Fabrizio Cariani - Judgment Aggregation

Abstract: Judgment aggregation studies how collective opinions arise from the aggregation of individual ones. This article surveys a variety of aggregation rules (possible ways of aggregating individual judgments into collective ones). Aggregation by majority opinion is known to satisfy some but not all the desiderata for an aggregation rule. More general impossibility results show that not all the natural desiderata can be satisfied by a single aggregation rule. To interpret these results, we focus here on some applications of judgment aggregation models in Epistemology. In particular, we explore their possible role in an account of collective belief and collective reason. The focus on these applications allows us to give precise answers concerning which of the natural desiderata must be abandoned.

1. INTRODUCTION.

Judgment Aggregation studies how collective judgments arise from the aggregation of individual opinions. Its motivating observation is that prima facie plausible rules for aggregating judgments do not (and cannot) have all the features we take to be desirable. Judgment Aggregation, then, aims to classify the various aggregation rules by means of the properties they do satisfy and to select those that are, in some sense, best.

Philosophically, the formal tools of Judgment Aggregation help clarify our understanding of:

- The nature of group belief and of collective decision making.

Groups do not, of course, have single physical minds, but they engage in intentional behavior of the sort that we explain (in the individual case) by ascribing attitudes such as beliefs. Presumably, group beliefs, if there are any, depend on the beliefs of the members of the group. What is the pattern of dependence?

- Group testimony and deference to a group.

Groups need not even act like single intentional subjects; a person who is deferring to independently consulted experts may need to aggregate their opinions into a single set of beliefs. How can this aggregation be accomplished? How should we respond to the testimony of a panel of experts when they disagree?

Both of these applications require us to establish conditions in which it makes sense to ascribe belief in a certain proposition to a group.

It will be useful to organize this essay around a particular philosophical goal: can we model, alongside a concept of collective belief, a notion of collective reason? That is can we make sense of group beliefs that function as reasons for other group beliefs. I start by sketching a very simple, interpretation-neutral, modeling framework. The remainder of this survey is organized around the changes we need to make to the simple framework to satisfy our epistemological aims.

2. A MODELING FRAMEWORK.

The task of the formal framework is to provide a background against which different aggregation rules (that is, ways of aggregating individual judgments) can be defined and evaluated. A very basic framework that can handle this task can be built out of these parameters:[1]

An issue I.[2] A finite set of propositions that is closed under negation (i.e., if $p \in I$, then so is $\sim p$).[3]

A finite set of judges J. For simplicity, assume that J contains an odd number of judges.[4]

Judgment-sets. A judgment-set is a non-empty subset of the issue. To each individual, associate a maximally consistent judgment-set (i.e. a consistent subset S of I such that every proper superset T of S (with $T \subseteq I$) is inconsistent). It is useful to have a term for maximally consistent judgment-sets: I call them *total descriptions*. An individual's opinion can be equivalently represented as a consistent pattern of *Yes/No* responses to the propositions in I (j_i denotes the judgment-set of judge i).[5]

A profile X. A vector (i.e. a sequence) of individual judgments. The i-th position in X corresponds to j_i. Given a profile X, let $X(p)$ be the sequence of Yes and No for p (in terms of the table below, $X(p)$ denotes the column under p).

An assignment of values to these parameters can be represented with a table:

		Issue		
		A	B	A & B
Judges	1	Yes	Yes	Yes
	2	Yes	Yes	Yes
	3	No	No	No

Example 1

An aggregation rule is a function that takes issue I and profile X as inputs and outputs another judgment-set—the collective judgment. In other words, an aggregation rule A takes as input a table, such as the one in Example 1, and outputs another line of the

[1] This is roughly the framework proposed in List and Pettit (2002) and is inspired by the literature on the aggregation of preferences stemming from the groundbreaking work of K. Arrow (1951).

[2] This is often called the *agenda*. However, I find the term 'agenda' to be only adequate if the framework is applied to group decision-making. For this reason, I prefer the term 'issue'.

[3] The relevant sense of 'proposition' here is unstructured. Think of propositions as the equivalence classes induced by the logical equivalence relation onto the sentences of a language. In the literature, it is not infrequent to see 'proposition' used to refer to a syntactic object (i.e., a sentence). This is abusive from a philosophy of language standpoint, but the potential confusions are effectively dodged by the focus on context-insensitive sentences and by the decision to group logically equivalent sentences together.

[4] Even-number sized J complicates the results below, without providing additional theoretical insight.

[5] 'Maximally' consistent means 'complete relative to I'. Although we will not prove this, it is worth remarking this is a harmless restriction at the level of the judgment sets that function as inputs. It is more significant if we require it at the level of the outputs (Dietrich and List, 2008 discuss the significance of this restriction).

table, corresponding to the collective judgment according to A (see Example 2 below). When a proposition p belongs to $A(I, X)$, I say that the aggregation rule A endorses p (on X). To avoid notational clutter, I leave the issue implicit, and just write $A(X)$.

An easy example of aggregation rule is the consensus-or-nothing (CON) rule:

For every p (in I), CON endorses p iff for every judge k, k accepts p.

CON is not very opinionated: in Example 1, CON treats the group as lacking a belief in any one of the propositions at issue. Giving *veto power* to every judge may, in some cases, be desirable; however, we often require an aggregation rule that is more sensitive to the existence of large majorities. For example, it makes sense to explain a certain group choice, by ascribing to the group a belief in even if it is not unanimous in supporting p. With an eye towards more generous aggregation rules, we can explore the majority rule (MAJ):

For every p, MAJ endorses p iff a majority of the judges in J accepts p.

MAJ is opinionated on every proposition for all inputs (recall that I'm excluding even sizes for J). Unfortunately, MAJ is not guaranteed to produce a logically consistent outcome: it can sometimes output a contradictory collective verdict (in spite of the consistency of the individual judgments). This phenomenon is known in the literature as the *discursive dilemma*. Here is a schematic example with issue {A, B, A&B, negations}:

		Issue		
		A	B	A & B
Judges	1	Yes	No	No
	2	Yes	Yes	Yes
	3	No	Yes	No
	MAJ	Yes	Yes	No

Example 2: Discursive dilemma

Although each judge's opinion is consistent, there is a majority for A, a majority for B, but the overlap of the majorities is not itself a majority, and there is a majority against the conjunction.

A dismissive reaction to the dilemma may be to hold that there is nothing wrong in countenancing incoherent group belief states. Nothing wrong, at least, for those philosophers who accept that incoherent belief states are rationally permissible to individual agents (many epistemologists find this to be an attractive line as a response to the Preface Paradox). In response, notice (with Pettit 2006) that individual believers are generally not allowed to rationally entertain patent inconsistencies: however, in the case of MAJ the inconsistent set of beliefs can be as simple as $\{A, B, \sim (A \& B)\}$. We need to find a more substantive solution to the dilemma.

The discursive dilemma is the central motivating point of Judgment Aggregation theory. MAJ seems to be the convergence of a number of desirable properties of aggregation rules: it treats every judge equally, and similarly does not bias some propositions over others (we will specify these properties precisely below). For simple issues such as $\{A, \sim A\}$, MAJ can be justified by reference to the analogue of a classic characterization theorem by K. May (1952), which shows that, provided that the issue is extremely simple, MAJ is the only rule that satisfies a range of such desirable properties. Is there an aggregation rule that shares the desirable properties of MAJ but lacks the undesirable ones? Are there other dimensions, besides not guaranteeing consistency, along which MAJ needs to be improved upon? These are the central questions that Judgment Aggregation aims to address.

Since the discursive dilemma cannot arise for simple issues, I focus in the rest of this survey on issues that exhibit a certain degree of logical complexity. We say that I is conjunctive (resp. disjunctive) just in case, for some distinct propositions p, q (both in I), I also contains the conjunction (resp. disjunction) of p and q.[6]

3. Generalizations: Abstract Impossibility Theorems.

The inadequacy of MAJ motivates the taxonomic aspect of Judgment Aggregation. Classifying aggregation rules by means of the desirable properties they satisfy enables us to evaluate their costs and virtues. Key instruments in this endeavor are impossibility theorems. These results show that, within our framework, desirable properties $P_1,\ldots,P_n$ cannot be jointly satisfied by a single aggregation rule. The relativity to a framework is quite significant: it is in principle possible to evade an impossibility result by moving to a more general framework. Impossibility theorems inform us about the relative costs of various constraints on aggregation (within a given framework). The more general the framework, the more significant the result.

Here is a very basic set of properties that are jointly unsatisfiable within our framework:

Universality. A is defined on every possible profile of individual opinions.

Consistency. For every input profile X, $A(X)$ is consistent.

Completeness. For every input profile X, $A(X)$ is complete relative to I.
[i.e. if p is in I, either p or $\sim p$ is in $A(X)$].

Anonymity. For every permutation h of J, $A(j_1,\ldots,j_n)=A(j_{h(1)},\ldots, j_{h(n)})$
[Informally, the identity of the judges does not matter to the aggregation process.]

[6] Since issues are closed under negation and propositions are identified up to logical equivalence, any conjunctive issue is also disjunctive and vice-versa.

The next condition requires one piece of notation: given a profile X, and a proposition q, let $X(p)$ denote the sequence of judgments on p in X.

> **Systematicity**. For every two propositions p and q and every two profiles X_1, X_2 if $X_1(p)= X_2(q)$, then $p \in A(X_1)$ iff $q \in A(X_2)$.

Systematicity is the conjunction of two conditions:

> 1. **Independence**. For every two profiles X_1 and X_2 and every p in I, if $X_1(p)=X_2(p)$ then $p \in A(X_1)$ iff $p \in A(X_2)$
> [Informally, the verdict on p depends exclusively on the pattern of judgments on p alone]
>
> 2. **Neutrality.** The precise statement of neutrality varies (there are distinct conditions that, if added to Independence, yield Systematicity), but informally it is the idea that an aggregation rule should treat every proposition alike.

An image can help visualize the conditions: suppose you have, for each proposition, a bucket labeled with the name of that proposition. In the *p*-bucket all the judgments on *p* are tossed. Aggregation rules correspond to functions that take the contents of a set of buckets as input and output a verdict on each bucket. Anonymous rules are guaranteed to operate without information concerning the identity of the voters (imagine the voters tossing only their judgments in the *p*-bucket without disclosing their identity). Independent rules are guaranteed to produce an output on *p* just given (i) the contents of the *p*-bucket and (ii) its label (that is, *p*). Neutral rules are guaranteed to provide a verdict irrespective of the label on the bucket.

The Basic Impossibility Theorem (proved in List and Pettit 2002)[7] states that, on conjunctive/disjunctive issues, these conditions cannot be jointly satisfied.[8] [9]

4. INDEPENDENCE OR NOT?

The key conceptual tension in Judgment Aggregation concerns the Independence constraint.[10] On the one hand, there is some motivation to retain Independence. The

---

[7] See also the comparison between this result and Arrow's Theorem in List and Pettit (2004).

[8] What about other issues? Majority is consistent as long as the logical connections among propositions in the issue are limited to the requirement of closure under negation.
An important development in Judgment Aggregation is the study of which combinations of properties are possible given a fixed degree of logical interrelation among the proposition in the issue. So-called *Agenda-characterization* results associate certain logical constraints on the *issue* with *possibility results*.
For an overview of the results, see section 3 of List and Puppe (2008).

[9] The Basic Impossibility Theorem is non-redundant in the sense that for every condition *C* that we invoked, there is an aggregation rule that satisfies all the other conditions but does not satisfy *C*. It is not, however, optimal in the sense that the conditions can be weakened while still retaining the impossibility.

case in favor of Independence relies on a theorem to the effect that assuming Independence is the only way of guaranteeing avoidance of manipulability by voters (an analogue of the social-theoretic condition studied in the Gibbard-Satterthwaite theorem). An aggregation *A* rule is *manipulable* at *X* by *i* on *p* just in case *i* 'disagrees' with *A(X)* on *p*, but can alter the group verdict on *p* merely by submitting an 'insincere' opinion on some other proposition. The theorem (which of course requires a precise definition of 'disagree' and 'insincere') amounts to the observation that, given Universality, Non-Manipulability entails Independence. If Manipulability were undesirable, and if Independence is the only way of avoiding it, we would have motivation in favor of Independence.[11] Of course, this argument turns on whether Manipulability (as characterized) is undesirable, and it is not obvious that it is (what if, for example, the revised opinion on *q* inductively supports *p?*).[12]

On the other hand, the conceptual and technical case against Independence is more compelling. We started off with the idea of aggregating judgments on a set of logically connected propositions, attempting to clarify concepts of collective belief and collective reason. Independence tells us to split this task into many separate instances of judgment aggregation on a single proposition. By doing this, we lose sight of the connection among the propositions. As a consequence, Independent rules are insensitive to the difference between spurious and genuine agreement. S_1 and S_2 may both believe p for compatible reasons; T_1 and T_2 may instead agree on *p* for incompatible (and even mutually undermining) ones. But Independent aggregation rules are insensitive to these differences. Intuitively, however, it should be possible for the difference between the S_i's and the T_i's to bring about a difference in aggregated opinion. Furthermore, it should make a difference to what sorts of considerations can be produced as reasons for the collective belief.[13]

Formal results dramatize this problem.[14] The Basic Impossibility Theorem says that (when the issue is conjunctive or disjunctive), we cannot consistently satisfy all five conditions. Pauly and van Hees (2006) show that we can still get the impossibility theorem if we weaken Anonymity to:

Non-Dictatorship: there is no individual *k* such that for all profiles *X* and issues *I*, $A(X)= j_i$.

Pauly and van Hees also initiated a sequence of results (documented in List and Puppe 2008) to the effect that, given some natural background conditions, and slightly more complicated issues we can get the impossibility without Neutrality. The

---

[10] The parallel condition from Arrow (1952) is also a central piece of the debate in the Social Choice literature.

[11] Dietrich and List (2007). Dietrich and List prove the theorem and observe that it can be put to this use, but do not endorse the argument.

[12] For a more sophisticated argument see Dowding and van Hees (2008).

[13] This disconnect has been pressed in various ways by Chapman (2002), Mongin (2008), Bovens and Rabinowicz (2003) and Pigozzi (2006).

[14] The strengthenings discussed in this section are based on the theorems in Pauly and van Hees (2006).

impossibility theorems that result from weakening these conditions increase the pressure on Independence: the weaker the surrounding conditions, the greater the presumption that Independence is responsible for the impossibility results. However, the most convincing technical argument against Independence is perhaps the simplest. Consider:

Global Majority Constraint: if a total description *K* is accepted by an absolute majority of the judges in *X*, then $A(X)=K$.

K here is an entire judgment set, not a single proposition. In Example 1 from above, for instance, a rule that respects the Global Majority Constraint would endorse the judgment set $\{p, q, p \,\&\, q\}$. Global Majority is an overwhelmingly appealing condition---especially for those who are willing to accept majority rules in the tamed environment of simple issues. However, and this is the key point, Independence and Global Majority can only be satisfied by MAJ, and hence are incompatible with Consistency.[15]

There are regions of logical space which a staunch supporter of Independence could try to occupy, and it is beyond the scope of this overview to close off every possible option. Ultimately, I believe that the negative case (together with the additional evidence that can be found in many of the references below) justifies the near-consensus that Independence is to be rejected.

There is however no consensus on how to relax Independence. Independence conditions look superficially like supervenience claims ("no difference in aggregated outcome on *p* without a difference in the pattern of individual opinions on *p*"). Some non-Independent aggregation rules try to stay as close as possible to this supervenience format. For example, Pettit (2006) holds that, in the application to group deference, we should aggregate group opinion by means of Supermajority testimony. Supermajority rules are just like the majority rule except that the threshold required for acceptance is greater than 50%. If the threshold were constant, supermajority rules would satisfy Independence. However, Pettit recommends a Supermajority rule with a threshold that varies as a function of some global properties in the issue. In particular, Pettit wants a Supermajority threshold *k* such that any set of propositions supported by more than k% of the judges is guaranteed to be consistent.[16] Variable-threshold Supermajority shares some of the conceptual problems that apply to Independent rules (though it evades the technical ones).

5. Collective Reasons.

[15] Proof: Suppose *A* satisfies Independence and Global Majority. Suppose that a profile *X* has a majority on *p*. In X the judges who support *p* may disagree on other propositions in the issue, so they may not accept the same total description. Now consider X' a profile that's identical to X in the distribution of opinion on p but possibly different in that all the judges who support p now submit exactly the same total description. By Global Majority, *A* endorses *p* on *X'* . By Independence, *A* endorses *p* on *X'* iff *A* endorses *p* on *X* (because *X* and *X'* have exactly the same pattern of opinion on *p*). So *A* endorses *p* on *X*. Conversely, if you suppose that there is a minority for *p*, the same reasoning proves that *A* does not endorse *p* on *X*. So *A* must be MAJ.

[16] List (2007) studies how to determine the appropriate threshold as a function of the issue.

In this section, we consider other departures from the supervenience format and consider how they apply to our central question of modeling collective reasons. None of the approaches is completely satisfactory, but they are worth exploring because progress in this area will result by combining the strengths of each approach.

One of the most intriguing programs in Judgment Aggregation (Bovens and Rabinowicz, 2006) ranks aggregation rules according to how well they score on a variety of epistemic measures. Hartmann and Sprenger (ms.) compare different aggregation rules with respect to how well they track the truth. One possibility is to identify the degree to which a rule A tracks the truth relative to a total description S with the probability that A's endorsement is S conditional on S being the correct judgment set (recall that total descriptions are maximally consistent judgment sets), that is:

$$\mathbf{Deg}(A, S)=\Pr(A \text{ endorses } S \mid S)$$

We can use **Deg** to define a partial order on aggregation rules by quantifying out S:

A is better$_{\mathrm{D}}$ than A' iff for every S, $\mathbf{Deg}(A, S) \geq \mathbf{Deg}(A', S)$,

[with strict inequality for some S]

Hartmann and Sprenger take better$_{\mathrm{D}}$ to be a notion of comparative goodness of aggregation rules that is sensitive not just to the need to make the right decision, but also to the need to make the right decision for the "right reasons".[17]

Restricted versions of MAJ prove worthy of special attention as a result of these investigations.[18] Consider a university committee faced with the decision of whether to approve a new minor (M). The university may lay down a set of criteria (i.e. premises) $C_1,\ldots,C_n$. Examples of premises could be that the minor should comply with university-wide regulations, that it does not duplicate existing minors, and so on. The premises are understood to relate to M (the conclusion) by a condition like:

$$M \equiv (C_1 \;\&\; \ldots \;\&\; C_n)$$

On the assumption that the C_i's are setwise (and not just pairwise) logically independent, taking majority on the individual premises (but not on the conclusion M) is a Consistent and Complete aggregation rule. This is premise-based majority (PB).[19] Formulating PB in general requires a small enrichment to our framework: our models must single out two

[17] See also Hartmann, Pigozzi and Sprenger (ms.) which explores which rule is more likely to get the conclusion right.

[18] Hartmann and Sprenger suggest that this program is more constructive than Impossibility-theorem methodology. In my view, this is mostly a difference of emphasis: Impossibility-driven taxonomy also has a positive upshot: it sets up the choice among different aggregation rules as a trade-off among different sets of desirable properties that cannot be jointly instantiated. The true difference is that Hartmann and Sprenger emphasize epistemic properties that come in degrees and ask which rules perform best relative to them.

[19] One can interpret some Judgment Aggregation problems (including the discursive dilemma) as a conflict between premise-driven and conclusion-driven aggregation. Brams, *et al.* (1998) observe that premise-driven and conclusion-driven voting can diverge so sharply that a conclusion endorsed (*qua* conclusion) by none of the voters can end up being recommended by a premise-driven approach.

designated subsets of the issue (the premises and the conclusion) as well as fix the nature of their logical connection.

The sense in which the C_i's function as 'premises' here is that, if all endorsed, they count as reasons for the group to accept M; if even one is rejected by the group, it will count as a reason for the group to reject M. The results I alluded to reveal that, under fairly specific conditions, PB fares better (in the sense of $better_D$) than its competitors.[20]

By contrast, there are two theoretical limits to adopting PB as an aggregation rule. To be inconsistency-proof, PB requires us to isolate, for each issue, a distinguished set of logically independent premises that settle the conclusion. In fact, this set must be more than just 'distinguished': it must be, in some sense, the best such set. No other set of premises can be equally good, since PB is highly unstable with respect to how the premises are identified.[21]

Second, it is not clear that PB latches in full generality onto the concept of collective reason. Consider this example:

> **The Bakery**
> An employee-owned bakery must decide whether to buy a pizza oven (P) or a fridge to freeze their outstanding Tiramisu (F). The pizza oven and the fridge cannot be in the same room. So they also need to decide whether to rent an extra room in the back (R). They all agree that they will rent the room if they decide to buy both the pizza oven and the fridge (P & $F \rightarrow R$)–but they are contemplating renting the room regardless of the outcome of the vote on the appliances.

The intuitive gloss on the case is that while P and F would work as supporting reasons for the 'conclusion' R, there is no sense in saying that their negations also count as reasons for $\sim R$. Rather, we would want to say that $\sim R$ together with P, can count as a reason for $\sim F$. The point is that there is no privileged way of naming premises and conclusions in the Bakery case, independently of an advisor's judgment set. In cases like this, the distinction between premises and conclusions is by no means clear and univocal.

The moral is not that PB is in principle objectionable, but rather that the conditions of its correct applicability are significantly limited. Relatedly, any defense of PB must

---

[20] Other aggregation rules that also presuppose the premise/conclusion distinction. Some examples are:
(i) the conclusion-based procedure CB (take majority on the conclusion).
(ii) situation-based procedures SB (take relative majority on total descriptions); Hartmann and Sprenger discuss this as a comparison point against PB.
(iii) distance based-procedure DB (provide a 'score' of how far away various judgment sets are from each profile of opinion; let the aggregated outcome be the judgment set that maximizes the score); this is defended by Pigozzi (2006).

Like PB, CB satisfies 'restricted' versions of Independence (that is, it satisfies Independence only relative to certain propositions). SB and DB embody a more significant holism, in not satisfying even these localized versions of Independence.

[21] Bovens and Rabinowicz (2006) flag the problem, Cariani, Pauly and Snyder (2008) elaborate on it in a more general framework.

implicitly concede a degree of pluralism about aggregation rules. Different specific aggregation problems may call for different aggregation rules.

Similar considerations apply to a framework studied in Dietrich (ms. a). Dietrich proposes to enrich the modeling framework with a binary relation (which I denote '◄') among the propositions in the issue. This binary relation is intended to be a relevance relation. In different applications, 'relevance' can take up a different meaning. For example 'p◄ q' could mean that p is causally relevant to q; or that, if both are endorsed, p would count as a reason for q; or that a decision on p is prior to a decision on q. In general, logical entailment is neither necessary nor sufficient for 'p◄ q' to hold. Nor is the relevance relation required to be symmetric or transitive. On the positive side, however, the relevance relation is set independently of the judges' opinions: it is fixed at the level of the issue and not allowed to vary from judge to judge.

The chief contributions of this generalized framework are unification (apparently distinct aggregation rules can be obtained from each other by merely changing the interpretation of '◄') and a finer logical space of aggregation procedures in which otherwise inexpressible aggregation rules can be formulated.

So, can an exogenous relevance relation (like ◄) adequately model the notion of collective reason? I think the general answer is negative and the Bakery example reveals why. As I observed in presenting the case, which propositions function as reasons for others partly depends on which other opinions one holds; so, the pattern of dependence cannot be set independently of the judges' individual opinions. In particular, F (getting a new fridge) can count as a reason for R (renting the extra room) or for $\sim P$ (not getting a new pizza oven); it is only as part of a particular belief set that certain beliefs function as inferential reasons for others. This point suggests a further generalization of Dietrich's framework: we could treat the relevance relations as indexed to particular judges. Instead of there being *one* exogenous relevance relation, there would be as many as there are judges. The problem at that point would be to investigate how these relations can contribute to the aggregated opinion.[22]

In related work, Dietrich (2010a) studies a framework in which issues are restricted so as to contain only two types of sentences: the first type comprises atomic sentences and their negations; the second comprises conditionals linking those first-type sentences. These conditionals are not material conditionals and are assumed to satisfy a version of the Stalnaker-Lewis logic. In this environment, conditionals and closure under negation are the only source of logical connections. Dietrich proves that, when the logical interrelation of the issues is so restricted, we can construct consistent *quota rules* (rules in which for every proposition there is a threshold that is sufficient for collective acceptance—possibly a different threshold for different propositions).

This framework might allow an account of collective reasons, if we interpret an advisor's acceptance of a subjunctive conditional '$p > q$' as representing the advisor's

[22] See Cariani (ms.) for more motivation and detail on how to generalize Dietrich´s framework. It must be observed that Dietrich does not take the modeling of collective reasons to be the central application of his framework.

taking the antecedent, if believed, to be a reason for the consequent. This interpretation runs into two problems. First, it is essential to Dietrich's possibility results that the conditionals only relate atomic sentences. So, even if it were correct, it would be limited in scope (for example, the framework would not apply to the Bakery case). Second, it is questionable to link so tightly one's reasons for belief and the subjunctive conditionals one accepts: I could believe, on the basis of my information about his biography and good education, that David must have studied at a college in California. I could also accept the subjunctive conditional 'If David went to college, he went to college in California' (and believe its antecedent). But it does not follow that my reason for believing that he went to college in California is (or even includes) my belief that he went to college.

6. CONCLUSION

The aim of this brief survey was to introduce Judgment Aggregation by emphasizing a particular epistemological application. As a consequence, some important alternative perspectives on the aggregation of opinions have been left in the background, but they should be mentioned.

We might be interested in whether we can apply to group agents a broadly Bayesian model of rational choice. The aggregation of graded beliefs and desires is an area of aggregation theory with a series of distinctive technical problems concerning the compatibility of various desirable conditions (Wagner 1984, 1985, Dietrich 2010b), and specific epistemological applications (Lehrer and Wagner 1981, Fitelson and Jehle 2009). List and Dietrich (ms.a) advance a unified theory of attitude aggregation, within which they formulate general results that apply to both graded and ungraded attitudes. Our epistemological focus on the notion of a collective reason naturally draws attention to qualitative aggregation models. However, qualitative aggregation models need not be (and often are not) framed in these terms. List and Puppe (2008) is a survey of Judgment Aggregation with much deeper technical emphasis.

Having acknowledged these points, the advantage of focusing on a specific application is that it provided us with an immediate spin on the impossibility theorems and on the resources of various modeling frameworks: we must reject Independence and the proposition-wise supervenience picture it imposes. The key open issues involve the extent of the required departure (what modeling resources are needed by non-Independent aggregation rules?) as well as the amount of pluralism we should endorse (can we adopt a single aggregation rule or must we, in modeling different particular situations, adopt different aggregation rules?).[23]

[23] Version of 07/2012. I excised one mistaken remark from the published version.

7. Bibliography

Arrow, K. (1951) *Social Choice and Individual Values*, New York, Wiley.

Bovens, L. and **Rabinowicz**, W. (2003) "Democratic Answers to Complex Questions. An Epistemic Perspective", Synthese, 150, pp. 131-153.

Brams, S.J., **Kilgour**, D. and **Zwicker**, W. (1996) "The Paradox of Multiple Elections", *Social Choice and Welfare* (1998), pp. 211-236.

Cariani, F. (ms.) "Collective Reasons *via* Judgment Aggregation", manuscript, Northwestern University.

Cariani, F., **Pauly**, M. and **Snyder**, J. (2008) "Decision Framing in Judgment Aggregation", Synthese, 163, pp.1-24.

Chapman, B. (2002) "Rational Aggregation." *Politics, Philosophy and Economics* 1(3): 337-354

Dietrich, F. (ms.) "Judgment Aggregation and the Relevance of Some Issues to Others", manuscript, Maastricht University.

Dietrich, F. (2006) "Judgment Aggregation: (Im)Possibility Theorems", Journal of Economic Theory, 126(1), pp. 286-298.

Dietrich, F. (2010a) "The Possibility of Judgment Aggregation on Agendas with Subjunctive Implications", *Journal of Economic Theory*, 145(2), 603-38.

Dietrich, F. (2010b) "Bayesian Group Belief", *Social Choice and Welfare*, 35(4), 595-626.

Dietrich, F. and **List**, C. (2007) "Strategy-Proof Judgment Aggregation", *Economics and Philosophy*, 23, pp.269-300.

Dietrich, F. and **List** C. (2008) "Judgment Aggregation without Full Rationality", *Social Choice and Welfare*, 31(1), pp. 15-39.

Dietrich, F. and **List**, C. (ms. a) "The Aggregation of Propositional Attitudes: Towards a General Theory", *forthcoming* in Oxford Studies in Epistemology.

Dowding, K. and **van Hees**, M. (2008), "In Praise of Manipulation", *British Journal of Political Science*, 38, pp. 1 - 15.

Fitelson, B. and **Jehle**, D. (2009) "What is the 'Equal Weight' View?", Episteme, 6(3), pp. 280-293.

Gibbard, A. (1973) "Manipulation of voting schemes: a general result." *Econometrica* 41, pp. 587-601.

Hartmann, S. and **Sprenger**, J. (ms.) "Judgment Aggregation and the Problem of Truth-Tracking", forhcoming in *Synthese*.

Hartmann, S., **Pigozz**i, G., and **Sprenger**, J. (ms.) "Reliable Methods of Judgment Aggregation", in *Journal of Logic and Computation*, 20, 603-617.

Hawthorne, J. (ms.) "Voting in Search of the Public Good", manuscript, University of Oklahoma.

Lehrer, K. and **Wagner**, C. (1981) *Rational Consensus in Science and Society*, Reidel, Dordrecht-Boston.

List, C. (2006) "The Discursive Dilemma and Public Reason", Ethics, 116(2), pp.362-402.

List, C. (ms.) "When To Defer to Supermajority Testimony and When Not To", manuscript LSE.

List, C. and **Pettit**, P. (2002) "Aggregating Sets of Judgments: An Impossibility Result", *Economics and Philosophy* 18, 89-110.

List, C. and **Pettit**, P. (2004) "Aggregating Sets of Judgments: Two Impossibility Results Compared", *Synthese* 140(1-2), 207-235.

List, C. and **Puppe**, C. (2008) "Judgment Aggregation: a Survey", in Oxford Handbook of Social Choice Theory.

May, K. (1952) "A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions", *Econometrica*, 20(4), pp.680-684.

Mongin, P. (2008) "Factoring out the Impossibility of Logical Aggregation*", Journal of Economic Theory*, 141(1), 100 – 113.

Pauly, M. and **van Hees**, M. (2006) "Logical Constraints on Judgment Aggregation", *Journal of Philosophical Logic*, 35(6), 569-585.

Pettit, P. (2006) "When to Defer to Majority—and When Not To", *Analysis*, 179-187.

Pettit, P. (2007) "Rationality, Reasoning and Group Agency", Dialectica, 61(4), pp.495-519.

Pigozzi, G. (2006) "Belief Merging and the Discursive Dilemma: an Argument-Based Account to Paradoxes of Judgment Aggregation", Synthese, 152(2), pp.285-298.

Wagner, C. (1984) "Aggregating Subjective Probabilities: Some Limitative Theorems", Notre Dame Journal of Formal Logic, 25(3), pp. 233-241.

Wagner, C. (1985) "On the Formal Properties of Weighted Averaging as a Method of Aggregation", Synthese 62, 97-108.