define a set of propositions $\{p_1, \dots, p_{2k-1}\}$ as follows: the proposition p_i contains the world w_s iff $i \in s$. Then, for every set of k distinct propositions from this set, the indices form a set s, and the world w_s will make all of these propositions true, so the set is k-consistent. However, any collection of $k+1$ propositions will contain at least one propositions from outside of s, and thus there is no world in which all of them are true, and so the set is not $k+1$-consistent.

Similarly, we can construct a set of $2k+1$ propositions that is k-consistent but not $k+1$-consistent. Since there is no world where at least half of the propositions are true, the set will be incoherent. Let S be the set of all subsets of $\{1, \dots, 2k+1\}$ of size exactly k. For each s in S, let there be a distinct world w_s, and let this be all the worlds the make up the boolean algebra. Define a set of propositions $\{p_1, \dots, p_{2k+1}\}$ as follows: the proposition p_i contains the world w_s iff $i \in s$. Then, for every set of k distinct propositions from this set, the indices form a set s, and the world w_s will make all of these propositions true, so the set is k-consistent.

[25] We thank an anonymous referee for this example.

However, any collection of $k+1$ propositions will contain at least one propositions from outside of s, and thus there is no world in which all of them are true, and so the set is not $k+1$-consistent.

References

Chandler, J. (2013). Probability, acceptance and aggregation. *Erkenntnis 78*(1), 201–217.

Christensen, D. (2004). *Putting Logic in its Place.* Oxford University Press.

Deza, M. and E. Deza (2009). *Encyclopedia of Distances.* Springer.

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

Dietrich, F. (2007). A generalised model of judgment aggregation. *Social Choice and Welfare 28*(4), 529–565.

Dietrich, F. and C. List (2007a). Arrow's theorem in judgment aggregation. *Social Choice and Welfare 29*(1), 19–33.

Dietrich, F. and C. List (2007b). Judgment aggregation with consistency alone. Online working paper (`http://eprints.lse.ac.uk/20110/`).

Dietrich, F. and C. List (2008). Judgment aggregation without full rationality. *Social Choice and Welfare 31*(1), 15–39.

Douven, I. and J. Romeijn (2007). The discursive dilemma as a lottery paradox. *Economics and Philosophy 23*(3), 301–319.

Duddy, C. and A. Piggins (2012). A measure of distance between judgment sets. *Social Choice and Welfare 39*(4), 855–867.

Easwaran, K. (2012). Dr. Truthlove, or how i learned to stop worrying and love Bayesian probabilities. Unpublished manuscript.

Easwaran, K. and B. Fitelson (2012). An "evidentialist" worry about Joyce's argument for probabilism. *Dialectica 66*(3), 425–433.

Easwaran, K. and B. Fitelson (2013). Accuracy, coherence and evidence. To appear in *Oxford Studies in Epistemology* (Volume 5), T. Szabo Gendler & J. Hawthorne (*eds.*), Oxford University Press.

Foley, R. (1992). *Working Without a Net.* Oxford University Press.

Friedman, J. (2013). Suspended judgment. *Philosophical studies 162*(2), 165–181.

Fumerton, R. (1995). *Metaepistemology and Skepticism.* Rowman & Littlefield.

Grossi, D. and G. Pigozzi (2012). Introduction to judgment aggregation. In N. Bezhanishvili and V. Goranko (Eds.), *Lectures on Logic and Computation*, pp. 160–209. Springer.

Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science 65*(4), 575–603.

Joyce, J. M. (2005). How probabilities reflect evidence. *Philosophical Perspectives 19*(1), 153–178.

Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In F. Huber and C. Schmidt-Petri (Eds.), *Degrees of Belief*. Springer.

Kaplan, M. (2013). Coming to terms with our human fallibility: Christensen on the preface. *Philosophy and Phenomenological Research 87*(1), 1–35.

Klein, P. (1985). The virtues of inconsistency. *The Monist 68*(1), 105–135.

Kolodny, N. (2007). How does coherence matter? In *Proceedings of the Aristotelian Society*, Volume 107, pp. 229–263.

Kornhauser, L. A. and L. G. Sager (1986). Unpacking the court. *The Yale Law Journal 96*(1), 82–117.

Kyburg, H. E. (1970). Conjunctivitis. In M. Swain (Ed.), *Induction, acceptance and rational belief*, pp. 55–82. Reidel.

Levi, I. (2004). List and Pettit. *Synthese 140*(1), 237–242.

List, C. (2013). When to defer to supermajority testimony–and when not to. This volume.

List, C. and P. Pettit (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy 18*(1), 89–110.

List, C. and C. Puppe (2009). Judgment aggregation. In P. Anand, P. Pattanaik, and C. Puppe (Eds.), *The handbook of rational and social choice*, pp. 457–482. Oxford University Press.

Miller, M. K. and D. Osherson (2009). Methods for distance-based judgment aggregation. *Social Choice and Welfare 32*(4), 575–601.

Mongin, P. (2012). The doctrinal paradox, the discursive dilemma, and logical aggregation theory. *Theory and decision 73*(3), 315–355.

Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *The Philosophical Review 109*(3), 373–409.

Pauly, M. and M. van Hees (2006). Logical constraints on judgment aggregation. *Journal of Philosophical Logic 34*(6), 569–585.

Pigozzi, G. (2006). Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation. *Synthese 152*(2), 285–298.

Pollock, J. L. (1990). *Nomic probability and the foundations of induction.* Oxford Univ. Press.

Ryan, S. (1991). The preface paradox. *Philosophical studies 64*(3), 293–307.

Ryan, S. (1996). The epistemic virtues of consistency. *Synthese 109*(2), 121–141.

Schervish, M., T. Seidenfeld, and J. Kadane (2009). Proper scoring rules, dominated forecasts, and coherence. *Decision Analysis 6*(4), 202–221.

Williamson, T. (2000). *Knowledge and Its Limits.* Oxford University Press.