Aggregating with Reason(s). *

Fabrizio Cariani

[This is a postprint of a paper published in *Synthese*, 190(15), pp. 3123-3147, 2013.
Submitted: 4/20/11. Revised: 2/3/12: Accepted: 5/10/12. First Published Online: 6/25/12;
Post-publication corrections: 7/15/12]

Judgment Aggregation studies procedures for aggregating coherent individual opinions into coherent collective opinions. The study of such procedures[1] is often motivated by the need to find a mathematical solution to the so-called *doctrinal paradox*. Start with the idea of Aggregation rules as functions that input individual judgments and output collective judgments on some salient propositions. Say that an aggregation rule $\mathcal{A}$ is *consistency preserving* just in case whenever all group members submit logically consistent opinions, $\mathcal{A}$ outputs a logically consistent set of opinions.[2]

The *doctrinal paradox* consists of the observation that for some sets of propositions the majority rule (albeit generally attractive) is not consistency-preserving. There are many instances of this phenomenon, with different kinds of logical connections. Suppose that there are three judges (j_1, j_2, j_3), three relevant propositions, ϕ, ψ and $\phi \vee \psi$.

	ϕ	ψ	$\phi \vee \psi$
j_1	Yes	No	Yes
j_2	No	Yes	Yes
j_3	No	No	No

Example 1

* I thank Branden Fitelson, Stephan Hartmann, Jeff Horty, Eric Pacuit, Gabriella Pigozzi, Sherri Roush, Jan Sprenger and an anonymous referee. Special thanks to Shyam Nair for delivering written comments on this paper at FEW 2011. I also thank for their feedback audiences at Tilburg University, FEW 2011, Workshop on Voting Methods and Judgment Aggregation (where I presented a descendent of this paper). Work on this paper was partly suported by Sherri Roush's NSF project *Fallibility and Revision in Science and Society*, Award No. SES - 0823418, whose support I gratefully acknowledge.

1 See List and Pettit (2002) for the initial spark and List and Puppe (2008) for a recent survey.

2 Throughout this paper, the relevant sense of consistency (viz. entailment) is truth-functional consistency (viz. entailment) as analyzed in sentential logic. Both notions often occur relativized to background knowledge.

This observation generalizes: there are simple properties such that any aggregation rule that satisfies all of them is not consistency-preserving.[3] Solving the *doctrinal paradox* consists in characterizing natural, principled and satisfiable classes of properties that are compatible with consistency-preservation. Call this the *coherence challenge*.

The importance of meeting this challenge cannot be underestimated, but we can also aim for something slightly different. Pettit (2001) frames the doctrinal paradox within the context of a discussion of deliberative democracy and gives a central place to reasons: "the problem in question is [...] tied [...] only to the enterprise of making group judgments on the basis of reasons" (p. 272). Similarly, in motivating her proposed aggregation rule, Pigozzi (2008) stresses the importance of collective reasons:

> A verdict in a court is a public act. Not only, if convicted, has a defendant the right to know the reasons for which she has been convicted, but also these reasons will guide future decisions — they are patterns for future verdicts. In other words, the final decision must be supported and justified by reasons. (p.289)

Call the problem of producing a collective judgment supported by collective reasons the *reasons challenge*.

The two challenges are clearly conceptually distinct: the central question of the *reasons challenge* ("what reasons can a group provide for or against a certain opinion or decision?") arises even when the group's opinion is perfectly coherent. The Judgment Aggregation framework is generally thought to incorporate a satisfactory approach to reason-based group choice and reason-based group opinion. In this paper, I argue that this approach is actually limited in scope and should be generalized in natural ways.

I develop this argument as follows. In section §1, I illustrate some specific issues surrounding the role of reasons in the context of Judgment Aggregation. Here, I stress some limitations of the dominant approach to collective reasons and I identify three specific tasks for a generalized framework. In section §2, I develop a generalization of the aggregation framework which better captures the role of reasons in judgment aggregation. In section §3, I define a class of previously unavailable rules and explain how we can understand collective reasons within this class. I prove some simple results that help understand the behavior of the rules in this class. In §4, I discuss how the rules I propose, together with the account of collective reasons, can help

[3] List and Pettit (2002), Pauly and van Hees (2006).

satisfy my *desiderata*. In §5, I motivate a further generalization. I argue, however, that there are limits to the extent to which this generalization can be formalized.

1. Motivation

1.1. The Premise-Conclusion Dichotomy.

The standard approach to the *reasons challenge* is to introduce a distinction between two kinds of propositions. Instead of thinking that all relevant propositions have equal roles, one could think that some propositions have *reasons*-roles (they function as reasons for others) and others have target-roles (they are the sorts of things that reasons are given for or against). This idea can be developed by designating a set of pairwise logically independent sentences as *premises* and designating a further sentence as the *conclusion*. Additionally, any distribution of truth-values on the premises must settle by entailment (relative to background knowledge) the truth-value of the conclusion.

Consider again Example 1. The *agenda* (i.e. the set of relevant propositions) is the set: $\{\phi, \psi, \phi \vee \psi, \textit{negations}\}$. We can naturally partition this agenda into a set containing the *premises* (e.g., the propositions $\{\phi, \psi, \textit{negations}\}$) and one containing the *conclusion* ($\{\phi \vee \psi, \textit{negations}\}$). A popular (but not the only) way to aggregate judgments in this framework is *Premise-Based Majority*: take majority on the premises and propagate by entailment to the conclusion.[4] Premise-Based Majority suggests a general answer to the coherence challenge. According to

[4] One alternative is to take the majority only on the conclusion and ignore the verdicts on the premises (this is called *Conclusion-Based Majority*). Conclusion-Based Majority shares with Premise-Based Majority the need for a specification of a Premise-Conclusion Dichotomy.

It is a more complex question whether *Distance-Based* rules (see, for instance, Konieczny*et. al.*, 2004, Pigozzi, 2008, Miller and Osherson, 2009, Chandler, forthcoming) also presuppose a Premise-Conclusion Dichotomy. The mathematical formulation of distance-based rules does not require a Premise-Conclusion Dichotomy. One may need to talk about the distinction between atomic and complex sentences, but one doesn't need to call some of them "premises".

However, standard distance-based rules, by themselves, are just answers to the *coherence* challenge, they do not address the *reasons* challenge. So say, that a certain distance-based rules outputs the judgment set $\{\phi, \psi, \phi \;\&\; \psi\}$: in order to explain what functions as a reason and what functions as a conclusion one must still invoke the Premise-Conclusion Dichotomy.

Premise-Based Majority, in Example 1, the group should accept the coherent judgment set: $\{\sim\phi, \sim\psi, \sim(\phi \vee \psi)\}$.

The Premise-Conclusion dichotomy may be thought to help with the reasons challenge, by setting up a link such as:

Premise-Reasons Link
The *reasons* for a collective judgment on a conclusion are just those judgments on the premises that entail the verdict on the conclusion.

So if the group judges $\{\phi, \sim\psi, \phi \vee \psi\}$, the reasons for its adopting the judgment $\phi \vee \psi$ are given by the set $\{\phi, \sim\psi\}$.[5] The Premise-Reasons Link does not require that Premise-Based Majority be used: whatever rule one is using, provided that the output is $\{\phi, \sim\psi, \phi \vee \psi\}$, the Premise-Reasons Link implies that the reasons for $\phi \vee \psi$ are $\{\phi, \sim\psi\}$.

The picture of reasons I just sketched is fairly common in the literature (with some exceptions I'll flag shortly). I don't consider it to be fundamentally mistaken, but I consider it to be insufficiently general on a number of dimensions. First, it limits the domain of applicability of judgment aggregation to cases that are rather artificial. We need pairwise logically independent premises and a conclusion that can be settled by *any* distribution of truth-value on the premises.[6]

Second, it requires a group to have antecedently agreed on a Premise-Conclusion dichotomy—that is to have decided which propositions count as reasons and which count as conclusions. This is not how reasons generally work. Which propositions count as reasons and which count as targets can (and usually does) vary from judge to judge. In fact, it can even vary among judges with identical opinions.

Let me illustrate these points with some examples.

Hard Decisions.
A university administration must decide whether to start a basketball program (B) and whether to start a volleyball program (V). The teams cannot train in the same gym, so the administration also needs to decide whether to buy a new space (N). They all agree that they will buy the new space if they decide to start both programs

[5] The fact that, on this account, $\sim\psi$ is part of the reason for accepting $\phi \vee \psi$ is a bit of an embarrassment, but one can complicate the criterion so as to circumvent this problem. I will not attempt to do this, since I am about to point to some further, independent, issues.

[6] This point is effectively pressed, among others, by Nehring and Puppe (2010), who also use it to motivate some new aggregation rules that are not premise based but that, like the rules I present below, require that a collectively accepted proposition be justified, when possible.

(B & $V \to N$)–but they are contemplating buying the new space regardless of the outcome of the vote on the program.

In cases like this, it would make sense for one judge to take B and V as reasons for N, while another judge takes B and $\sim N$ as reasons for $\sim V$, and yet another takes $\sim N$ and V as reasons for $\sim B$. No independently fixed pattern of premisehood can capture this flexibility.

This phenomenon can occur even when two judges agree on every proposition at issue. For instance, consider biconditional agendas—agendas of the form $\{\phi, \psi, \phi \equiv \psi, \textit{negations}\}$. Imagine two judges (let's call them *he* and *she*) who agree that $\{\phi, \psi, \phi \equiv \psi\}$. Such judges may still disagree on which judgments ground which others. Perhaps, *he* accepts ψ on the basis of ϕ and $\phi \equiv \psi$; *she*, on the other, hand accepts ϕ because she accepts ψ and $\phi \equiv \psi$. This difference is unlikely to affect what propositions they should collectively accept, but it does matter to how they should present their reasons.[7]

In sum, the Premise-Reasons Link requires two problematic assumptions:

(i) if an agent i in a group accepts ϕ and ϕ functions as a reason (rather than as a conclusion) in i's epistemic state, then if any other agent k in the same group accepts $\sim\phi$, $\sim\phi$ must function as a reason (rather than as a conclusion) in k's epistemic state.

(ii) any two agents with the same opinions, must share the same pattern of reasons.

Generally speaking, these assumptions are false—both for individuals and for groups.[8]

The first task of my paper is, then, to produce a framework that generalizes the judgment aggregation framework without assuming that reason-roles and conclusion-roles are externally fixed. The way to

[7] The point is completely general: for example, a conjunction could be accepted on the basis of its conjuncts, or, vice-versa, the conjuncts could be accepted on the basis of the conjunction.

[8] For the same reason, I am also inclined to avoid addressing the *reason challenge* within a framework due to Dietrich. Dietrich (2008) studies a framework in which, in addition to all the standard tools, we single out a relation of dependence $\blacktriangleleft$ holding among propositions in the agenda. Dietrich does not impose a particular interpretation on $\blacktriangleleft$, suggesting that it could vary from application to application. He does remark that, given some structural constraints, it could be interpreted as a generalization of the *premise/conclusion* dichotomy. If that is meant to suggest that it can help model the *reasons-relation* in the sense I have been discussing, I offer exactly the same criticism: $\blacktriangleleft$ is fixed externally and the *reasons-relation* is not.

do this, in my view, is to start with finer inputs: instead of aggregating states that are just constituted by patterns of judgment, I aggregate states that are constituted by patterns of judgments together with a *reasons-relation*.

1.2. Conflicting Reasons

Once we acknowledge the possibility of representing individual judges as having different patterns of reasons, further issues emerge. For example, I argue elsewhere (Cariani, ms.) that the nature of the relationships among the reasons in support of ϕ can make a difference as to whether or not ϕ is accepted. I contrast the two following cases:[9]

> The Stockbrokers
> You have invested equal amounts of money in two companies: *Cookbooks* and *Shoes*. You have no other stocks. Your stockbroker, Bob, thinks that next week *Cookbooks* will improve by 10%, while *Shoes* will break even. Your other stockbroker, Jim, thinks that *Shoes* will improve by 10% while *Cookbooks* will break even. Should your stockbrokers agree that your portfolio will improve by 5%?

> Little Kids
> Jodie and Karl are at the park, trying to figure out the features of a person they see in the distance. Based on her assessment of the person's size and speed, Jodie thinks the person is a child C. Karl agrees with Jodie. However, Jodie thinks she has additional information that leads her to conclude that the person is a boy (B), while Karl thinks he has evidence for the claim that the person is a girl (G). They both accept the background claims: $(B \vee G) \rightarrow C, B \equiv \sim G$.

The problems have a similar logical structure. There is a disjunctive proposition on which the judges agree, but they have opposite opinions on the disjuncts (for a similar structure with an odd number of judges see the next subsection):

	ϕ	ψ	$\phi \vee \psi$
1	Yes	No	Yes
2	No	Yes	Yes

Example 2

What is striking, however, is that despite the parallel structure, the intuitions do not seem parallel:

[9] The first case was inspired by a structurally similar case in Horty (2002).

- In Stockbrokers, there is stronger (though not irresistible) pressure to think that the group should not accept C (this would be the verdict of Premise-Based Majority).

- In Little Kids, there is far less pressure in the same direction. Even though Jodie and Karl disagree on B and G, they accept C on the basis of independent (perhaps non-deductive) evidence.

I stress that the intuition in Stockbrokers is not irresistible. Sometimes, especially when under demanding practical circumstances, we want, do, and should reach *incompletely theorized agreements* in the sense of Sunstein (1995):

> People often reach *incompletely theorized agreements* on a general principle. Such agreements are incompletely theorized in the sense that people who accept the principle need not agree on what it entails in particular cases. People know that murder is wrong, but they disagree about abortion. They favor racial equality, but they are divided on affirmative action. (p.1739)

An incompletely theorized agreement seems definitely right in Little Kids. As far as Stockbrokers, however, intuition can go both ways. It may certainly be practically useful for them to reach an incompletely theorized agreement; but it's an open question whether logical or epistemological principles *mandate* such an agreement.

The standard framework, however, has an odd consequence. Suppose that the two cases are given isomorphic formalizations—e.g. as in Example 2. Then, any aggregation rule must output isomorphic verdicts; after all, aggregation rules are *functions* of their inputs, and in particular they are functions that do not make distinctions between problems with the same logical structure. This assumption means that it's impossible to have a rule that accepts the disjunction in Little Kids but rejects it in Stockbrokers. Using my framework, I am going to represent the cases as corresponding to non-isomorphic inputs, and hence show that we define rules that distinguish them.

1.3. Reasons vs. Unanimity

Examples with the structure of Stockbrokers and Little Kids have been first identified and discussed in the Judgment Aggregation literature by Nehring (2005), who considered some strengthened versions of the doctrinal paradox with this form.

	ϕ	ψ	θ	$\phi \vee \psi \vee \theta$
1	Yes	No	No	Yes
2	No	Yes	No	Yes
3	No	No	Yes	Yes

Example 3

Let's refer to $\phi \vee \psi \vee \theta$ as *the conclusion*. Suppose that everyone grounds their judgment on the conclusion on their judgment on ϕ, ψ and θ (the premises, in the standard approach). What is striking about Nehring's example is that, although the judges unanimously support the conclusion, accepting it would conflict with the majority verdict on the premises.

In Example 3, Premise-Based Majority does not merely decline to accept the unanimously accepted disjunction. It actually accepts its negation. Conclusion-Based Majority respects Unanimity (at least on the conclusion), but at the cost of completely ignoring the judgments on the reasons. Some Distance-Based rules[10] decline to accept the conclusion (hence violate the unanimity principles) but stop short of accepting its negation: they produce *incomplete outputs*. These three verdicts can be summarized in the following conditional statement:

Reasons vs. Unanimity
If the rule is sensitive to the disagreement on the reasons, then the unanimity on the conclusion is not a sufficient basis for accepting it.

My third task in the paper will be to define natural rules that are sensitive to reasons (in a sense to be made precise) while respecting *propositions-wise unanimity* on all propositions (hence these rules will behave like Conclusion-Based Majority in Example 3). Such rules can also be defined in the standard framework, but that they are more naturally defined and motivated within my generalization.

1.4. Relevant Work.

Miller (2008) discusses a formal model that is specifically designed to handle variability in what judges take to be their reasons.[11] In Miller's model, judges are assigned opinions drawn from an agenda (as in standard judgment aggregation), but in addition each judge may have

[10] In particular the rule in Pigozzi (2008).
[11] I thank an anonymous referee for pointing me to Miller's paper.

a different perspective on how one should decide the vote (what Miller calls a 'decision rule').

Despite these similarities, there are three major differences between the approach I develop here and Miller's. First, Miller *holds* the conclusion fixed, and supposes that the judges may disagree on how to settle it. I refrain from holding the conclusion fixed. There are applications of Judgment Aggregation for which there is no settled fact of the matter as to what should count as a conclusion. For example Pettit (2006), who otherwise looks favorably on aggregation methods appealing to the Premise-Conclusion dichotomy, argues convincingly that in some cases of modeling deference to a group of experts it is unattractive to prioritize some propositions over others.

Secondly, Miller and I have very different goals: he is concerned with generalizing some impossibility results and more generally with questions surrounding what I have called the *coherence challenge*. I am concerned with setting up a formal framework that can satisfy some philosophically motivated constraints related to the *reasons challenge*.

Finally–partly as as result of these differences–Miller's impossibility results require both Independence and Completeness. Within a reasons-sensitive framework of the kind that I have in mind, there are compelling reasons to reject both assumptions.[12] I discuss these reasons at various points in §§3-4: for now, I stress that my rejection of both conditions means that I work within a very different corner of logical space than Miller.

2. Formal Framework

I start by reviewing the standard aggregation framework. This is a familiar drill, but it's worth repeating. After that, I present my generalization:

2.1. Standard Framework

Start with a *modeling language* $\mathcal{L}$ (a standard sentential language), generated by a finite set of atomic sentences $\mathcal{L}_0$.[13] Let $\mathcal{G}$ be a finite *group* of advisors/judges/voters: for simplicity, identify $\mathcal{G}$ with the set $\{1, ..., n\}$;

[12] There are other compelling reasons, both technical and conceptual, to drop Independence, regardless of which framework we're working in (Chapman, 2002; Mongin, 2005; Cariani, Pauly and Snyder, 2008).

[13] Notice that we use *sentences* of a formal language to stand for *propositions* (the contents of agent's beliefs). The potential confusions that this engenders are dodged

n is typically assumed to be odd (until noted, I adopt this restriction). The *agenda* $\mathcal{I}$ is a subset of $\mathcal{L}$ that is closed under negation (i.e. if $\phi \in \mathcal{I}$, then $\sim\phi \in \mathcal{I}$). A *judgment set* j is a subset of $\mathcal{I}$. An *epistemic state* is a maximally consistent judgment set. A *profile* $\langle j_1, ..., j_n \rangle$ is a vector of epistemic states (I write this as $\vec{j}$). While referring to individuals, I will informally say that i accepts (viz. rejects) ϕ, meaning that $\phi \in j_i$ (viz. $\phi \notin j_i$). An *aggregation rule* $\mathcal{A}$ is a function[14] from profiles to judgment sets.[15]

For future reference, we need to define two important properties of aggregation rules, Completeness and Independence.

DEFINITION 1 (Completeness). *$\mathcal{A}$ is complete (relative to $\mathcal{I}$) iff for every $\vec{j}$, and every $\phi \in \mathcal{I}$, $\phi \in \mathcal{A}(\vec{j})$ or $\sim\phi \in \mathcal{A}(\vec{j})$.*

Independence requires a preliminary definition: $\vec{j}$ and $\vec{k}$ are **ϕ-matching** iff for all i, $\phi \in j_i \Leftrightarrow \phi \in k_i$. Informally, two profiles are ϕ-matching if the vectors of opinion on ϕ are identical. Next:[16]

DEFINITION 2 (Independence). *$\mathcal{A}$ is independent iff for every $\phi \in \mathcal{I}$ and any two profiles $\vec{j}$ and $\vec{k}$ that are **ϕ-matching***

$$\phi \in \mathcal{A}(\vec{j}) \Leftrightarrow \phi \in \mathcal{A}(\vec{k})$$

2.2. Enriched Framework

I argued in §1 that this framework's representation of epistemic states is too coarse to correctly model the role of reasons. To make the framework finer, I adopt a richer representation of individual epistemic

---

by the fact that we are focusing on context-insensitive sentences and that typically aggregation rules treat logically equivalent sentences in exactly the same way.

[14] I suppose that these functions are defined on all possible profiles, thus effectively building a Universal Domain assumption in the concept of an aggregation rule.

[15] **Further Notation:**

- $\mathcal{A}$ (sometimes with numerical subscripts) is used as a metalinguistic variable over aggregation functions
- B, C, D, E, F, G are used (as needed) as object language atomic sentences (i.e. sentences of $\mathcal{L}_0$). These always appear with a specified intended interpretation. A is *not used* in this way to avoid confusion with $\mathcal{A}$ that I use as a variable over aggregation rules.
- ϕ, ψ, θ are metalinguistic variables ranging over propositions in the object language of any complexity;
- M is a metalinguistic variable that ranges over subsets of $\mathcal{G}$.

[16] This is not the most succinct way of presenting the Independence assumption, but it is useful for the purposes of my later revisions.

states: i's epistemic state J_i (note the capitalization: J_i refers to an epistemic state in my enriched framework, j_i refers to an epistemic state in the standard framework, i.e. a maximally consistent judgment set) is a pair of the form:

$$\langle Opinions_i, Reasons_i \rangle$$

where $Opinions_i$ is a maximally consistent judgment set and $Reasons_i$ is a relation on $\mathcal{P}(\mathcal{I}) \times \mathcal{I}$ representing claims about how i bases his/her opinions on other opinions.[17] To streamline the presentation, I use a piece of notation for reasons-relations: I write $\Sigma \hookrightarrow_i \phi$ to mean that $\langle \Sigma, \phi \rangle \in Reasons_i$ (this notation does *not* belong to the formal language $\mathcal{L}$ itself: it's an abbreviation for claims in my set-theoretic metalanguage).

If a judge i accepts ϕ and no reasons are officially recorded in $Reasons_i$, then either i accepts ϕ non-inferentially, or on the basis of propositions that lie outside of the agenda. Judges can also accept a proposition for multiple independent reasons. For example, i may accept $\phi \vee \psi$ on the basis of independent beliefs in the disjuncts: in these cases, I write $\{\phi\} \hookrightarrow_i \phi \vee \psi$ and $\{\psi\} \hookrightarrow \phi \vee \psi$.[18]

Logical entailment from Σ to ϕ is neither necessary nor sufficient for $\Sigma \hookrightarrow_i \phi$ to hold.[19] Furthermore, I do not take logical complexity to be a guide to the reasons-relation: atomic sentences have no more of a right to be called 'reasons' than complex ones.[20] All I require of the reasons-relation is:

Acceptance: If $\Sigma \hookrightarrow_i \phi$, then i accepts ϕ and every member of Σ.[21]

Acceptance together with the requirement of logical consistency on individual opinion implies:

[17] A complete story about collective reasons should embed an account of reasons that captures more of their structure and dynamics—e.g. the *default theories* discussed by Horty (2007a), (2007b). However, for my current purposes, such a structured level of analysis is not necessary.

[18] Contrast with $\{\phi, \psi\} \hookrightarrow_i \phi \vee \psi$, which means something different—namely, that each of ϕ and ψ is only a *part* of the agent's reasons for $\phi \vee \psi$.

[19] It is not necessary because agents can accept ϕ on the basis of a proposition ψ that only supports ϕ inductively. It is not sufficient because otherwise we would be forced to have both $\{\phi \,\&\, \psi\} \hookrightarrow_i \phi$ and $\{\phi, \psi\} \hookrightarrow_i (\phi \,\&\, \psi)$—which would introduce a kind of circularity in the reasons-relation.

[20] This is a point on which I differ from Mongin (2005).

[21] One could also explore a variant of this framework without Acceptance. The reasons-relation would not capture what reasons an agent *actually* has, but what patterns of reasons an agent recognizes. That is, those patterns such that *if* i accepted their 'premises' then i would take those premises to be reasons to accept the conclusion. This would be a rather different framework from what I propose here.

Consistency: for all $i \in \mathcal{G}$, $(\Sigma \hookrightarrow_i \phi)$ only if $\Sigma \cup \{\phi\}$ is consistent.

As in the standard framework, a profile $\vec{J}$ is a vector of epistemic states and an aggregation rule $\mathcal{A}$ maps profiles to pairs of the form

$$\langle Opinions_c, Reasons_c \rangle$$

(just as in the standard framework, this pair need not qualify to be an epistemic state because $Opinions_c$ may be incomplete or inconsistent). For reference, call the first element of this pair $Opinions[\mathcal{A}(\vec{J})]$ and the second $Reasons[\mathcal{A}(\vec{J})]$. Similar notation is also useful for individual states: given $\vec{J}$, I write $Opinions_{Ji}$ for the judgment set in the i-th coordinate of $\vec{J}$ (similarly for $Reasons_{Ji}$). Finally, it will be helpful to be able to retrieve the vector of opinions from a profile $\vec{J}$, defining:

$$\vec{Opinions}[\vec{J}] = \langle Opinions_{J1}, ..., Opinions_{Jn} \rangle$$

$$\vec{Reasons}[\vec{J}] = \langle Reasons_{J1}, ..., Reasons_{Jn} \rangle$$

Given any two vectors $\vec{v}_1$ and $\vec{v}_2$, $\vec{v}_1 = \vec{v}_2$ iff they are point-wise identical. I call this enriched framework Rich.

2.3. Representation

Rich generalizes the standard aggregation framework in the following sense: every 'standard' aggregation rule corresponds to a class of rules in Rich. Consider an aggregation rule $\mathcal{A}$ in Rich. $\mathcal{A}$ maps a vector of pairs of the form $\langle Opinions_{Ji}, Reasons_{Ji} \rangle$ to a pair of the same form. Some rules, however, are insensitive to the value of the *Reasons* coordinate: their output (at least in its *Opinions* part) is completely determined by the *Opinions* coordinate of the individual states. I say that these rules 'represent' the standard rules in Rich.

DEFINITION 3 (Representation).
A rule $\mathcal{A}_1$ in Rich *represents a standard rule $\mathcal{A}_2$ iff for all $\vec{J}$ and $\vec{j}$ s.t. $\vec{Opinions}[\vec{J}] = \vec{j}$,* $\text{Opinions}[\mathcal{A}_1(\vec{J})] = \mathcal{A}_2(\vec{j})$

The following, then, is evidently true:

(a) For every standard rule $\mathcal{A}$, there is a rule in Rich that represents $\mathcal{A}$.

(b) There are rules available in Rich that do not represent any standard rule.[22]

[22] I provide some examples in the next section.

This captures a minimal sense in which RICH is more general than the standard framework. Every rule that we can define in the standard framework is represented in RICH.

The additional expressive power does not help with respect to the *coherence challenge*. Many of the impossibility results in the standard framework can be replicated in RICH by lifting the relevant conditions. For example, to lift Independence, we need only modify the definition of what it is to be ϕ-matching. $\vec{J}$ and $\vec{K}$ are ϕ-matching (in the revisited sense) iff for all i, $\phi \in Opinions_{Ji} \Leftrightarrow \phi \in Opinions_{Ki}$

DEFINITION 4 (Independence (revisited)). *$\mathcal{A}$ is independent iff for any two profiles $\vec{J}$ and $\vec{K}$ that are ϕ-matching (in the revisited sense),*

$$\phi \in Opinions[\mathcal{A}(\vec{J})] \Leftrightarrow \phi \in Opinions[\mathcal{A}(\vec{K})]$$

The following simple result (proof in the appendix) connects the revisited and the original notion of Independence:

THEOREM 1. *If $\mathcal{A}_1$ is independent (in the revisited sense), then $\mathcal{A}_1$ represents a standard rule $\mathcal{A}_2$ that is independent (in the standard sense).*

It is immediately clear that if standard Independence is incompatible with other properties, then the revisited Independence is also incompatible with the revisited versions of those properties. Insofar as standard independence is involved in many impossibility results, this result should cast skepticism on the prospects of saying something new about the *coherence challenge*. This is unfortunate, but not really counter to my motivation: we must ask how RICH can help with the *reasons challenge*.

3. RICH and Reasons.

One of the most significant complaints against Independence is that the individual pattern of acceptance on ϕ is too thin a basis to settle whether or not ϕ should be collectively accepted. What could be the motivation for aggregating over multiple propositions if we are then splitting the aggregation problem into separate, isolated, chunks (Chapman, 2002; Mongin, 2005)?

In my interpretation, this complaint hinges on the intuition that it makes a difference if the judges that support ϕ are *cohesive* (that is, support ϕ for the same reason, or mutually reinforcing reasons, or at least compatible ones) rather than *non-cohesive* (they support ϕ for incompatible or mutually undermining reasons).[23]

3.1. Cohesive Rules

First, we want to define what it is for a group of judges to support a proposition cohesively.

DEFINITION 5 (Strong Cohesiveness). *$M \subseteq \mathcal{G}$ strongly cohesively supports ϕ (on $\vec{J}$) iff there is a set Σ of propositions such that:*

(i) every member of M accepts (on $\vec{J}$) every member of Σ as well as ϕ and

(ii) for every member i of M, $\Sigma \hookrightarrow_i \phi$ (on $\vec{J}$).

I call this notion *strong* cohesiveness, because a more intuitive notion of cohesiveness does not require that every judge accept ϕ for precisely the *same* reasons. Arguably, something weaker is enough—namely that the reasons not be mutually undermining. In §5, I explore how to produce a more liberal definition of cohesiveness. For now, it is easier to proceed with this simplified definition.

Next, we can define a class of aggregation rules that are responsive to the difference between cohesive and non-cohesive support. $\mathcal{A}$ is a *Cohesive Rule* just in case *Opinions*$[\mathcal{A}(\vec{J})]$ can be characterized by setting the parameters in the following schema:

DEFINITION 6. $\phi \in \text{Opinions}[\mathcal{A}(\vec{J})]$ *iff there is a set $M \subseteq \mathcal{G}$, such that M strongly cohesively supports ϕ (on $\vec{J}$) and*

$$\frac{|M|}{|M \cup \textsf{\textit{contrast}}_{\phi}(\vec{J})|} > \textsf{\textit{threshold}}_{\phi}(\mathcal{I})$$

$\textsf{threshold}_{\phi}$ and $\textsf{contrast}_{\phi}$ are parameters (different Cohesive Rules might assign different values to them) subject to the following constraints:

– $\textsf{threshold}_{\phi}$ is a function mapping agendas into [1/2, 1).[24]

[23] At the very least, non-cohesiveness should *sometimes* undermine apparent agreement.

[24] In principle, the parameters could behave differently on different formulas. I want to keep the option open, although for this paper, I won't take advantage of this option.

– $\mathsf{contrast}_\phi$ is a function that maps profiles (e.g. $\vec{J}$) into $\mathcal{P}(\mathcal{G})$ subject to the following additional condition:
  (#) If $Opinions[\vec{J}] = Opinions[\vec{K}]$, then $\mathsf{contrast}_\phi(\vec{J}) = \mathsf{contrast}_\phi(\vec{K})$

Let me give some examples. I will start by considering rules with a majority threshold—i.e. rules with $\mathsf{threshold}_\phi$ constant at 1/2 (for all agendas).

One possible way of defining $\mathsf{contrast}_\phi$ is by letting it map $\vec{J}$ to the entire group $\mathcal{G}$.[25] Call the resulting rule *cohesive majority* (or $\mathcal{CM}$ for short).

Cohesive Majority: $\phi \in Opinions[\mathcal{CM}(\vec{J})]$ iff there is a set $M \subseteq \mathcal{G}$, such that M strongly cohesively supports ϕ (on $\vec{J}$) and

$$\frac{|M|}{|M \cup \mathcal{G}|} = \frac{|M|}{|\mathcal{G}|} > \frac{1}{2}$$

Informally, $\mathcal{CM}$ accepts ϕ on a profile $\vec{J}$ iff there is a majority of judges M that supports ϕ and does so (strongly) cohesively.

Without changing the threshold, an interesting alternative can be defined by letting $\mathsf{contrast}_\phi$ map $\vec{J}$ to the set

$$\{i \in \mathcal{G} \mid i \text{ accepts } \sim\phi \text{ on } \vec{J}\}$$

(For reference, call this set $\mathcal{G}^{\sim\phi,\vec{J}}$). The result is:

Discounting Majority: $\phi \in Opinions[\mathcal{DM}(\vec{J})]$ iff there is a set $M \subseteq \mathcal{G}$, such that M strongly cohesively supports ϕ (on $\vec{J}$) and

$$\frac{|M|}{|M \cup \mathcal{G}^{\sim\phi,\vec{J}}|} > \frac{1}{2}$$

Since $M \cap \mathcal{G}^{\sim\phi,\vec{J}} = \emptyset$, the threshold condition can be more vividly (but equivalently) stated as:

$$|M| > |\mathcal{G}^{\sim\phi,\vec{J}}|$$

The intuition behind $\mathcal{DM}$ is that we want to first find the largest set of cohesive supporters of ϕ (or one of the largest, if there are ties). Having found such a set, we check that it outnumbers the set of *opposers* of ϕ. Intuitively, we are discounting those judges whose acceptance of ϕ is

[25] All we really need for the following rule is that $\mathsf{contrast}_\phi$ map $\vec{J}$ to any superset of $\mathcal{G} - M$. However, since $\mathsf{contrast}_\phi$, as I defined it, does not take M as a parameter, we can't just set $\mathsf{contrast}_\phi = \mathcal{G} - M$.

not cohesive with the judges in M, and then checking whether M is large 'enough'.

$\mathcal{CM}$ is more stringent than $\mathcal{DM}$.[26] Both rules are more stringent than simple majority (e.g. if a cohesive majority supports ϕ, then a majority supports ϕ, but the converse is not true).

I have elaborated how $\mathcal{CM}$ and $\mathcal{DM}$ derive collective opinions. But aggregation rules in Rich must also produce, where available, a pattern of collective reasons. Cohesive Rules in general (and $\mathcal{CM}/\mathcal{DM}$ in particular) allow for a natural account of collective reasons. Notice that Cohesive Rules require for acceptance of ϕ that there be a *cohesive* group of ϕ-supporters. Let $\hookrightarrow_c$ be the collective reasons relation, and $\mathcal{A}$ an arbitrary Cohesive Rule.

DEFINITION 7 (Collective Reasons). $\langle \Sigma, \phi \rangle \in Reasons[\mathcal{A}(\vec{J})]$ *iff*

(i) $\phi \in Opinions[\mathcal{A}(\vec{J})]$

(ii) $\Sigma \subseteq Opinions[\mathcal{A}(\vec{J})]$

(iii) there are no ties for "largest subset of $\mathcal{G}$ *that cohesively supports* ϕ *(on* $\vec{J}$*)".*

(iv) for every i *in the largest subset of* $\mathcal{G}$ *that cohesively supports* ϕ, $\Sigma \hookrightarrow_i \phi$.

The first two conditions require that the appropriate propositions belong to the aggregated opinion. In this way, group reasons satisfy the Acceptance requirement.

One of the most interesting features of Cohesive Rules is precisely that they allow this sort of account of collective reasons. Moreover, this account delivers some intuitively correct veridicts about what should count as the group's reasons (see §4).

Neither $\mathcal{CM}$ nor $\mathcal{DM}$ is *consistency-preserving*. On those profiles on which every judge i accepts every proposition non-inferentially, both rules collapse onto the Majority rule. Since the Majority rule is not guaranteed to be consistency-preserving, neither are $\mathcal{CM}$ and $\mathcal{DM}$. Consider for instance a version of the doctrinal paradox (Example 1) with a profile $\vec{J}$ in which for all $i \in \mathcal{G}$, $Reasons_{Ji} = \emptyset$ In such a case, $\mathcal{MA}$, $\mathcal{CM}$ and $\mathcal{DM}$ all yield the same verdicts.

[26] The two rules only differ in $M \cup \mathsf{contrast}_\phi$. The difference is, for $\mathcal{CM}$, this is the entire $\mathcal{G}$, while, for $\mathcal{DM}$, sometimes it is a proper subset of $\mathcal{G}$. Hence there are cases in which $|M|/|M \cup \mathsf{contrast}_\phi|$ is large enough according to $\mathcal{DM}$ but not according to $\mathcal{CM}$.

In §4, I show that, restricted to some appropriately defined classes of profiles, $\mathcal{CM}$ is consistency-preserving and in fact coincides with a version of Premise-Based Majority. Still, there is an interesting question of which cohesive rules are consistency-preserving across *all* profiles; I turn to that question now.

3.2. Cohesive Supermajority

Pettit (2006) informally describes a standard super-majority rule that is guaranteed to produce a consistent output. This is a supermajority rule in which the acceptance threshold varies with some global logical properties of the agenda. List (2007) analyzes this insight formally (based on results in Dietrich and List, 2007) and observes that the threshold $t_{\mathcal{I}}$ must be at least $x - 1/x$ where x is the size of the largest minimally inconsistent subset of the agenda.[27]

This threshold also works to make Cohesive Rules consistency-preserving. Let $\mathsf{threshold}_{\phi}$ be the function that maps each agenda $\mathcal{I}$ onto $t_{\mathcal{I}}$ defined as in the previous paragraph. Then we can define two new rules $\mathcal{CS}$ and $\mathcal{DS}$ according to the two ways of specifying the value of $\mathsf{contrast}_{\phi}$

Cohesive Supermajority: $\phi \in Opinions[\mathcal{CS}(\vec{J})]$ iff there is a set $M \subseteq \mathcal{G}$, such that M strongly cohesively supports ϕ and

$$\frac{|M|}{|\mathcal{G}|} > t_{\mathcal{I}}$$

Discounting Supermajority: $\phi \in Opinions[\mathcal{DS}(\vec{J})]$ iff there is a set $M \subseteq \mathcal{G}$, such that M strongly cohesively supports ϕ and

$$\frac{|M|}{|M| + |\mathcal{G}^{\sim\phi,\vec{J}}|} > t_{\mathcal{I}}$$

Both are cohesive, consistency-preserving and can inherit the account of collective reasons I described earlier. Both rules violate Completeness and Independence (see the appendix for examples of violations). I take these to be design features rather than problems. In any case, we must drop Completeness once we make harmless generalizations such as:

(a) dropping the restriction that the numbers of judges must be odd

[27] For any thresholds below this value we can construct standard profiles on which the appropriate supermajority rules would be inconsistent.

(b) dropping the irrealistic assumption that individual judges must have complete beliefs (relative to the agenda).

Furthermore, the requirement of cohesive acceptance introduces points of incompleteness anyway.[28]

One may worry that $\mathcal{DS}$ and $\mathcal{CS}$ aren't just Incomplete, but also fail to be *deductively closed*. List (2008) advances some objections to the supermajoritarian proposal of Pettit (2006) that trade on this point. I fail to see the bite of these complaints: provided that Σ is consistent, one can simply take the deductive closure $Th(\Sigma)$ and let that be the output of the aggregation rule.

As for Independence, Cohesive Rules do not satisfy it by design. Crucially, however, all Cohesive Rules (and hence in particular $\mathcal{CS}$ and $\mathcal{DS}$) satisfy a weaker version of Independence. As usual, we first define a relation between profiles: $\vec{J}$ and $\vec{K}$ are *deeply ϕ-matching* iff for every judge i,

(i) $\phi \in Opinions[J_i] \Leftrightarrow \phi \in Opinions[K_i]$ and

(ii) for every set of propositions Σ, $\Sigma \hookrightarrow_i \phi$ holds in $\vec{J}$ iff it holds in $\vec{K}$.

DEFINITION 8 (Weak Independence). *$\mathcal{A}$ is weakly independent iff for every proposition ϕ and any two profiles $\vec{J}$ and $\vec{K}$ that are deeply ϕ-matching,*

$$\phi \in Opinions[\mathcal{A}(\vec{J})] \Leftrightarrow \phi \in Opinions[\mathcal{A}(\vec{K})]$$

THEOREM 2. *All Cohesive Rules are Weakly Independent.*

In the next section I discuss the significance of this theorem and why it matters with regards to the concerns that motivated my framework.

The final result I want to mention is that there is a useful characterization of $\mathcal{DS}$ (at least among the Cohesive Rules). This characterization makes vivid how $\mathcal{DS}$ is special.[29] Say that a Cohesive Rule is:

[28] See Gärdenfors (2006) for further resistance against Completeness in judgment aggregation. If one really wanted, one can always restore Completeness with some kind of further procedure that can break ties, as suggested in Pigozzi (2008). In my view, any such procedure can only be justified on pragmatic grounds (e.g. the urgency of making a collective decision), and not on principled logical or epistemological grounds. Dietrich and List (2008) show that a core of impossibility results is still available if we relax completeness to deductive closure. Their results, however, still presuppose Independence, which I reject.

[29] I do not claim on the basis of the following characterization that $\mathcal{DS}$ is in an absolute sense better than other rules. For example, whether or not we should require that an aggregation rule be anonymous essentially depends on the modeling application. What I claim, rather, is that $\mathcal{DS}$ sits at the intersection of several important properties and that these properties can be simultaneously desirable.

Good iff for all $\vec{J}$, $\mathcal{G}^{\sim\phi,\vec{J}} \subseteq \mathsf{contrast}_\phi(\vec{J})$

Threshold-Neutral iff for all $\phi, \psi \in \mathcal{I}$, $\mathsf{threshold}_\phi(\vec{J}) = \mathsf{threshold}_\psi(\vec{J})$

Proper iff it is Good, Anonymous, Consistency-preserving, and Threshold-Neutral.

Additionally, letting C be an arbitrary class of rules, define:

DEFINITION 9 (Maximality). *$\mathcal{A}$ is maximal in C iff $\mathcal{A} \in C$ and for every $\mathcal{I}$, $\phi \in \mathcal{I}$ and $\vec{J}$,*

$$\phi \in \mathit{Opinions}[\mathcal{A}(\vec{J})] \iff \text{there is a rule } \mathcal{A}' \text{ in } C,\ \phi \in \mathit{Opinions}[\mathcal{A}'(\vec{J})].$$

THEOREM 3. *$\mathcal{DS}$ is maximal in the class of Proper Cohesive Rules.*

That is to say that $\mathcal{DS}$ is the 'closest' to Completeness one can get within the class of Proper Cohesive Rules. Even though Completeness does not seem desirable on purely logical/epistemic grounds, it would be nice to have rules that are more liberal than $\mathcal{DS}$. Theorem 3 shows that this is impossible without stepping outside of the class of Proper Cohesive Rules.[30]

4. Rich in Action.

In this section I address some of the philosophical issues surrounding the framework I have just developed, including how to deal with the motivating points of the paper.

4.1. Weak Independence

The significance of Theorem 2 is that it allows us to retrieve some of the advantages of Independence, without some of its costs. Standard Independence is sometimes glossed as the claim that the aggregated outcome on ϕ depends only on the individual opinions on ϕ *and not on opinions on any other proposition* ϕ. There is something persuasive about the idea that a judgment on ϕ shouldn't be sensitive to irrelevant information. However, nothing in the standard definition of Independence characterizes 'relevant information'. If interpreted as ruling out 'irrelevant information', Independence would imply the absurd claim that no other propositions can be relevant to ϕ (a point

[30] I discuss a couple of suggestions in §4.4 and §5

also highlighted by Dietrich, 2008). Our slogan should rather be that the aggregated outcome on ϕ should depend only on the opinions on ϕ and on whatever other propositions individual agents take to be relevant to ϕ. The distinctive feature of RICH is that it involves a *subjective* interpretation of 'relevance': what is relevant to ϕ is simply what individual judges base opinions for or against ϕ on.

As Douven and Romeijn (2007) point out, there is some reason to explore rules satisfying Independence: the prospects of saying something substantive and general about aggregation rules in the absence of Independence conditions may appear bleak. This is changing as part of a general trend towards the study of non-Independent rules (the trend is witnessed by the attention to distance-based approaches). RICH provides a terrain on which certain kinds of intuitive non-independent aggregation rules can be mapped out and studied precisely.

I must add that Independence conditions are not simply justified by the tractability that they allow. One of the central results in the standard framework (Dietrich and List, 2007b) shows that Independence is a necessary condition of Non-Manipulability (the requirement that the outcome of the vote not be manipulable by insincere voting). An important direction of investigation for the current approach is to explore analogues of this result for the condition of Weak-Independence. I defer these results and their interpretation to future work.

4.2. The Three Tasks

I had set three goals for the framework. The first was to generalize the standard approach to reasons without assuming that reasons-relations are fixed externally. The second was to define rules that can assign non-isomorphic inputs to STOCKBROKERS from LITTLE KIDS, and treat them differently. The third was to show that Proposition-wise Unanimity Preservation can be implemented within reasons-sensitive rules. In this section, I discuss (in reverse order) how we fared on each of these goals.

I have little to say about the first goal, it is evident that we have constructed a framework in which we do not need to regard reasons relations as externally fixed. Rather, collective reasons-relations arise out of individual reasons together with my account of Collective Reasons (which is applicable at least to the Cohesive Rules).

Premise-Based approaches are sensitive to reasons but fail to preserve unanimity. Conclusion-Based approaches satisfy Propositions-wise Unanimity (at least on the conclusion), but are reasons-insensitive.

$\mathcal{DM}$ and $\mathcal{DS}$ are responsive to reasons while at the same time satisfy Propositions-wise Unanimity (on all propositions).

Moving on to the second task, recall the logical structure which I ascribed to both STOCKBROKERS and LITTLE KIDS:

	ϕ	ψ	$\phi \vee \psi$
1	Yes	No	Yes
2	No	Yes	Yes

Example 2 (repeated)

The difference between the two cases is that in one case (STOCKBROKERS) the disjunction is based on the disjuncts (so we should add the further facts: $\{\phi\} \hookrightarrow_1 \phi \vee \psi$ and $\{\psi\} \hookrightarrow_2 \phi \vee \psi$). In LITTLE KIDS these basing facts do not apply (all propositions are based at least in part on opinions that lie outside of the agenda). In this way, $\mathcal{CS}$ can accept the disjunction in LITTLE KIDS, without accepting it on STOCKBROKERS (this is a point of difference between $\mathcal{CS}$ and $\mathcal{DS}$, since the latter satisfies Propositions-wise Unanimity).

The last point was whether Propositions-wise unanimity was compatible with a reasons-sensitive view. We now have a clear understanding of what it is for a rule to be reasons-sensitive: it means that the rule satisfies Weak Independence (individual opinions on ϕ and reasons in support of ϕ fix collective opinions) but not Strong Independence (individual opinions on ϕ alone do not fix collective opinions on ϕ). All Cohesive Rules are reasons-sensitive in this sense. Additionally, $\mathcal{DM}$ and $\mathcal{DS}$ satisfy Proposition-wise Unanimity (for any ϕ, if all judges accept ϕ, the group accepts ϕ). So the discounting rules occupy a middle ground between Conclusion-Based approaches and the Premise-Based approaches.

4.3. GENERALIZING PREMISE BASED MAJORITY.

There is a further payoff in taking the perspective I suggested: some Cohesive Rules in RICH generalize familiar rules in the standard framework. In particular, despite its failure to preserve consistency, $\mathcal{CM}$ generalizes Premise-Based Majority (which, in this section, I abbreviate as *PB*).

When I say 'generalize' I mean that in a restricted, but natural, class of profiles $\mathcal{CM}$ and *PB* agree on their verdicts. These profiles are exactly those profiles on which the judges agree that the reasons-relation flow from the atomic sentences (and their negations) to the conclusion.

Let $\mathcal{I}$ be an agenda with propositions $\{\phi_1, ..., \phi_z, \psi, \textit{negations}\}$ (here the ϕ_i's range over atomic sentences). Let $\textit{Premises}(\mathcal{I})$ denote all atomic sentences and negations of atomic sentences in $\mathcal{I}$. Let ψ be the conclusion and let it be related in the appropriate way to $\{\phi_1, ..., \phi_k\}$, that is, let it have at least these two properties:

(i) Any distribution of truth-values on $\textit{Premises}(\mathcal{I})$ determines truth-functionally the truth-value of ψ.

(ii) For every premise ϕ_x there is a distribution of truth-value d on the other premises, such that d together with $\phi_x = T$ forces the conclusion to be true and d together with $\phi_x = F$ forces the conclusion to be false or vice-versa (informally: no premise is superfluous).

Consider an individual epistemic state J_i. Let $|\psi|_i$ be whichever of ψ and $\sim\psi$ is accepted in J_i. The notion of minimal entailment is required here: Σ minimally entails ϕ iff $\Sigma \models \phi$ and there is no proper subset Σ' of Σ such that $\Sigma' \models \phi$. Let:

$$\Pi_i = \{\Sigma \subseteq \mathcal{I} \mid \Sigma \text{ minimally entails } |\psi|_i\}$$

DEFINITION 10 (Premise-centered).
J_i *is* premise-centered *iff For all* $\Sigma \subseteq \mathcal{I}, [\Sigma \hookrightarrow_i \phi \Leftrightarrow \Sigma \in \Pi_i]$[31]
$\vec{J}$ *is* premise-centered *iff for every i,* J_i *is premise-centered.*

In other words, premise-centered states are those in which the reasons-relation mimicks the relation between premises and conclusions.

THEOREM 4. *If* $\vec{J}$ *is* premise-centered*, then:*

$$\phi \in \text{Opinions}[\mathcal{CM}(\vec{J})] \textit{ iff } \phi \in PB(\text{Opinions}[\vec{J}])$$

Informally, this means that, when the profile is premise-centered, the propositions accepted by $\mathcal{CM}$ are exactly those propositions that would be accepted by $\mathcal{PB}$ on the equivalent profile.

[31] Some examples can help illustrate this definition. Consider the agenda $\{A, B, A \,\&\, B\}$, and suppose that $A \,\&\, B$ is designated as the conclusion. Here are some example of premise-centered epistemic states:

- $\{A, B, A \,\&\, B, \{A, B\} \hookrightarrow (A \,\&\, B)\}$
- $\{\sim A, B, \sim(A \,\&\, B), \{B\} \hookrightarrow \sim(A \,\&\, B)\}$
- $\{\sim A, \sim B, \sim(A \,\&\, B), \{\sim A\} \hookrightarrow \sim(A \,\&\, B), \{\sim B\} \hookrightarrow \sim(A \,\&\, B)\}$

Moreover, *Reasons*[$\mathcal{CM}(\vec{J})$] turns out to be exactly what we should expect. If $\phi \vee \psi$ is accepted because a cohesive majority accepts $\{\phi, {\sim}\psi\}$, we have that every i in that cohesive majority has $\{A\} \hookrightarrow_i A \vee B$, and hence that $\{A\} \hookrightarrow_c A \vee B$. Theorem 4 captures the sense in which Premise-Based Majority is a special case of $\mathcal{CM}$, the case in which the reasons-relation flows from the atomic sentences to the conclusion.

4.4. Extending $\mathcal{DS}$.

Theorem 3 shows that $\mathcal{DS}$ is the Proper Cohesive Rule that most closely approaches Completeness. It seems desirable to find some ways of extending $\mathcal{DS}$ to produce more liberal rules.[32]

In this connection, it is natural to take a more liberal position with respect to the thresholds required for acceptance. Consider this example:

	ϕ	ψ	$\phi \vee \psi$	Reasons
1	Yes	Yes	Yes	$\{\phi\} \hookrightarrow_1 \phi \vee \psi$
2	Yes	Yes	Yes	$\{\phi\} \hookrightarrow_2 \phi \vee \psi$
3	No	No	No	$\{{\sim}\phi, {\sim}\psi\} \hookrightarrow_3 {\sim}(\phi \vee \psi)$

Example 4

In Example 4, we have a cohesive majority for $\phi \vee \psi$, but not a cohesive super-majority. In particular, while $\mathcal{DM}$ recommends acceptance of $\{\phi, \psi, \phi \vee \psi\}$, $\mathcal{DS}$ recommends suspending on all propositions. Still $\mathcal{DM}$ is unattractive because it isn't consistency-preserving on all profiles. One would want *both* consistency-preservation *and* a rule that is more liberal than $\mathcal{DS}$.

There are at least two possible alternative approaches here, both involving changing how we deal with the threshold. Currently $\mathsf{threshold}_\phi$ is a function from agendas to threshold. Instead we could treat it as a function of *both* the agenda $\mathcal{I}$ and the profile $\vec{J}$. The first approach aims to characterize the profiles on which the lower threshold (i.e. 1/2) is safe. Example 4 instantiates a natural sufficient condition for 'safety': there is a majority in favor of an entire maximally consistent judgment set. Clearly, this condition is not also necessary, however, and I do not know of necessary and sufficient conditions for 'safety' that can work here.

[32] 'Extending' here means defining a rule $\mathcal{A}$ such that for all agendas $\mathcal{A}$, profiles $\vec{J}$ and propositions $\phi \in \mathcal{A}$, $\phi \in$ *Opinions*[$\mathcal{DS}(\vec{J})$] $\Rightarrow \phi \in$ *Opinions*[$\mathcal{A}(\vec{J})$].

The other approach would be to adopt a rule on which we set the threshold (for an agenda $\mathcal{I}$ and given a profile $\vec{J}$) at the lowest value z that guarantees consistency (when the rule is given that threshold). So, in Example 4, it is set at 1/2, but on some other profiles it may be as high as $t_{\mathcal{I}}$. The fact that $t_{\mathcal{I}}$ gives us an upper bound for z is of course essential to guarantee that such a z must always exist. I won't develop either approach further in this paper, but they strike me as promising directions to improve the 'coverage' of the rules in RICH.

5. Amending Cohesiveness.

Strong cohesiveness is too strong. Two (or more) judges do not need to cite the exact same reasons for them to count as cohesive in their support of a proposition. Instead, we want to say that two (or more) judges can count as cohesive even if they adopt different reasons—provided that those reasons are not mutually undermining.[33]

In response to this problem, I suggest treating the notion of 'being mutual undermining' schematically. Consider an analogy: Judgment Aggregation makes extensive use of the concepts of consistency and entailment. These concepts can be analyzed in various ways, but many of the central facts of Judgment Aggregation do not depend on one particular analysis of consistency.[34] The same modular architecture can be used for 'being mutually undermining': the details of this concept are difficult and no doubt there are multiple possible explications. However, insofar as we can sketch broad structural features of these concepts and draw results out of them, we can leave the complexities for separate treatment.

Moreover, unlike the concepts of consistency and entailment, I am inclined to think that 'being mutually undermining' is in part dependent on features that are specific to the modeling application. These facts should make us doubt of the prospect of a purely logical rendering of what it means to be 'mutually undermining'. In light of this, in every specific modeling application, I introduce a relation $\sim_\phi$

[33] At the same time, Strong Cohesiveness might also be regarded as too weak—because it is only sensitive to the reasons for a particular proposition, but not to the reasons that support those reasons. Sam and Abe might support ϕ for the same reason ψ but be wildly non-cohesive in how they support ψ. I leave discussion of this generalization for separate treatment. But basically we might want to say that for j_1 and j_2 to cohesively support a proposition ϕ they must have cohesive *chains* of reasons that support ϕ.

[34] For a strong result to this effect, see Dietrich (2007).

(for each ϕ), holding between sets of reasons Σ and Σ' just in case they do not undermine each other(as far as support for ψ is concerned).

DEFINITION 11 (General Cohesiveness). *M cohesively supports ϕ iff there are sets $\Sigma_1, ..., \Sigma_k$ of propositions (drawn from $\mathcal{I}$) such that:*

(i) for all $i \in M$, i accepts every member of Σ_i, ϕ and $\Sigma_i \hookrightarrow_i \phi$

(ii) for each $i, j \in \{1, ..., k\}$, $\Sigma_i \sim_\phi \Sigma_j$.

In the special case in which $\sim_\phi$ is the identity relation for all ϕ, General Cohesiveness is obviously equivalent to Strong Cohesiveness. In defining Cohesive Rules, we can replace 'strong cohesiveness' with 'general cohesiveness' and obtain more liberal rules.

One cost of general cohesiveness is that it makes it more difficult to identify a collective reason. When I required Strong Cohesiveness for ϕ, I could, in most well-behaved cases, point to a unique set of collective reasons—the reason of the largest cohesive group of ϕ supporters (if there was one). Once we relax to General Cohesiveness, even if there is a largest cohesive group of ϕ supporters, we are not guaranteed to have a single reason on the basis of which they all support ϕ. The resulting problem is that we may not in general point to a collective reason for a given collective opinion.

I cannot offer a formal solution to this problem. Lest this discourages you, I do note, first, that the same problem applies also to the accounts of collective reasons that I have argued against. Second, that something more substantive (although not formal) about the problematic cases can be said. When a group of advisors M, Generally (but not Strongly) Cohesively supports ϕ, we can take the reasons for ϕ to be some set of propositions Σ' such that:

- for all $i \in M$, $\Sigma' \sim_\phi \Sigma_i$.
- for all $i \in M$, if i entertained Σ, and $\Sigma' \hookrightarrow \phi$, they would accept both

Such a Σ' may not exist in which case the group lacks a collective reason for their judgment.

Imagine a group of scientists each having their own datasets and data analyses pointing to the conclusion that ϕ. Each scientist may produce a slightly different justification for ϕ, so that they count as Generally but not Strongly Cohesive. The above proposal suggests that they will count as having a collective reason if there is a 'neutral' description of the experimental evidence that they are potentially disposed to accept and that they are disposed to take as supporting ϕ. For obvious reasons, a condition like this is hard to write into the formal framework. But it is a possible point at which we can find a natural role for collective deliberation.

6. Conclusion

I have argued that if we want an aggregation models to address the *reasons-challenge* (to explain how a group can come up with a collective judgments supported by reasons), we must adopt a model that represents reasons-relations already at the individual level. Collective reasons in favor of a verdict can be provided when judges are sufficiently cohesive and the Cohesive Rules I have defined and studied allow to express combinations of properties that are otherwise hard to capture (such as combinations of Weak Independence with other properties). The rules I have singled out are often incomplete, partly because they make collective acceptance quite demanding.

As a parting thought, it seems unlikely to interpret these rules as *voting* rules. They are best understood, in my view, as characterizing conditions at which a deliberating group can come to a collective verdict supported by collective reasons.

Appendix

Proof of Theorem 1

Proof. The theorem claims that if $\mathcal{A}$ is independent (in the revisited sense), then $\mathcal{A}$ represents a standard rule that is independent (in the standard sense). See section 2 for the salient definitions.

Suppose that $\mathcal{A}_1$ is independent (in the revisited sense). Let $\vec{J}$ and $\vec{j}$ be such that $\overrightarrow{Opinions}[\vec{J}] = \vec{j}$. Define a standard rule $\mathcal{A}_2$ by setting, for all standard profiles $\vec{j}$,

$$\mathcal{A}_2(\vec{j}) = \mathcal{A}_2(\overrightarrow{Opinions}[\vec{J}]) = Opinions[\mathcal{A}_1(\vec{J})]$$

We must check that $\mathcal{A}_2$ is indeed definable in this way. Particularly, that for any two profiles $\vec{J}$ and $\vec{K}$ in Rich,

(*i*)If $\overrightarrow{Opinions}[\vec{J}] = \overrightarrow{Opinions}[\vec{K}]$, then $Opinions[\mathcal{A}_1(\vec{J})] = Opinions[\mathcal{A}_1(\vec{K})]$

This is guaranteed by the fact that $\mathcal{A}_1$ is independent in the revisited sense. The assumption that the vectors of opinion are completely identical, means that $\vec{J}$ and $\vec{K}$ are ϕ-matching for any proposition in the agenda. Then obviously $\mathcal{A}_1$ must represent $\mathcal{A}_2$.

Additionally, we must show that $\mathcal{A}_2$ is independent in the standard sense. To do this, consider two ϕ-matching (in the standard sense) profiles $\vec{j}$ and $\vec{k}$. Now let $\vec{J}$ and $\vec{K}$ be any rich profiles such that:

(i) $\vec{Opinions}[\vec{J}] = \vec{j}$

(ii) $\vec{Opinions}[\vec{K}] = \vec{k}$

It follows immediately that $\vec{J}$ and $\vec{K}$ are ϕ-matching (in the revisited sense). Since $\mathcal{A}_1$ is independent in the revisited sense, this implies that,

$$\phi \in \mathcal{A}_2(\vec{j}) \text{ iff } \phi \in \mathit{Opinions}[\mathcal{A}_1(\vec{J})] \text{ iff } \phi \in \mathit{Opinions}[\mathcal{A}_1(\vec{K})] \text{ iff } \phi \in \mathcal{A}_2(\vec{k})$$

which implies that $\mathcal{A}_2$ is also independent. □

Violations of Completeness and Independence

Both $\mathcal{CS}$ and $\mathcal{DS}$ violate Completeness: on some profiles they fail to return a verdict on some propositions. For example, assuming that ϕ and ψ are distinct, we can have:

	ϕ	ψ	$\phi \vee \psi$	reasons-relation
1	Y	N	Y	$\{\phi\} \hookrightarrow_1 \phi \vee \psi$
2	Y	Y	Y	$\{\psi\} \hookrightarrow_2 \phi \vee \psi$
3	N	N	N	
$\mathcal{CM}, \mathcal{DM}$	Y	N	–	
$\mathcal{MA}$	Y	N	Y	

Example 3

In addition to violating Completeness, both rules also violate Independence. We can change the collective outcome on $\phi \vee \psi$ in Example 3 merely by intervening on the reasons-relation and adding $\{\phi\} \hookrightarrow_2 \phi \vee \psi$.

Proof of Theorem 2

Proof. Let $\vec{J}$ and $\vec{K}$ be profiles that are deeply ϕ-matching (see Definition 8), and let $\mathcal{A}$ be a Cohesive rule. Suppose $\phi \in$ *Opinions*$[\mathcal{A}(\vec{J})]$. Then there is a set $M \subseteq \mathcal{G}$ that strongly cohesively supports ϕ on $\vec{J}$ and satisfies the inequality:

$$\frac{|M|}{|M \cup \mathsf{contrast}_\phi(\vec{J})|} > \mathsf{threshold}_\phi(\mathcal{I})$$

For the purposes of this proof, the specific values of $\mathsf{contrast}_\phi(\vec{J})$ and $\mathsf{threshold}_\phi(\mathcal{I})$ are irrelevant.

We want to show that $\phi \in$ *Opinions*$[\mathcal{A}(\vec{K})]$, so we must ascertain that the same conditions are satisfied on $\vec{K}$. Since, the profiles are deeply ϕ-matching, the members of M must have exactly the same reasons and exactly the same opinions in $\vec{K}$ with respect to all of the propositions that can affect the collective judgment on ϕ. The threshold must be unchanged because it cannot vary with the profile. The only item in the inequality that is sensitive to the profile is $\mathsf{contrast}_\phi$. Perhaps we might have:

$$\mathsf{contrast}_\phi(\vec{J}) \neq \mathsf{contrast}_\phi(\vec{K})$$

But since $\vec{J}$ and $\vec{K}$ are deeply ϕ-matching, condition (#) (one of the criteria of admissibility for values of $\mathsf{contrast}_\phi$ from §3, just after Definition 6) rules this out.

□

Proof of Theorem 3

Proof. Let C be the class of Proper Cohesive Rules. The first conjunct of the definition of maximality is clearly satisfied, as $\mathcal{DS}$ is itself a Proper Cohesive Rule. Consider now an arbitrary agenda $\mathcal{I}$, $\phi \in \mathcal{I}$ and profile $\vec{J}$. We must prove:

$$\phi \in \textit{Opinions}[\mathcal{DS}(\vec{J})] \Leftrightarrow \text{there is a rule } \mathcal{A} \in C \text{ s.t., } \phi \in \textit{Opinions}[\mathcal{A}(\vec{J})].$$

The left-to-right direction is obvious. I argue for the right-to-left direction in two steps. Consider the subset $C' \subseteq C$ consisting of Proper Cohesive Rules with $\mathsf{contrast}_\phi = \mathcal{G}^{\sim\phi,\vec{J}}$. Our goal can then be split in:

$$(*) \text{ there is an } \mathcal{A} \in C' \text{ s.t. } \phi \in \textit{Opinions}[\mathcal{A}(\vec{J})] \Rightarrow \phi \in \textit{Opinions}[\mathcal{DS}(\vec{J})]$$

(**) there is an $\mathcal{A} \in C$ s.t., $\phi \in Opinions[\mathcal{A}(\vec{J})] \Rightarrow$ there is an $\mathcal{A} \in C'$ s.t., $\phi \in Opinions[\mathcal{A}(\vec{J})]$.

For part (*), fix a rule $\mathcal{A}$ in C' and a profile $\vec{J}$ s.t. $\phi \in Opinions[\mathcal{A}(\vec{J})]$. Suppose $\phi \notin Opinions[\mathcal{DS}(\vec{J})]$. By definition $\mathcal{A}$ can only differ from $\mathcal{DS}$ in $\mathsf{threshold}_\phi$. So $\mathcal{A}$ must have a lower threshold than $t_{\mathcal{I}}$. But $t_{\mathcal{I}}$ is the lowest threshold that guarantees consistency-preservation, so $\mathcal{A}$ must not be consistency-preserving and must hence not be Proper.
To complete the argument, we must show (**), but this evidently follows: let $\mathcal{A}_1$ be a Proper Cohesive Rule with $\mathsf{contrast}_\phi \supseteq \mathcal{G}^{\sim\phi,\vec{J}}$, and let $\mathcal{A}_2$ have the same threshold but $\mathsf{contrast}_\phi = \mathcal{G}^{\sim\phi,\vec{J}}$, then it is by definition the case that

$$Opinions[\mathcal{A}_1] \subseteq Opinions[\mathcal{A}_2]$$

which is enough to establish (**).

□

Proof of Theorem 4

Proof. ($PB \Rightarrow C\mathcal{M}$) Suppose that $\phi \in PB(\overrightarrow{Opinions}[\vec{J}])$. Suppose that ϕ is a premise. Then there is a majority on ϕ. Given that the profile is premise-centered, members of this majority must support ϕ non-inferentially, and hence cohesively. So all the premises are accepted by $C\mathcal{M}$.

Suppose that ϕ is the conclusion: then ϕ is accepted in virtue of the fact that a set M of judges accepts a set of premises Σ, such that for all $i \in M$, $\Sigma \hookrightarrow_i \phi$, with every member of M accepting Σ. Since the profile is premise-centered,

(i) $\Sigma \subseteq Premises(\mathcal{I})$

(ii) $\Sigma \hookrightarrow_i \phi$ means that Σ minimally entails ϕ (and so $\Sigma \models \phi$).

Since every $i \in M$ accepts all of Σ, there is a majority in favor of all of Σ. Since (by (i)) Σ contains only premises, and on premise-centered profiles premises are not supported inferentially, $\Sigma \subseteq Opinions[C\mathcal{M}(\vec{j})]$. But then since $\Sigma \subseteq Opinions[C\mathcal{M}(\vec{j})]$, and $\Sigma \hookrightarrow_i \phi$ we must have $\phi \in Opinions[C\mathcal{M}(\vec{j})]$. I note here that in general the judgments outputted by Cohesive Rules are not closed under $\models$, but they are closed under $\hookrightarrow_c$. In the particular case of Premise-Centered profiles, both types of closure hold.

($\mathcal{CM} \Rightarrow PB$) Suppose conversely that $\phi \in Opinions[CM(\vec{J})]$ for $\vec{J}$ a premise-centered profile. If ϕ is a premise, then ϕ is accepted non-inferentially by a majority of judges and hence it is accepted by PB. If ϕ is a conclusion, it is accepted on the basis of the existence of a cohesive majority that accepts ϕ on the basis of a set Σ such that:

(i) Σ minimally entails ϕ.

(ii) Σ contains only premises (or negations of premises).

(iii) every member of M accepts all of Σ.

Conditions (i)-(iii) suffice to guarantee that the $\phi \in PB[\vec{Opinions}[\vec{J}]]$

□

References

CARIANI, Fabrizio (ms.) "Aggregation Rules and Epistemic Norms", manuscript, Northwestern University.

CARIANI, Fabrizio, PAULY, Marc and SNYDER, Josh, "Decision Framing in Judgment Aggregation", *Synthese*, 163, pp. 1-24.

CHANDLER, Jake, (forthcoming) "Acceptance and Aggregation", forthcoming in *Erkenntnis*.

CHAPMAN, Bruce, (2002) "Rational Aggregation", Politics, Philosophy and Economics, 1(3), pp. 337-354.

DIETRICH, Franz (2007) "A Generalized Model of Judgment Aggregation", *Social Choice and Welfare*, 28, pp. 529-565.

DIETRICH, Franz (2008) "Judgment Aggregation and the Relevance of Some Issues to Others", manuscript, LSE.

DIETRICH, Franz (forthcoming) "The Possibility of Judgment Aggregation on Agendas with Subjunctive Implications", *Journal of Economic Theory*.

DIETRICH, Franz, LIST, Christian. (2007) "Judgment aggregation by quota rules: majority voting generalized", *Journal of Theoretical Politics*.

DIETRICH, Franz, LIST, Christian (2007b) "Strategy-proof judgment aggregation". *Economics and Philosophy*, 23(3).

DIETRICH, Franz and LIST, Christian (2008) "Judgment Aggregation without Full Rationality", Social Choice and Welfare, 31, pp. 15-39.

DOUVEN, Igor and ROMEIJN, Jan-Willem, "The Discursive Dilemma as a Lottery Paradox", Economics and Philosophy, 23, 301-319.

GÄRDENFORS, Peter (2006) "A Representation Theorem for Voting with Logical Consequences", *Economics and Philosophy*, 22(2), pp. 181-190.

HORTY, John F. (2002) "Skepticism and Floating Conclusions" *Artificial Intelligence*, 135, pp. 55-72.

HORTY, John F. (2007a) "Reasons as Defaults", *Philosophers' Imprint*, 7(3), pp. 1-28.

HORTY, John F. (2007b) "Defaults with priorities", *Journal of Philosophical Logic*, 36, pp. 367-413.

KONIECZNY Sébastien, LANG Jerome, and MARQUIS, Pierre, "DA^2 Merging Operators", in *Artificial Intelligence*, 157(1-2), 49-79. Artificial Intelligence, 157(1-2):49-79.

LIST, Christian (2005) "Group Knowledge and Group Rationality: a Judgment Aggregation Perspective", Episteme 2(1), pp. 25-38.

LIST, Christian (2007) "When to Defer to Super-Majority Testimony — and When Not To", manuscript, London School of Economics.

LIST, Christian (2008) "Group Deliberation and the Revision of Judgments: an Impossibility Result", in *Journal of Political Philosophy* (forthcoming)

LIST, Christian and PETTIT, Philip (2002)"Aggregating Sets of Judgments: An Impossibility Result", *Economics and Philosophy*, 18, pp. 89-110.

LIST, Christian and PUPPE, Clemens (2008) "Judgment aggregation: a survey", in P. Anand, C. Puppe and P. Pattaniak (eds.), *Oxford Handbook of Rational and Social Choice*, Oxford, OUP.

MILLER, Michael K. (2008), "Judgment Aggregation and Subjective Decision-Making", *Economics and Philosophy*, 24, pp. 205-231.

MILLER, Michael K. and OSHERSON, Daniel (2009), "Methods for Distance-Based Judgment Aggregation", *Social Choice and Welfare*, 32, pp. 575-601.

MONGIN, Philippe (2005) "Factoring Out the Impossibility of Judgment Aggregation", *Journal of Economic Theory*, 141(1), pp. 100 - 113.

NEHRING, Klaus (2005) "The (Im)possibility of a Paretian Rational", manuscript, UC Davis.

Nehring, Klaus and Puppe, Clemens (2010) "Justifiable Group Choice", *Journal of Economic Theory*, 145, pp. 583-602.

Pauly, Marc (2006) "Logical Constraints on Judgment Aggregation", *Journal of Philosophical Logic*, 65, pp. 132-138.

Pettit, Philip (2001), "Deliberative Democracy and the Discursive Dilemma", in *Philosophical Issues*, 11, 2001.

Pettit, Philip (2006) "When to Defer to Majority Testimony and When Not To", *Analysis*, 65, pp. 132-138.

Pigozzi, Gabriella (2008) "Belief Merging and the Discursive Dilemma", *Synthese*, 152(2), pp. 285-298.

Sunstein, Cass (1995) "Incompletely Theorized Agreements", *Harvard Law Review*, 108(7), pp. 1733-1772.