

Introduction to Judgment Aggregation

Christian List and Ben Polak*

February 22, 2010

Abstract

This introduces the symposium on judgment aggregation. The theory of judgment aggregation asks how several individuals' judgments on some logically connected propositions can be aggregated into consistent collective judgments. The aim of this introduction is to show how ideas from the familiar theory of preference aggregation can be extended to this more general case. We first translate a proof of Arrow's impossibility theorem into the new setting, so as to motivate some of the central concepts and conditions leading to analogous impossibilities, as discussed in the symposium. We then consider each of four possible escape-routes explored in the symposium.

JEL classification: D70, D71

Keywords: Judgment aggregation, Arrow's theorem; escape routes

1 Introduction

Ever since the publication of Arrow's classic book, *Social Choice and Individual Values* [1], the theory of aggregation has been a thriving area of research. Arrow's book focused on the aggregation of preferences, understood as the aggregation of multiple individual orderings over a set of mutually exclusive alternatives into a corresponding collective ordering or collective choice. Since the need to aggregate preferences arises in many economic and political contexts, Arrow's work struck a chord with scholars across the social sciences. In fact, the interest in preference aggregation goes back at least to Condorcet and, less formally, to Rousseau in the 18th century, but possibly even to medieval scholars such as Ramon Llull (circa 1232-1315) and Nicolaus Cusanus (1401-1464) (McLean [40]). The broad relevance of the problem of preference aggregation and especially the power and elegance of Arrow's axiomatic approach may also explain why the bulk of almost six decades of social-choice-theoretic research has focused on the aggregation of preferences or on the closely related problems of aggregating individual utilities or individual welfare.

*C. List, Departments of Government and Philosophy, London School of Economics; B. Polak, Department of Economics and School of Management, Yale University. The present symposium is based on a Cowles Foundation workshop on judgment aggregation, held at Yale, September 2006. We are very grateful to Itzhak Gilboa for co-organizing the workshop and helping us launch this project. We also thank the symposium contributors, the referees and other workshop participants for their input, advice and patience. This introduction further benefitted from discussions with Don Brown, Franz Dietrich, Klaus Nehring and Clemens Puppe. Finally, we thank Karl Shell and Susan Schulze for their editorial support and encouragement. List's work was supported by a Philip Leverhulme Prize of the Leverhulme Trust and a Laurance S. Rockefeller Visiting Fellowship at Princeton University.

But there are many contexts in which we want to aggregate individuals' opinions or judgments apart from mere preferences. For example, consider a court of several judges who must decide a case. We may be interested not only in the court's final judgment in the case (e.g., whether the defendant is liable) but also in its findings on the facts of the case (did the defendant do some harm) and on the question of the relevant legal obligation (was the defendant obliged not to do this harm). Thus we may want to aggregate the judges' individual opinions on each of these issues. Elsewhere, consider a search committee charged to select among job candidates. Here, we may be interested not only in the committee's final selection of which candidates to hire, but also in their findings on whether each candidate meets certain criteria for the job, and maybe even in their judgment on which criteria must be met by any successful candidate. Again, we may want to aggregate the committee members' opinions on each of these issues. Or thirdly, consider an expert panel that is asked to give advice on a set of complex scientific questions. Here, we may be interested not only in the panel's view on whether there will be a critical global temperature increase, but also in their judgments on whether various kinds of emissions are above certain thresholds, and what the significance of those thresholds is. Other types of opinions that we might want to aggregate but which are distinct from preferences include categorizations, say, of animals into species or collections of symptoms into a defined disease.

These aggregation problems differ from preference aggregation not just in their interpretation, but also in the constraints governing them. Aggregation problems are generally governed by three types of constraints: constraints on the individual inputs (such as universal domain), constraints on the collective outputs (such as transitivity and completeness), and constraints on the relationship between the inputs and outputs (such as Arrow's independence or Pareto conditions). In Arrovian preference aggregation problems, the inputs and outputs are constrained to be complete and transitive preference orderings. Once we move from preference aggregation to more general aggregation problems, input and output constraints can take the form of requiring sets of accepted positions to be logically consistent (relative to an appropriate criterion of consistency). In the legal setting, for example, we typically want to exclude the combined collective finding that 'the defendant did it', 'he was obliged not to do it', but 'he is not liable' (following the background constraint that harm and obligation are necessary and sufficient for liability). In an expert panel, we want the overall collective viewpoint to be logically coherent. Elsewhere, input and output constraints may come from feasibility conditions. In our hiring-committee setting, there may be a maximum number of candidates we can hire. Alternatively, more complicated constraints might be imposed on the search committee by a higher authority: for example, they may have to hire at least two candidates of which one, at least, must be a woman. In problems of placing items into categories, we may want to restrict the inputs and outputs to partitions or perhaps nested collections of the items.

This paper – and the symposium it introduces – is devoted to a recent effort to develop a general theory that can address these and other problems of aggregation: the theory of judgment aggregation. The theory's central question can be summarized as follows. Suppose a group of individuals is faced with a set of binary questions, each admitting a 'yes/no' or 'true/false' answer. Suppose, further, the questions are interconnected, in that the answers given to some of them constrain the answers that can consistently be given to others. How, then, can the group arrive at a consistent collective set of answers to these questions, based on its members' individual answers? This problem, formulated

in the present form by List and Pettit [36, 37] with precursors in Guilbaud [27], Wilson [61] and Rubinstein and Fishburn [56], is surprisingly general and, as will become clear, subsumes the aggregation problems mentioned above, including the problem of preference aggregation.

The focus of this introductory paper will be to show how ideas from preference aggregation translate into new ideas in the more general setting of judgment aggregation. In part, we hope this offers a bridge to the new literature for those familiar with the old. We also hope to illustrate what is and what is not immediately generalizable from classical social choice theory to judgment aggregation, and what those generalizations look like. For example, the recent literature on judgment aggregation started with a generalization of Condorcet’s paradox (List and Pettit [36, 37]): Just as majority voting can yield cyclical collective preferences, so majority voting on any set of binary issues with certain logical interconnections can lead to inconsistent collective judgments.¹

To illustrate, return to the expert-panel example above. Let p be the proposition that CO₂ emissions are above a given threshold; let $p \rightarrow q$ be the proposition that if CO₂ emissions are above that threshold, then there will be a critical temperature increase; and let q be the proposition that there will be a critical temperature increase. Consider the two matrices in Table 1, one representing the opinions of an illustrative panel of experts and the other representing the preferences in Condorcet’s classic paradox. In each case, all three individuals hold perfectly consistent opinions, in the expert-panel case understood in terms of propositional logic and in Condorcet’s case understood in terms of the standard transitivity and completeness conditions on preferences. In both cases, majority voting fails to preserve consistency from the individual to the collective level: The majority opinions are inconsistent by the relevant criteria. A similar pattern can occur even with the intuitively simplest of all logical connections, such as those between two propositions (p, q) and their conjunction ($p \wedge q$).

	p	$p \rightarrow q$	$\neg q$		$x \succ y$	$y \succ z$	$z \succ x$
Individual 1 $\{p, p \rightarrow q, q\}$	✓	✓	×	Individual 1 $x \succ y \succ z$	✓	✓	×
Individual 2 $\{\neg p, p \rightarrow q, \neg q\}$	×	✓	✓	Individual 2 $y \succ z \succ x$	×	✓	✓
Individual 3 $\{p, \neg(p \rightarrow q), \neg q\}$	✓	×	✓	Individual 3 $z \succ x \succ y$	✓	×	✓
Majority	✓	✓	✓	Majority	✓	✓	✓

Table 1: The problem of majority inconsistency

Given this relationship between Condorcet’s paradox and the more general problem of majority inconsistency, it is natural to ask whether Arrow’s impossibility theorem can be generalized as well. Just as Arrow’s theorem shows that Condorcet’s paradox is the tip of the iceberg of a deeper impossibility result on preference aggregation, so one may ask whether the problem of majority inconsistency is an instance of a more general difficulty

¹This observation – often called the *discursive dilemma* (see also Pettit [53], Brennan [3]) – was inspired by a related but distinct observation in the study of decision making in collegial courts, the so-called *doctrinal paradox* (Kornhauser and Sager [30, 31], Kornhauser [29]). The doctrinal paradox consists in the fact that in decisions on a conclusion, such as whether a defendant is liable, based on several logically related premises, such as whether he did some harm and whether he was obliged not to do it, majority voting on the premises may support a different decision from majority voting on the conclusion.

in judgment aggregation. The literature now contains a sequence of results in response to this question, some of which are included in the present symposium and will be reviewed in detail below. The gist of these results is that, for a large class of aggregation problems, there exist no aggregation rules satisfying some Arrow-inspired conditions such as a universal domain condition, a collective rationality condition, an informational condition akin to Arrow’s independence of irrelevant alternatives, and some minimal democratic conditions such as unanimity preservation, non-dictatorship or anonymity.

These results, in turn, can be traced back to three different intellectual origins: some work on abstract algebraic aggregation,² some logic-based work on judgment aggregation,³ and some related work on strategy-proof social choice.⁴ However, as will become clear, the literature has not been confined to proving impossibility results, but has also considered how to avoid them. Arguably, the key contribution of the impossibility results is precisely to give us some insights into which requirements on aggregation can be upheld, and which must be relaxed, in order to arrive at compelling solutions to various aggregation problems. (For literature reviews, see List and Puppe [38] and List [35].)

The present paper is structured as follows. In Section 2, we introduce the logic-based model of judgment aggregation, following List and Pettit [36] and Dietrich [6]. In Section 3, we discuss how Arrow’s impossibility theorem generalizes from preference aggregation to general judgment-aggregation settings, thereby introducing the first two symposium papers (Nehring and Puppe [49], Dokow and Holzman [23]) and some related results (e.g., Dietrich and List [9]). In Section 4, we turn to several positive results on how to avoid the Arrovian impossibility in judgment aggregation and introduce the following five symposium papers (Dietrich and List [16], Dokow and Holzman [24], Dietrich and Mongin [20], Nehring and Puppe [50], Dietrich [8]). In Section 5, we comment on the relationship between the logic-based model of judgment aggregation and the alternative, abstract algebraic model. In Section 6, we offer some concluding remarks on future avenues of research.

2 The logic-based model

Let $N = \{1, 2, \dots, n\}$ be a finite set of individuals ($n \geq 2$) faced with a judgment aggregation problem. The set of *propositions* on which judgments are to be made is called the *agenda* and defined as a set X of sentences from a suitable logic, where X is non-empty, closed under negation (i.e., if $p \in X$, then $\neg p \in X$, where \neg stands for ‘not’) and, for our purposes, finite. Propositions are denoted p, q, r , and so on. We assume that double negations cancel each other out. The set X might contain not just atomic propositions but also compound propositions such as $p \wedge q$ (p and q) or $p \rightarrow q$ (‘if p then q ’). In the

²Wilson [61], followed by Rubinstein and Fishburn [56] and Dokow and Holzman [23] in this symposium, developed an abstract algebraic framework for studying the aggregation of attributes distinct from preferences and proved a theorem that generalizes Arrow’s theorem in this framework.

³List and Pettit [36, 37] introduced a logic-based model of judgment aggregation and proved an impossibility theorem that also has a corollary for preference aggregation. This work was strengthened and extended by Pauly and van Hees [52], Dietrich [5, 6] and others, and generalized to subsume Arrow’s theorem (Dietrich and List [9]).

⁴Nehring and Puppe [47] proved some results on strategy-proof social choice, which have several corollaries for judgment aggregation, including Arrovian-style results, as presented in this symposium (Nehring and Puppe [49]).

introductory expert-panel example, the agenda is

$$X = \{p, \neg p, p \rightarrow q, \neg(p \rightarrow q), q, \neg q\},$$

and in a preference aggregation problem over three alternatives x , y and z , it is

$$X = \{x \succ y, y \succ x, x \succ z, z \succ x, y \succ z, z \succ y\},$$

with the stipulation that $x \succ y$ is the negation of $y \succ x$ and so on.

Each individual i 's *judgment set* is the set of propositions $J_i \subseteq X$ that he or she accepts (e.g., endorses or believes to be true). A judgment set is called *consistent* if it is a consistent set in the relevant sense of the logic (more on this below). It is *complete* if it contains a member of each proposition-negation pair $p, \neg p \in X$. An n -tuple of judgment sets across the individuals in N , (J_1, \dots, J_n) , is called a *profile*.

An *aggregation rule* is a function F that assigns to each profile of individual judgment sets (J_1, \dots, J_n) in some domain a collective judgment set $J = F(J_1, \dots, J_n)$, interpreted as the set of propositions accepted by the group N as a whole. Usually, we require the domain of F to be the set of all profiles of consistent and complete judgment sets on X (*universal domain*), and we require the collective judgment set to be consistent and complete (sometimes called *collective rationality*). These requirements, however, are not built into the model itself and are sometimes relaxed.

An example of an aggregation rule is *majority voting*, where, for every profile (J_1, \dots, J_n) in the domain,

$$F(J_1, \dots, J_n) = \{p \in X : |\{i \in N : p \in J_i\}| > n/2\}.$$

Another example is a *dictatorship*, where there exists some fixed individual $i \in N$ such that, for every profile (J_1, \dots, J_n) in the domain,

$$F(J_1, \dots, J_n) = J_i.$$

Other examples will be considered later.

While all these definitions should be fairly straightforward, we have glossed over a central element of the model: the logic, which captures the notion of consistency. Generally, a *logic* is a non-empty set \mathbf{L} of sentences, called *propositions*, that is endowed with a *negation operator* \neg ('not') and a notion of *consistency*, according to which some subsets of \mathbf{L} are *consistent* and the others *inconsistent*, subject to some standard conditions.⁵ In standard propositional logic, for example, a set of propositions is consistent if all its elements can be simultaneously true. Thus the set $\{p, q, p \rightarrow q\}$ is consistent (where p and q are atomic propositions) whereas the sets $\{p, p \rightarrow q, \neg q\}$ and $\{p, \neg p\}$ are not. In the logic we have used to express preference aggregation problems, a set of propositions is consistent if it does not breach the rationality conditions on preferences (such as asymmetry, transitivity and connectedness). Thus the set $\{x \succ y, y \succ z\}$ is consistent while the set $\{x \succ y, y \succ z, z \succ x\}$ is not. We can embed feasibility constraints on individual and collective judgments into the notion of consistency in the logic. For example, in the case of a hiring committee, the propositions p and q might be 'hire Paul' and 'hire Quincy',

⁵First, every proposition-negation pair $\{p, \neg p\} \subseteq \mathbf{L}$ is inconsistent. Second, subsets of consistent sets $S \subseteq \mathbf{L}$ are consistent. Third, the empty set \emptyset is consistent, and every consistent set $S \subseteq \mathbf{L}$ has a consistent superset $T \subseteq \mathbf{L}$ containing a member of each proposition-negation pair $\{p, \neg p\} \subseteq \mathbf{L}$. For details, see Dietrich [6].

respectively. If there is only one job opening, then the set $\{p\}$ is consistent while the set $\{p, q\}$ is inconsistent.

A proposition $p \in \mathbf{L}$ for which neither $\{p\}$ nor $\{\neg p\}$ is inconsistent is called *contingent*. For convenience, we include only contingent propositions in the agenda. Further, we say that a set $S \subseteq \mathbf{L}$ *logically entails* a proposition $p \in \mathbf{L}$, written $S \vdash p$, if $S \cup \{\neg p\}$ is inconsistent. For example, in standard propositional logic, the set $\{p, p \rightarrow q\}$ logically entails the proposition q . In the logic used to express preference aggregation problems, the set $\{x \succ y, y \succ z\}$ entails the proposition $x \succ z$. For our hiring committee with only one opening, the proposition p logically entails $\neg q$. By substituting different logics into the model, we can capture a great variety of different aggregation problems.

3 Generalizing Arrow's theorem

To see how Arrow's theorem generalizes from preference aggregation to aggregation problems involving other logical relationships or feasibility constraints, we will state and prove Arrow's theorem and try to translate the proof to more general settings. Although the result is not the most elegant proof of a generalized impossibility theorem (the existing, direct proofs fare better on that count), the aim is to learn something from the translation. First, we will see which notions in Arrow's result and proof generalize easily and what these generalizations look like. Second, where the translation leaves holes, we will see what extra conditions are needed to fill those holes. This will motivate some of the main conditions seen in the symposium and in the judgment-aggregation literature beyond.

For our purposes, we consider Arrow's theorem in the case in which all individual and collective preferences are strict. The preference aggregation problem can then be summarized as follows. Each individual $i \in N$ holds a preference order \succ_i (an asymmetric, transitive and connected binary relation) over a set $K = \{x, y, z, \dots\}$ of alternatives. We are looking for a *preference aggregation rule*, \mathcal{F} , that assigns to each profile $(\succ_1, \dots, \succ_n)$ of individual preference orders a collective preference order \succ . We require \mathcal{F} to satisfy *universal domain* (every profile of strict orders is admissible as input), *collective rationality* (the output is a strict order), *independence of irrelevant alternatives* (the collective preference over any pair $x, y \in K$ depends only on individual preferences over that pair), and the *weak Pareto principle* (if everyone prefers x to y , then x is collectively preferred to y). Arrow's theorem states that, when $|K| \geq 3$, only dictatorships satisfy these conditions, i.e., there is some fixed individual $i \in N$ such that, for every profile $(\succ_1, \dots, \succ_n)$, $\mathcal{F}(\succ_1, \dots, \succ_n) = \succ_i$.

The intended generalization can be motivated by reinterpreting preferences as judgments on special kinds of propositions. For instance, the preference order $x \succ y \succ z$ encodes the judgments that x is preferable to y , y is preferable to z , and x is preferable to z . More generally, any preference order encodes a set of judgments on all binary ranking propositions of the form ' v is preferable to w ', where v and w are distinct alternatives in K . Thus the agenda is the set $X = \{v \succ w : v, w \in K \text{ with } v \neq w\}$. Judgment sets here are consistent and complete when these constraints are understood in terms of the standard constraints on preferences. This reinterpretation of preferences as judgments on binary ranking propositions raises the prospect of proving a result similar to Arrow's theorem for more general agendas of propositions.

What do Arrow's conditions on the aggregation rule look like in this general case? Uni-

versal domain becomes the condition that every profile of consistent and complete judgment sets is admissible as input to the aggregation rule. Collective rationality becomes the condition that the output is a consistent and complete judgment set. The weak Pareto principle becomes the *propositionwise unanimity principle*, requiring that if everyone accepts a proposition p within the agenda, then p is collectively accepted. And independence of irrelevant alternatives becomes the requirement – here called *independence* – that the collective judgment on any proposition p within the agenda depend only on individual judgments on p . We want to know for which agendas of propositions dictatorships are the only aggregation rules satisfying these conditions. This question is non-trivial, since we can easily construct agendas for which non-dictatorial aggregation rules such as majority voting work perfectly well. A simple example is $X = \{p, \neg p\}$.

Let us revisit the key steps in the proof of Arrow’s theorem and generalize them in turn. Throughout the proof, we assume that the aggregation rule satisfies Arrow’s four conditions, and our definitions and lemmas take these conditions as given. Many classic proofs of the theorem consist of two lemmas: the ‘contagion’ or ‘field-expansion’ lemma and the ‘group-contraction’ lemma. Both make use of the fact that, under independence of irrelevant alternatives, a preference aggregation rule can be represented in terms of its ‘winning coalitions’, as defined in a moment.

3.1 The contagion or field-expansion lemma

The literature contains several versions of the first lemma, which refer to weaker or stronger notions of ‘decisiveness’. We write $K \times_{\neq} K$ to denote the set of all ordered pairs of distinct alternatives from K . Recall these standard definitions from social choice theory.

Definition A group $S \subseteq N$ is called *almost decisive over the ordered pair* $(x, y) \in K \times_{\neq} K$ if $[x \succ_i y \text{ for all } i \in S, \text{ and } y \succ_j x \text{ for all } j \in N \setminus S]$ implies $x \succ y$. A group that is almost decisive over all pairs in $K \times_{\neq} K$ is called *almost decisive*. For each pair (x, y) , the set of almost decisive groups over (x, y) , denoted $\mathcal{W}_{(x,y)}$, is called the set of *winning coalitions over* (x, y) .

Definition A group $S \subseteq N$ is called *decisive over the ordered pair* $(x, y) \in K \times_{\neq} K$ if $[x \succ_i y \text{ for all } i \in S]$ implies $x \succ y$. A group that is decisive over all pairs in $K \times_{\neq} K$ is called *decisive*. If a single individual i is decisive, then i is a *dictator*.

The difference between these two notions is that an almost decisive group wins if everyone outside the group opposes it, whereas a decisive group wins regardless of others’ support or opposition. Clearly, decisiveness implies almost decisiveness, and the two notions are equivalent if we accept the monotonicity property that more ‘votes’ in favor of a position cannot hurt collective acceptance of that position. Formally:

Definition An aggregation rule is *monotonic* if, for any $(x, y) \in K \times_{\neq} K$, the set of winning coalitions is superset-closed: $S \subseteq T \subseteq N$ and $S \in \mathcal{W}_{(x,y)}$ imply $T \in \mathcal{W}_{(x,y)}$.

While monotonicity is natural in almost all social-choice and judgment-aggregation settings (when independence is satisfied), it is not always assumed and not required for Arrow’s theorem.

We can now state three standard versions of the contagion lemma.⁶

Lemma 1 (Contagion) *If the group $S \subseteq N$ is*

- (a) *almost decisive over some $(x, y) \in K \times_{\neq} K$, then S is almost decisive.*
- (b) *decisive over some $(x, y) \in K \times_{\neq} K$, then S is decisive.*
- (c) *almost decisive over some $(x, y) \in K \times_{\neq} K$, then S is decisive.*

Notice that the third version is stronger in that it converts almost decisiveness into decisiveness. The proofs of these three claims are very similar, but it turns out that, when we translate to general judgment-aggregation problems, small differences matter.

Proof sketch. For now, let us focus on contagion from the pair (x, y) to the pair (x, z) for all $z \in K$. Consider preference profiles that are consistent with Table 2 (for the moment, ignoring the first row).

	p	Y	$\neg q$
	$x \succ_i y$	$y \succ_i z$	$z \succ_i x$
$i \in S$	✓	✓	×
$i \in N \setminus S$	—	✓	—

Table 2: A profile in the proof of the contagion lemma

The ‘—’ in the first and third columns for row ‘ $i \in N \setminus S$ ’ means that these preferences have not yet been specified. That is, everyone prefers y to z , and everyone in S prefers x to y and x to z . By independence, the collective preference between x and z can only depend on individual preferences over this pair.

Towards showing (a): Suppose that S is almost decisive over (x, y) but not almost decisive over (x, z) . Then, since everyone in S prefers x to z , if everyone in $N \setminus S$ prefers z to x , z must also be collectively preferred to x . But consider profiles in which everyone in $N \setminus S$ has preferences $y \succ_i z \succ_i x$ (that is, their row reads $[\times, \checkmark, \checkmark]$). In this case, since S is almost decisive on (x, y) , x is collectively preferred to y ; and since everyone prefers y to z , y is collectively preferred to z , by the weak Pareto principle. Therefore, by transitivity, x is collectively preferred to z , a contradiction.

Towards showing (b): Suppose that S is decisive over (x, y) but not decisive over (x, z) . Then, since everyone in S prefers x to z , there must exist a combination of preferences over x and z for those in $N \setminus S$ such that z is collectively preferred to x . But consider profiles in which all those $i \in N \setminus S$ with preference $z \succ_i x$ (as before) have preferences $y \succ_i z \succ_i x$ (their row reads $[\times, \checkmark, \checkmark]$); and all those $i \in N \setminus S$ with preference $x \succ_i z$ have preferences $x \succ_i y \succ_i z$ (their row reads $[\checkmark, \checkmark, \times]$). In this case, since S is decisive on (x, y) , x is collectively preferred to y ; and since everyone prefers y to z , y is collectively preferred to z , by the weak Pareto principle. Therefore, by transitivity, x is collectively preferred to z , a contradiction.

Towards showing (c): Suppose that S is almost decisive over (x, y) but not decisive over (x, z) . Then, since everyone in S prefers x to z , there must exist a combination

⁶The first of these used to be in the standard textbook proof. For example, it is the version in Mas-Colell et al. [39]. The second emerged in Sen [59] and is now perhaps standard. For example, it is used in Campbell and Kelly’s chapter on Arrovian impossibility theorems in the Handbook of Social Choice and Welfare [4]. The third, stronger version comes from the classic treatment in Sen [58].

of preferences over x and z for those in $N \setminus S$ such that z is collectively preferred to x . But consider profiles in which all those $i \in N \setminus S$ with preference $z \succ_i x$ (as before) have preferences $y \succ_i z \succ_i x$ (their row reads $[\times, \checkmark, \checkmark]$); and all those $i \in N \setminus S$ with preference $x \succ_i z$ have preferences $y \succ_i x \succ_i z$ (their row reads $[\times, \checkmark, \times]$). In this case, since S is almost decisive on (x, y) , x is collectively preferred to y ; and since everyone prefers y to z , y is collectively preferred to z , by the weak Pareto principle. Therefore, by transitivity, x is collectively preferred to z , a contradiction.

Similar steps show contagion from the pair (x, y) to the pair (w, y) for all $w \in X$; and putting these and similar steps together yields the desired conclusions. \square

3.2 Generalizing the contagion lemma

Let us deconstruct this proof, translate each piece to a general judgment-aggregation setting, and see if we arrive at the building blocks of an analogous lemma. Following our assumption in the case of preferences, we assume that the judgment aggregation rule satisfies the generalized versions of Arrow's four conditions; again, our definitions and lemmas take these conditions as given. First let us translate the notions of decisiveness. As in the case of preference aggregation, any judgment aggregation rule satisfying independence can be represented in terms of its winning coalitions.

Definition A group $S \subseteq N$ is called *almost decisive over the proposition* $p \in X$ if $[p \in J_i$ for all $i \in S$, and $p \notin J_j$ for all $j \in N \setminus S]$ implies $p \in F(J_1, \dots, J_n)$. A group that is almost decisive over all propositions in X is called *almost decisive*. For each p , the set of almost decisive groups over p , denoted \mathcal{W}_p , is called the set of *winning coalitions over* p .

Definition A group $S \subseteq N$ is called *decisive over the proposition* $p \in X$ if $[p \in J_i$ for all $i \in S]$ implies $p \in F(J_1, \dots, J_n)$. A group that is decisive over all propositions in X is called *decisive*. If a single individual i is decisive, then i is a *dictator*.

We will highlight three pieces of the argument.

(1) Conditional entailment. In each of (a), (b) and (c), we ‘leveraged’ a collective preference for x over y to a collective preference for x over z using unanimity about (y, z) as the fulcrum. This works because, *conditional* on a collective preference for y over z , a collective preference for x over y *entails* a collective preference for x over z , since the converse, $z \succ x$, would be inconsistent with transitivity. Intransitivity is the main notion of inconsistency in preference-aggregation settings. But we can translate the notion of *conditional entailment* to any judgment-aggregation setting using that setting's own notion of what is inconsistent or infeasible. Indeed, something close to the following notion is found in most papers in the symposium.

Definition Proposition p *conditionally entails* proposition q , written $p \vdash^* q$, if there is a (possibly empty) subset Y of the agenda X such that $\{p\} \cup Y \cup \{\neg q\}$ is a *minimal inconsistent set* and (for non-triviality) $p \neq \neg q$.⁷

⁷Recall that an inconsistent set is *minimal* if every proper subset is consistent. This implies that replacing any individual proposition in the set with its negation produces a consistent set. In particular, in this case, this means that the conditioning set Y is consistent with p alone and with $\neg q$ alone.

In the setting above, the minimal inconsistent set consists of $p := (x \succ y)$, $Y := (y \succ z)$ and $\neg q := (z \succ x)$, or $[\checkmark, \checkmark, \checkmark]$ in Table 2. Thus, p conditionally entails q : Given unanimity on $(y \succ z)$, if $(x \succ y)$ is collectively accepted, then $(x \succ z)$ must also be collectively accepted. More generally, Y need not be singleton but could be another set of propositions (possibly even empty). In the example from our opening section, $p :=$ ‘CO₂ emissions are above a given threshold’, $Y :=$ ‘if CO₂ emissions are above that threshold, then there will be a critical temperature increase’, and $\neg q :=$ ‘there will *not* be a critical temperature increase’. Here, again, p conditionally entails q : Given unanimity about the temperature-increasing effects of emissions above the threshold, if the panel judges that emissions are above that threshold, then it must also judge that there will be a temperature increase.

Thus, in general settings, if p entails q conditional on some (possibly empty) set of propositions Y , we can always leverage a collective judgment in support of p to a collective judgment in support of q using unanimous acceptance of every proposition in Y as the fulcrum (see the top row of Table 2, where a \checkmark in the Y column is read as acceptance of every proposition in Y).

(2) Using consistent individual profiles to obtain a contradiction. Each leverage argument involved constructing particular profiles of individual preferences and using these to obtain a contradiction. In particular, we constructed profiles for those individuals outside the group S over (x, y) and (x, z) . For this step to work, the constructed preferences had to be consistent (in particular, with the assumed unanimous preference for y over z). For part (a), we used the preferences $y \succ_i z \succ_i x$, or $[\times, \checkmark, \checkmark]$, for all in $N \setminus S$. For (b), we used these preferences again and also $x \succ_i y \succ_i z$, or $[\checkmark, \checkmark, \times]$. These were both consistent preferences. In fact, in both cases, the constructed preferences differed from the relevant minimal inconsistent set $[\checkmark, \checkmark, \checkmark]$ (the preference cycle) in just one place.

The analogous argument in a general setting requires constructing profiles of individual judgments concerning p and $\neg q$ where there is unanimity about every proposition in Y , and where $\{p\} \cup Y \cup \{\neg q\}$ is a minimal inconsistent set (again, see the top row of Table 2). For parts (a) and (b), the analogous constructed judgments for $N \setminus S$, $[\times, \checkmark, \checkmark]$ and $[\checkmark, \checkmark, \times]$, will again differ from the minimal inconsistent set, $[\checkmark, \checkmark, \checkmark]$, in just one place and hence, by minimality, they too must be consistent.

In part (c), however, we used the preferences $y \succ_i x \succ_i z$, or $[\times, \checkmark, \times]$, to obtain a contradiction. This differs from the relevant minimal inconsistent set $[\checkmark, \checkmark, \checkmark]$ in *two* places. In the preference case, this was consistent, but in the general case, this step will only translate if the relevant minimal inconsistent set, $\{p\} \cup Y \cup \{\neg q\}$, has the property that the ‘pair-negated’ variant $\{\neg p\} \cup Y \cup \{q\}$ is itself consistent. More generally:

Definition A minimal inconsistent set Y' is *pair-negatable* if it has a subset of size two, say $\{p', q'\} \subseteq Y'$, such that $(Y' \setminus \{p', q'\}) \cup \{\neg p', \neg q'\}$ is consistent.⁸ In our proof, $p' = p$ and $q' = \neg q$.

In the preference example above, $p' := (x \succ y)$, $Y' \setminus \{p', q'\} := (y \succ z)$, $q' := (z \succ x)$, and $(Y' \setminus \{p', q'\}) \cup \{\neg p', \neg q'\}$ can be summarized by the preferences $y \succ_i x \succ_i z$. Thus the minimal inconsistent set defined by the preference cycle is indeed pair-negatable.

Recall that part (c) of the lemma is not needed if the aggregation rule is assumed to satisfy monotonicity (since, then, almost decisiveness and decisiveness are equivalent).

⁸In general, the set $Y' \setminus \{p', q'\}$ could be empty.

But without assuming monotonicity, to obtain decisiveness from almost decisiveness (i.e., to obtain monotonicity), we need to assume the existence of at least one *pair-negatable* minimal inconsistent subset of the agenda. Indeed, all of the impossibility results in this symposium that do not impose monotonicity assume a condition close to this.⁹

Thus, in general settings, if p entails q conditional on some $Y \subseteq X$, and S is almost decisive (respectively decisive) on p , then S is almost decisive (respectively decisive) on q . If, in addition, $\{\neg p\} \cup Y \cup \{q\}$ is consistent, then S 's almost decisiveness on p implies S 's decisiveness on q .

(3) Connectedness and Total Blockedness. The last piece of the proof of the lemma involved repeating the leverage step using other potential inconsistencies with transitivity. For example, the potential cycle $w \succ x$, $x \succ y$ and $y \succ w$ is used to leverage a collective preference for x over y to one for w over y using unanimity about (w, x) as the fulcrum. Eventually this process reaches every ordered pair $(w, z) \in K \times_{\neq} K$. If we translate the argument to a general judgment-aggregation setting, we find that, just as a virus spreads from one individual to another along a network of infectious contacts, a group's decisiveness on some proposition p spreads to all other propositions that can be reached from p via a sequence of conditional entailments. That is, a group's decisiveness travels across a directed graph of propositions connected by conditional entailments, each associated with an appropriate minimal inconsistent set. Thus, the last extra condition needed to extend the contagion lemma to a general setting is that this directed graph is (strongly) connected: The pattern of minimal inconsistent subsets of the agenda is such that we can travel from any proposition to any other via a sequence of conditional entailments. Indeed, this condition, introduced by Nehring and Puppe [47, 49], or something close to it is found in all the symposium papers establishing impossibility results.

Definition An agenda X is *totally blocked* (sometimes also called *path-connected*) if, for any pair of propositions $p, q \in X$, there exists a sequence of propositions $p_1, \dots, p_k \in X$ such that $p = p_1 \vdash^* p_2 \vdash^* \dots \vdash^* p_k = q$.

Putting these translated pieces of the argument back together, using the two new conditions we have discovered, yields a new general contagion lemma.

Lemma 2 (General contagion) *Suppose that the agenda is totally blocked. Then, if the group $S \subseteq N$ is:*

- (a) *almost decisive over some proposition p in the agenda, then S is almost decisive;*
- (b) *decisive over some proposition p in the agenda, then S is decisive.*
- (c) *almost decisive over some proposition p in the agenda and, in addition, the agenda has at least one pair-negatable minimal inconsistent subset, then S is decisive.*

Proof sketch. For parts (a) and (b), the proof follows from the discussion above. For part (c), let S be almost decisive over some p . By (a), S is almost decisive over all propositions, including those that form a pair-negatable minimal inconsistent set (which

⁹This is equivalent to the existence of an *even-negatable* minimal inconsistent subset of the agenda (i.e., which has a subset of even size whose propositionwise negation makes the set consistent), introduced in Dietrich [6] and Dietrich and List [9], and also to the algebraic condition of non-affineness, introduced in Dokow and Holzman [23] in this symposium.

must exist by assumption). Therefore, there exists at least one proposition over which S is decisive. By part (b), S is decisive. \square

This lemma has a couple of corollaries of independent interest. First, the lemma as a whole tells us that if the agenda is totally blocked and has at least one pair-negatable minimal inconsistent subset, then the aggregation rule is *monotonic*. Second, part (a) on its own tells us that if an agenda is totally blocked, then the aggregation rule is *neutral* in the sense that it aggregates judgments on each proposition in the same way.

Definition A judgment aggregation rule is *neutral* if, for all propositions in the agenda, the set of winning coalitions is the same: $\mathcal{W}_p = \mathcal{W}_q$ for all $p, q \in X$. The strengthening of independence by imposing neutrality is called *systematicity*.

A similar definition also applies to the case of preferences, with propositions replaced by ordered pairs of distinct alternatives.

3.3 The group-contraction lemma

We can now turn to the second lemma used to prove Arrow's theorem. It too comes in several forms, but we will just need one.¹⁰

Lemma 3 (Group contraction) *Suppose the aggregation rule is neutral. Then there exists an individual who is almost decisive over some ordered pair $(x, y) \in K \times_{\neq} K$.*

Proof sketch. Suppose not. We already know that there exist almost decisive groups; e.g., N is such a group, by the weak Pareto principle. Among all groups that are almost decisive over some pair, let S be among the smallest. Let (y, z) be a pair over which S is almost decisive. Fix some $j \in S$ and consider preference profiles that are consistent with Table 3 (for now, ignoring the first row).

	p	Y	$\neg q$
	$x \succ_i y$	$y \succ_i z$	$z \succ_i x$
$i = j$	✓	✓	×
$i \in S \setminus \{j\}$	×	✓	✓
$i \in N \setminus S$	✓	×	✓

Table 3: A profile in the proof of the group-contraction lemma

The groups $N \setminus S$ and $S \setminus \{j\}$ could be empty. Since $S \setminus \{j\}$ is smaller than S , it cannot be almost decisive over (x, y) , and thus x must be collectively preferred to y . Since S is almost decisive on (y, z) , y must be collectively preferred to z . Thus, by transitivity, x must be collectively preferred to z , and hence j is almost decisive on (x, z) . \square

The combination of the two lemmas yields Arrow's theorem: From part (a) of the contagion lemma, we know that the aggregation rule is neutral. The group-contraction lemma then implies the existence of an almost decisive individual on some pair in $K \times_{\neq} K$. By part (c) of the contagion lemma, this individual is decisive over all pairs and hence a dictator. \square

¹⁰This version is essentially the one in Sen [58].

3.4 Generalizing the group-contraction lemma

Again, let us try to translate this proof to a general judgment-aggregation setting. Once again, the proof relies on conditional entailments (in fact, the same conditional entailment as before) and on constructing consistent profiles of individual preferences to induce a contradiction. A quick inspection of Table 3 confirms that each of the individual preferences used in the construction differ from the relevant minimal inconsistent set, $[\checkmark, \checkmark, \checkmark]$, in just one place. Thus, by the same argument as used above, the analogous profile of individual judgments (see the top row of the table) will be consistent. There is, however, one slightly hidden assumption. Unlike in the proof of the general contagion lemma, the minimal inconsistent set used here needs to have three elements: If Y were empty and $\{p, \neg q\}$ were itself a minimal inconsistent set, then the judgments held by the individuals in $N \setminus S$ would be inconsistent. Therefore, the analogous lemma requires the following additional property on the agenda (which, it is easy to check, is implied by total blockedness).

Definition An agenda X is *non-simple* if it has a minimal inconsistent subset with three or more elements.

Adding this extra condition yields the following straightforward generalization of the group contraction lemma.

Lemma 4 (General group contraction) *Suppose that the agenda is non-simple and the aggregation rule is neutral. Then there exists an individual who is almost decisive over some proposition p .*

The proof follows immediately from the discussion above.

3.5 The theorem

We can now state the analogue of Arrow's theorem in judgment aggregation.

Theorem 1 (a) *If the agenda is totally blocked and has at least one pair-negatable minimal inconsistent subset, the only aggregation rules satisfying universal domain, consistency and completeness, the unanimity principle and independence are the dictatorships.*

(b) *If the agenda does not have both of these properties and $|N| \geq 3$, there exist non-dictatorial aggregation rules satisfying the specified conditions.*

Proof sketch. Part (a) follows immediately from the general contagion and contraction lemmas. Its converse (b) (for $|N| \geq 3$) is established through explicit examples, as shown in Nehring and Puppe [49] (for the violation of total blockedness) and Dokow and Holzman [23] (for the violation of the pair-negatability property). \square

The literature contains several variants of this theorem. Dokow and Holzman [23], in this symposium, proved both parts. Part (a) alone was also proved by Dietrich and List [9]. Both of these works build on Nehring and Puppe's earlier variant [49], also in this symposium, in which monotonicity is additionally imposed on the aggregation rule but the agenda condition of pair-negatability is not required. Recall that, with monotonicity,

almost decisiveness implies decisiveness, and hence part (c) of the general contagion lemma is not needed. Nehring and Puppe [49] and Dietrich and List [9] further considered the case in which the aggregation rule is required to be neutral (i.e., independence is strengthened to systematicity). Here total blockedness can be weakened to non-simplicity and the general contagion lemma is not needed (apart from a variant of part (c) in the case without monotonicity, to get from almost decisiveness to decisiveness, using pair-negatability). Interestingly, with neutrality, the unanimity principle becomes largely redundant.¹¹ Precursors of the result with neutrality, under more restrictive agenda assumptions, include the impossibility results by List and Pettit [36] and Pauly and van Hees [52].¹²

The two lemmas underlying the proof above can be found in various places in this symposium and elsewhere. In essence, part (a) of the general contagion lemma occurs in Dokow and Holzman [23] and Dietrich and List [9], and an adapted version also occurs in Dietrich and Mongin’s symposium paper [20], discussed below. The conclusion of part (b) is similar to that of Nehring and Puppe’s contagion lemma [49], but the argument is different. They use what they call the ‘intersection property’, whereas the proof above can be traced back to classical social choice theory. The conclusion of part (c) and its use of pair-negatability can also be found in Dokow and Holzman [23] and Dietrich and List [9], though they each give a separate argument that the pair-negatability property (‘non-affineness’ in Dokow and Holzman) yields monotonicity of the aggregation rule. However, part (c) in the present form, deriving full decisiveness from almost decisiveness on a single proposition, is not commonly used in the judgment-aggregation literature.

The general group-contraction lemma essentially occurs in Dokow and Holzman [23]. Other papers, including Dietrich and List [9] and Dietrich and Mongin [20], replace it by a two-step argument proving, first, that the set of winning coalitions is intersection-closed and, secondly, that the intersection of all winning coalitions is singleton. This proof strategy resembles those proofs of Arrow’s theorem which show that the set of winning coalitions forms an ‘ultra-filter’.¹³ We return to this point when we discuss relaxations of the completeness requirement on collective judgments.

How demanding are the agenda conditions of Theorem 1? Total blockedness turns out to be quite strong. While it is paradigmatically satisfied by the agenda of pairwise preference rankings over three or more alternatives, some of our other illustrative agendas violate it. In the law-court example, where p and q are the premises for liability and r is the proposition about liability (subject to $r \leftrightarrow (p \wedge q)$), total blockedness is violated (e.g., there is no sequence of conditional entailments from the proposition that he did not do it to the proposition that he is liable). Indeed, unanimity voting on each of p , q and r (with $\neg p$, $\neg q$ and $\neg r$ being the ‘default’ judgments in the absence of unanimity) satisfies all of the Arrovian conditions. In the expert-panel example, where the agenda consists of p , $p \rightarrow q$, q and their negations, total blockedness is again violated, and a non-dictatorial aggregation rule satisfying the specified conditions is given by accepting p if it is unanimously accepted and accepting each of $p \rightarrow q$ and q unless it is unanimously

¹¹In the neutral case, even without the unanimity principle, the other conditions imply either (i) dictatorship – when the agenda has at least one inconsistent subset whose propositionwise negation is consistent (Dietrich [6]) – or (ii) the disjunction of dictatorship and inverse dictatorship (Dietrich and List [9]).

¹²Under stronger agenda assumptions, Pauly and van Hees [52] and Dietrich [5] also proved impossibility results without neutrality in which the unanimity principle is weakened to the requirement that the aggregation rule be non-constant. For other discussions of the relationship between preference and judgment aggregation, see also List and Pettit [37] and Nehring [43].

¹³See, for example, Campbell and Kelly [4] and references therein.

rejected (with ‘default’ judgments $\neg p$, $p \rightarrow q$ and q otherwise).

However, despite the demandingness of total blockedness, we can gain important insights from the role played by conditional entailments in making decisiveness contagious. In particular, whenever we can get from p to q via a sequence of conditional entailments, any winning coalition for p must also be winning for q ; and, further, whenever the agenda contains a strongly connected segment, the aggregation rule must be neutral within that segment. In this way, the structure of conditional entailments within the agenda severely restricts the class of admissible aggregation rules, sometimes forcing them to be locally dictatorial, oligarchic or subject to individual veto power. Many of these results, along with a general family of ‘blockedness’ conditions, are in Nehring and Puppe [49] in this symposium. See also Nehring [44] and Dietrich and List [15].

The pair-negatability property is interpretationally less demanding than total blockedness. It is satisfied in all of the examples above, and the only agendas violating it are those that are essentially isomorphic to a set of propositions in standard propositional logic whose only logical connectives are negation (\neg) and material bi-implication (\leftrightarrow) (Dokow and Holzman [23]). For such agendas, the so-called ‘parity rules’, which accept any proposition if and only if it is accepted by an odd number of individuals (in an odd-sized subset of N), are the only aggregation rules satisfying the Arrovian conditions.

Conditions (plus universal domain, consistency & completeness)	Unanimity principle	Unanimity principle & monotonicity
Independence	totally blocked & pair-negatable	totally blocked
Independence & neutrality	non-simple & pair-negatable	non-simple

Table 4: Agendas for which only dictatorships are possible

Table 4 summarizes the theorem’s main variants we have commented on. The rows and columns indicate the conditions imposed on the aggregation rule, and the corresponding entries show the classes of agendas for which any such rule must be a dictatorship.

4 Escape routes from impossibility

As noted in our opening remarks, the most important contribution of the impossibility results on judgment aggregation, as in preference aggregation, is not so much to demonstrate that there are no compelling aggregation rules, but to give us insights into which requirements on aggregation must be relaxed to find such rules. The generalization of Arrow’s theorem points to four different escape routes from the impossibility of non-dictatorial judgment aggregation.

The first three correspond to the three types of constraints on aggregation mentioned at the beginning. We can restrict the domain of admissible inputs to the aggregation rule; relax the constraints on the outputs; or weaken the informational restrictions imposed by independence (we set aside the less compelling route of dropping ‘democratic’ conditions such as the unanimity principle or monotonicity). The fourth route uses the fact that, unlike Arrow’s theorem, which applies to all non-trivial preference aggregation problems

(i.e., all with more than two alternatives), its generalization does not apply to all agendas of substantive interest. We can thus focus on agendas that are substantively interesting but fall outside the ‘problematic’ classes leading to dictatorship. While the first three routes generalize familiar responses to Arrow’s theorem, the last exploits the additional structure that can be introduced by the move from preference aggregation to judgment-aggregation settings.

4.1 Restricting the domain

The impossibility theorem we have discussed takes the aggregation rule to be defined on the domain of all profiles of consistent and complete judgment sets and thus places no restriction on the admissible disagreement between individuals. However, it is sometimes reasonable to expect individual judgments to fall into a more restricted domain, which captures greater agreement or cohesion among individuals. The symposium paper by Dietrich and List [16], building on List [32], asks whether there are plausible domain restrictions that guarantee that propositionwise majority judgments are consistent. Dietrich and List identify several such restrictions, which can be partially ordered by strength. They are similar in spirit to some classic domain restrictions in preference aggregation (such as single-peakedness) and generalize some of them, thereby shedding light on which abstract features of those classic conditions are responsible for their effects on majority consistency.

We here review some examples, beginning with the first and strongest domain-restriction condition in judgment aggregation, called *unidimensional alignment* (List [32]). It is based on the idea that judgments are constrained by the individuals’ location on some left-right axis, which represents their positions on some cognitive or normative dimension. A profile (J_1, \dots, J_n) is *unidimensionally aligned* if the individuals in N can be ordered from left to right such that, for every proposition p in the agenda, the individuals accepting p (i.e., those with $p \in J_i$) are either all to the left, or all to the right, of those rejecting it (i.e., those with $p \notin J_i$). For example, consider the agenda containing the following propositions and their negations:

- p : ‘A budget deficit is acceptable.’
- q : ‘We should increase defense spending.’
- $p \rightarrow q$: ‘If a budget deficit is acceptable, we should increase defense spending.’

Plausibly, individuals on the left of the political spectrum find a budget deficit acceptable, while those on the right do not; those on the right endorse an increase in defense spending, while those on the left do not; and everyone from the center to the right shares the view that, at least conditional on the acceptability of a budget deficit, defense spending should be increased. The result is a unidimensionally aligned profile, as illustrated in Table 5.

	Ind. 1	Ind. 2	Ind. 3	Ind. 4	Ind. 5
p	✓	✓	×	×	×
$p \rightarrow q$	×	×	✓	✓	✓
q	×	×	×	✓	✓

Table 5: A unidimensionally aligned profile

When this condition is satisfied, any proposition in the agenda is accepted by a majority if and only if it is accepted by the median individual on the given left-right order

(ignoring ties). Thus, in the present example, the majority judgments coincide with those of individual 3. But then the majority judgments inherit the consistency of the median individual's judgments, given individual consistency. Applied to binary preference rankings, unidimensional alignment reduces to the established condition of *order-restriction* (Rothstein [55]), which has *single-crossing* as a special case.

Dietrich and List [16] generalize this condition in several ways. They also show that the left-right order of individuals producing the required acceptance-rejection patterns need not be the same for all propositions in the agenda, but can vary across different sub-agendas. For instance, majority judgments are also consistent if the individual judgments restricted to the propositions in each minimal inconsistent subset of the agenda (and their negations) are unidimensionally aligned. This 'localization' of the domain restrictions resembles a familiar move in preference aggregation, where Condorcet cycles are already ruled out by requiring restrictions such as single-peakedness to hold for every triple of alternatives.

In preference aggregation, majority consistency can also be assured by a weaker local condition known as *value-restriction* (Sen [57]): In every triple of alternatives, the individuals unanimously rank one alternative not top, not middle, or not bottom. Dietrich and List show how to generalize this condition to judgment aggregation and give it an alternative form: Any minimal inconsistent set $Y \subseteq X$ contains a pair of propositions p, q not jointly accepted by anyone. To see why this generalizes Sen's condition, notice that if everyone agrees, for instance, that y is not middle among x, y and z , no-one will jointly accept $x \succ y$ and $y \succ z$. While weaker than the earlier domain-restriction conditions, this is still sufficient for majority consistency: If the majority judgments were inconsistent, at least one minimal inconsistent set $Y \subseteq X$ would have to be majority-accepted. But then the pair of propositions $p, q \in Y$ from the definition of value-restriction would be majority-accepted as well, and at least one individual $i \in N$ would have to accept both, contradicting value-restriction.

Whether value-restriction can be relaxed further depends on whether we wish to formulate domain restrictions as constraints on admissible individual judgment sets or as constraints on admissible profiles. In the former case ('product domains'), value-restriction cannot generally be relaxed further. In the latter case ('non-product domains'), a weaker necessary and sufficient condition for majority consistency can be given. Unlike unidimensional alignment or value-restriction, it specifies not only which patterns within each individual's judgment set are or are not permitted to occur, but also how often those patterns may occur across the profile: For every minimal inconsistent subset of the agenda, at least one proposition must be majority-rejected, the verification of which requires explicit counting.

4.2 Relaxing the constraints on the outputs

Instead of restricting the domain of inputs to the aggregation rule, we can also relax the constraints on its outputs. The impossibility theorem above requires the aggregation rule to generate a complete, not just consistent, collective judgment set for every profile of individual judgment sets. But there are many contexts in which completeness seems too much to ask of collective judgments; we may be able to agree or at least resolve our disagreements on some issues while agreeing to disagree on others.

Two natural kinds of aggregation rules violating completeness are unanimity rules and

supermajority rules (List and Pettit [36]).¹⁴ The former accept any proposition (likewise, its negation) collectively if and only if it is accepted by all individuals; the latter, if and only if it is accepted by a sufficiently large supermajority, such as two thirds, three quarters, or more. In both cases, no judgment is made when neither the proposition nor its negation receives the required support. An important difference between the two kinds of rules lies in the rationality conditions on collective judgments they secure. Notice that when a judgment set is consistent and complete, it is also *deductively closed*: Any proposition in the agenda that is logically entailed by the given judgment set is also contained in it. In the expert-panel example, the set $\{p, p \rightarrow q, q\}$ is deductively closed, while the set $\{p, p \rightarrow q\}$ is not. Accordingly, when we relax completeness, we must choose whether or not to preserve deductive closure. Since the intersection of several deductively closed sets of propositions is still deductively closed, unanimity rules guarantee deductive closure; supermajority rules do not.

Philosophers are divided on whether deductive closure is a reasonable requirement of rationality. On the one hand, the acceptance of any given set of propositions may be taken to commit someone to the acceptance of its implications. On the other hand, it may be computationally difficult to identify these implications, and the evidential support available for any set of propositions need not carry over to its implications. In the case of strict preference rankings, the conjunction of consistency and deductive closure is equivalent to the conjunction of asymmetry and transitivity, and consistency by itself is equivalent to acyclicity.

Dokow and Holzman [24], in this symposium, and Dietrich and List [13] investigate the general case in which completeness is dropped but deductive closure is preserved, strengthening related results by Gärdenfors [26]. The key finding is that the dictatorship result discussed above then turns into an oligarchy result. If (and, for $|N| \geq 3$, only if) the agenda is totally blocked and has at least one pair-negatable minimal inconsistent subset, the only aggregation rules generating consistent and deductively closed collective judgments and satisfying Theorem 1's other conditions are the *oligarchies*: There exists a fixed non-empty subset $M \subseteq N$ such that, for every profile (J_1, \dots, J_n) in the domain, $F(J_1, \dots, J_n) = \bigcap_{i \in M} J_i$. Unanimity rules are the limiting cases in which the set of 'oligarchs' M contains everyone. Dictatorships are the opposite limiting cases in which M is singleton.

To see why this oligarchy result holds, notice that the general contagion lemma above continues to hold even under the present weakening of the output constraints. In particular, the aggregation rule is still neutral and monotonic under the given conditions (i.e., the set of winning coalitions is the same across the agenda and superset-closed). The group-contraction lemma, by contrast, ceases to hold and must be replaced by another argument. As in the ultra-filter proofs of Arrow's theorem, the set of winning coalitions can be shown to be intersection-closed. However, without completeness, the intersection of all winning coalitions need no longer be singleton, but can be a larger set $M \subseteq N$, which consists precisely of the oligarchs. When applied to binary preference rankings, this oligarchy theorem becomes a variant of Gibbard's classic oligarchy result for the case of quasi-transitive collective preferences (see Sen [58]). Dokow and Holzman [24] and Dietrich and List [13] show further that this oligarchy result continues to hold if completeness is weakened to deductive closure not only on the output side but also on the input side of

¹⁴A third kind is given by the 'conclusion-based' rules that deliver judgments only on conclusions, but not on premises.

the aggregation rule.

If we give up deductive closure, non-oligarchic possibilities open up. For example, let each proposition in the agenda be collectively accepted whenever it is accepted by a supermajority of more than $\frac{k-1}{k}$ of the individuals, where k is the size of the largest minimal inconsistent subset of the agenda. The resulting collective judgments are always consistent (Dietrich and List [10], extending List and Pettit [36]). The reason is that k or fewer supermajorities of this size must have a non-empty intersection, and so the supermajority judgments on any set of k or fewer propositions inherit their consistency from the judgments of the individuals in this intersection. Thus no minimal inconsistent set, and by implication no other inconsistent set, can be accepted.

Supermajority rules implement different criteria for the acceptance and the rejection of each proposition; it is harder to accept a proposition than to reject it. Dietrich and List [19] show that such an asymmetry is necessary for the avoidance of dictatorship, even when the only requirement on judgment sets is consistency. Whenever the agenda is non-simple and cannot be partitioned into two logically independent sub-agendas, the only aggregation rules mapping profiles of consistent individual judgments to consistent collective ones satisfying a condition of *acceptance-rejection neutrality* (requiring symmetrical treatment of acceptance and rejection of each proposition) are dictatorships.¹⁵

4.3 Relaxing independence

A third escape-route from impossibility is to relax propositionwise independence. In preference aggregation, independence of irrelevant alternatives is usually defended by appealing to strategy-proofness. A defense of independence by appealing to non-manipulability can also be given in some judgment-aggregation settings (Dietrich and List [11], Nehring and Puppe [47, 49]). But even when we ignore strategic issues, we usually want a court's collective judgment on the factual question of whether the defendant did it to be independent of individuals' judgments on the legal question of whether he was obliged not to do it. However, it seems less natural (on normative or, absent cynicism, strategic grounds) to require the court's collective judgment on the defendant's liability to be independent of the individuals' judgments on whether he did it or whether he was obliged not to. If we relax independence, several interesting classes of aggregation rules become available. The most widely discussed one, the class of *premise-based rules*, is the topic of Dietrich and Mongin's symposium paper [20]. These rules go back, under the name *issue-by-issue voting*, to Kornhauser and Sager's work in a legal context [30, 31, 29], where they were contrasted with *conclusion-based rules*, or *case-by-case voting*, and more formally to List and Pettit [36], Bovens and Rabinowicz [2], Dietrich [5] and others.

The basic idea is to designate a subset of the agenda as a set of premises and to aggregate judgments on them by some propositionwise independent rule, but to allow the collective judgments on all other propositions (the conclusions) to depend on the resulting collective judgments on these premises. In the court setting, the issues of whether he did it and whether he ought not to have done it might be thought of as the premises, and we might decide these, for example, by majority voting. The issue of liability might be thought of as the conclusion, and we might decide this, for example, by the rule that the defendant is liable if and only if we collectively judge that he did it and ought not to have done

¹⁵Acceptance-rejection neutrality is related to *neutrality within issues* or *unbiasedness*, discussed in Nehring and Puppe [49] in this symposium and Dietrich and List [18].

so. (*Conclusion-based rules*, by contrast, apply a propositionwise independent aggregation rule directly on the conclusions.) Earlier contributions have usually assumed logically independent premises – ideally chosen so as to constitute a logical basis for the entire agenda, so that their adjudication settles all other propositions – and have focused mainly on majority voting on these premises. Dietrich and Mongin [20] generalize the notion of a premise-based aggregation rule by defining it in terms of two properties: propositionwise independence restricted to the premises and unanimity preservation also restricted to the premises. This permits, among other things, logically interdependent premises.

Dietrich and Mongin use this framework to generalize the so-called *doctrinal paradox*, the fact that premise-based and conclusion-based rules can deliver different outcomes. Just as Dietrich and Mongin define premise-based rules in terms of two weak properties, they define conclusion-based rules in terms of the weak requirement of unanimity preservation on all non-premises. Putting these two definitions together, they identify the class of agendas for which premise-based and conclusion-based aggregation are in conflict, i.e., for which only dictatorships meet both requirements (under universal domain, consistency and completeness). The identified agenda conditions generalize the agenda conditions of Theorem 1 and reduce to them when the premise set is the entire agenda. The bottom line is that, in many contexts, to solve the doctrinal paradox one must choose between privileging premises and privileging conclusions. Earlier related results were given by Mongin [42] and Nehring [45].

Another class of aggregation rules that give up independence is the class of *sequential priority rules* (List [33]). Any such rule is defined relative to a particular *order of priority* among the propositions in the agenda X . For each profile (J_1, \dots, J_n) , the propositions in X are then considered one-by-one in the specified order and the collective judgment on each proposition p is made as follows. If the majority judgment on p (or the judgment made by some other propositionwise criterion) is consistent with the collective judgments on propositions considered earlier, then that judgment becomes the collective judgment on p ; but if it is inconsistent with those earlier judgments, then the collective judgment on p is determined by the implications of those earlier judgments. In the expert-panel example of Table 1 above, the propositions might be considered in the order $p, p \rightarrow q, q$ (with negations interspersed), which would lead to the acceptance of p and $p \rightarrow q$ by majority voting and the acceptance of q by logical inference.

Just as the collective judgments produced by any premise-based rule depend on which propositions are chosen as the premises, the collective judgments produced by a sequential priority rule depend on the order in which propositions are considered. Sometimes there may be a natural choice of premises, or a natural order of priority, as in the court example, but often there can be reasonable disagreement about which propositions should be considered prior to others. The problem of majority inconsistency thus resurfaces as a problem of *path-dependence* (List [33]).

A third class of aggregation rules that are possible without independence is the class of *distance-based rules* (Pigozzi [54], Miller and Osherson [41]), originally introduced in the area of belief merging in computer science (Konieczny and Pino Pérez [28]). Any such rule is defined relative to a particular *distance metric* between judgment sets. An example is the *Hamming distance*, which defines the distance between any two judgment sets $J, J' \subseteq X$ as the number of propositions in the agenda on which J and J' disagree, i.e., $d(J, J') = |\{p \in X : p \in J \not\leftrightarrow p \in J'\}|$. The aggregation rule then maps each profile (J_1, \dots, J_n) to a consistent and complete collective judgment set J that minimizes

the total distance from the individual judgment sets, $\sum_{i \in N} d(J, J_i)$. Distance-based rules capture the idea of reaching a compromise between different individuals' judgment sets. If the Hamming distance is used, they further have the property of delivering the majority judgments whenever these are consistent. Applied to binary preference rankings, this *Hamming rule* reduces to what is known as the *Kemeny rule*.

To analyze the possibilities opened up by relaxing independence more systematically, Dietrich [7] introduces a condition of *independence of irrelevant information* defined in terms of a *relevance relation* between propositions. An aggregation rule satisfies this condition if and only if the collective judgment on each proposition depends only on individual judgments on propositions deemed 'relevant' to it. The condition reduces to the standard independence condition when each proposition is deemed relevant only to itself, but, for example, the kind of dependency structure we find in premise-based rules can also be captured: premises are deemed relevant to conclusions, but nothing is deemed relevant to any premise except that premise itself. Similarly, a sequential priority rule corresponds to a relevance relation that takes the form of a linear order over the propositions; and a distance-based rule corresponds to one where every proposition is relevant to every other. Dietrich's framework allows us to explore how the logical structure of the agenda, the relevance relation and the conditions on aggregation jointly determine the space of possibilities. However, while aggregation rules that give up independence have some attractive properties, their drawback can be shown to be manipulability, both by strategic voting (Dietrich and List [11], Nehring and Puppe [47, 49]) and by strategic agenda setting (List [33], Dietrich [5]).

4.4 Focusing on special agendas

As noted, total blockedness is a demanding condition on the interconnectedness between propositions that is by no means satisfied in all judgment-aggregation settings of interest. Several papers in the literature look at natural agendas that are not totally blocked, and ask whether any of the aggregation rules that become possible here are also plausible. While we have already considered non-independent rules that exploit a special structure of the agenda by prioritizing some propositions over others – e.g., 'premises' over 'conclusions' – our focus is now on aggregation rules that retain propositionwise independence.

Consider again the class of agendas that permit a natural division between a set of issues that are the 'criteria' for a particular decision and a *single* issue that is the 'decision' itself: In the court example, the 'decision' is whether or not the defendant is liable, and the 'criteria' are whether or not he did it and whether or not he was obliged not to do it. Another example could be our academic hiring committee if a decision is made only on one candidate, but the committee must also form collective judgments on several criteria that the candidate might or might not satisfy: say, whether her ability is sufficient, whether her field of expertise fits departmental needs, and so on. In the cleanest such examples, there are no (non-trivial) logical connections between the criteria except those that run via the decision; that is, any combination of judgments on the criteria alone is consistent (provided no proposition-negation pair is accepted). We can then summarize the constraints on the problem by asking which combinations of judgments on the criteria force the decision to go one way, say, 'acceptance'; and which force it to go the other, say, 'rejection'. Following Nehring and Puppe [50] in the symposium, call the former set of combinations of judgments on the criteria the *acceptance region*; and the latter set the *rejection region*.

In some cases, the acceptance and rejection regions exactly partition the set of possible (complete and consistent) judgments on the criteria; that is, if all of the criteria are settled, then a particular decision is entailed. These cases are called *truth-functional*. For example, in the court case, the defendant is liable if and only if he did it and he was obliged not to do it. Nehring and Puppe [48] and Dokow and Holzman [21] analyze truth-functional agendas like these. Nehring and Puppe [48] show that, for such agendas, given the other Arrovian conditions (including independence) plus monotonicity, dictatorship can be avoided if and only if either the acceptance or the rejection region contains essentially only a single combination of judgments on the criteria.¹⁶ For example, in the court case, the acceptance region contains only the combination of criterion-judgments ‘he did it’ and ‘he was obliged not to do it’. Even in cases like these, the only possible aggregation rules are *oligarchic with default*: There is a particular (consistent and complete) default judgment set $J_0 \subset X$ and a non-empty set of oligarchs $M \subseteq N$ such that any collective departure from the default judgment in J_0 requires unanimous consent among the oligarchs on that particular departure. The natural default in the court setting would be that the defendant did not do it, was not obliged not to do it, and is not liable. An admissible aggregation rule would be to pick two specific judges and only to depart from the default on each issue if both agreed to do so on that issue. The only anonymous such oligarchy is a *unanimity rule with default*, where $M = N$ (which is distinct from the unanimity rules we encountered in our discussion of relaxing completeness). A feature of oligarchies with default is that each oligarch has the power to *veto* any move away from the default. Thus, although we get a possibility result with independence here, the possibility is limited.

Many settings, however, are not truth-functional: Some combination of judgments on the criteria still leave room for the decision to go either way. The hiring committee above might be constrained to accept the candidate if they judge her to meet both criteria, and be constrained to reject her if they judge her to fail on both criteria, but they might be able to accept or reject her if they judge her to meet just one criterion. Many such problems have a natural ‘monotonicity in criteria’: With an appropriate relabeling of criterion-propositions and their negations, accepting more criteria cannot push you out of an acceptance region or into a rejection region, and accepting fewer criteria cannot push you into acceptance or out of rejection. In their second symposium paper, Nehring and Puppe [50] show that, in such cases, there are aggregation rules satisfying the Arrovian conditions plus monotonicity that avoid veto power – and thereby avoid oligarchy with default – if and only if either the acceptance region or the rejection region is empty.¹⁷ The hiring-committee example above violates this stipulation but if the committee’s only constraint were to reject the candidate if she failed on ability and field – i.e., the committee were never *forced* to accept the candidate – then veto power could be avoided.

Some of the possible aggregation rules in this case seem natural given the structure of the constraints. In particular, there are admissible aggregation rules that require supermajorities (but not unanimity) on any criterion that might push the collective judgment into the non-empty acceptance or rejection region. In the example, supermajorities would be required to conclude that the candidate has insufficient ability and to conclude that the candidate does not meet the needs of the department. The final decision to hire or

¹⁶The qualification ‘essentially’ is needed because there could be ‘redundant’ criteria that do not affect the decision either way. Formally, the condition is sometimes called ‘conjunctive’: Either acceptance or rejection of the decision-proposition is logically equivalent to a particular (consistent) conjunction of acceptances or rejections on each of the non-redundant criterion-propositions.

¹⁷This result requires the number of voters to be large relative to the number of criteria.

not could then be made by simple majority rule. The general intuition here is that if a particular proposition is critical in that accepting it might severely constrain collective acceptance of other propositions, then we need to hold it to a higher acceptance standard; in this case, a supermajority requirement.

This theme returns in Dietrich’s symposium paper [8]. Dietrich discusses, among other things, what he calls *simple implication agendas*, which contain only atomic propositions (and their negations) and propositions (and their negations) asserting an implication such as $p \rightarrow q$ involving just two atomic propositions. In the expert-panel example of Section 1, the implication $p \rightarrow q$ was that if CO₂ emissions are above a given threshold, then there will be a temperature increase. Some care is required in interpreting the ‘ \rightarrow ’ here. Under the standard material interpretation of propositional logic, $p \rightarrow q$ is equivalent to ‘ $\neg p$ or q ’. This would rule out (i.e., render inconsistent) the natural sounding judgments favored by big oil companies – ‘CO₂ emissions have not exceeded a given threshold, and, even if they had, there still would not be a temperature increase’ – since this combines $\neg p$ and $\neg(p \rightarrow q)$. Dietrich argues that when the implication $p \rightarrow q$ is itself open to judgment independent of the evaluation of the atomic propositions, it is more natural to adopt a ‘subjunctive’ interpretation of the implication. Under such an interpretation, $p \rightarrow q$ rules out combining p and $\neg q$ but is consistent with any other pair of judgments on p and q .

With this interpretation of implication, Dietrich provides a startling possibility result for simple implication agendas using just *quota rules*: Each proposition $p \in X$ is collectively accepted if and only if the number of individuals accepting p is greater than or equal to a proposition-specific quota m_p (where, to ensure completeness and to avoid trivial inconsistencies, $m_p + m_{\neg p} = n + 1$).¹⁸ Dietrich shows that, for any simple implication agenda, a quota rule ensures consistency if and only if $m_q \leq m_p + m_{p \rightarrow q} - n$ for all $p \rightarrow q$ in the agenda. For intuition, suppose that both p and $p \rightarrow q$ are collectively accepted. Then the number of individuals who accept both p and $p \rightarrow q$ is at least equal to the right side of the inequality. Each of these individuals must also accept q , and the inequality says that there are enough of them to ensure the collective acceptance of q , which avoids inconsistency.

An immediate consequence of this result is that if the agenda contains $p \rightarrow q$, then $m_q \leq m_p$. In our expert-panel example, the quota to accept the putative ‘cause’ (CO₂ emissions are above the threshold) must be at least as high as that to accept its putative ‘effect’ (temperatures will rise). More generally, if an atomic proposition occurs early in a chain of reasoning – for example, if it refers to a more fundamental cause – it must have a weakly higher quota. The only way to maintain neutrality across atomic propositions is to set the quota $m_{p \rightarrow q}$ on each implication-proposition equal to n ; i.e., to require unanimity. In particular, if there are ‘cycles’ of implications in the agenda, then every $p \rightarrow q$ in such a cycle must be held to a unanimity requirement.

Taken together, these papers suggest that specific agendas admit specific aggregation rules tailored to the setting at hand. A more general lesson, however, is that critical propositions need to be held to higher acceptance standards. Understanding this general lesson takes us full-circle back to the notion of contagion introduced in Section 3. Recall that, if we can reach q from p by a sequence of conditional entailments, then any winning coalition for p is also winning for q . In the case of quota rules, this means that the quota for q is lower than that for p : perhaps, a simple majority versus a supermajority. More

¹⁸For a discussion of quota rules on general agendas, see Dietrich and List [10].

generally, if there is possible contagion from p to q , then consistent aggregation rules that maintain independence must make it weakly harder to accept p than to accept q .

5 The abstract algebraic model

Some papers in the symposium use the logic-based framework outlined in Section 2 above. Others, however, use the abstract algebraic framework introduced by Wilson [61] and extended by Rubinstein and Fishburn [56], in which judgments take the form of 0 ('false') or 1 ('true') evaluations on a finite non-empty set $K = \{1, 2, \dots, k\}$ of binary issues. An *evaluation vector* is an assignment of 0s and 1s to these issues, formally an element of $\{0, 1\}^k$.

Just as some sets of propositions in the logic-based model are deemed inconsistent, so some evaluation vectors are deemed infeasible. Formally, let $Z \subseteq \{0, 1\}^k$ be the set of evaluation vectors that are deemed feasible. For instance, in the expert-panel example of Section 1, there are three issues, corresponding to the three proposition-negation pairs $\{p, \neg p\}$, $\{p \rightarrow q, \neg(p \rightarrow q)\}$, and $\{q, \neg q\}$. The set of feasible evaluation vectors is either

$$Z = \{(1, 1, 1), (1, 0, 0), (0, 1, 0), (0, 1, 1)\}$$

or

$$Z = \{(1, 1, 1), (1, 0, 0), (0, 1, 0), (0, 1, 1), (0, 0, 0), (1, 0, 1), (0, 0, 0)\},$$

depending on whether the conditional \rightarrow is interpreted as a 'material' conditional (as in standard propositional logic) or as a 'subjunctive' conditional (as in conditional logics, like those discussed in Dietrich's symposium paper [8]).

As the name suggests, the abstract model abstracts from some details of the real-world settings. In particular, the same feasible set can arise from very different contexts. For example, on the one hand, consider the preference aggregation problem with three alternatives x , y and z , and three issues: $\{x \succ y, y \succ x\}$, $\{y \succ z, z \succ y\}$, and $\{x \succ z, z \succ x\}$. On the other hand, consider a faculty committee charged to hire some candidates. Suppose there is an (as yet) unspecified number of slots and only three candidates: two women, A and B, and one man, C. Suppose further that, for some reason, the university administrators decide to put just three propositions (and their negations) on the agenda for the faculty committee to decide: 'hire at least A', 'hire at least B', and 'hire at least two of the candidates including at least one woman'. (Notice that this agenda, perhaps by design, does not allow the direct expression of support for C only.) Both of these very different aggregation problems yield the same set of feasible evaluations:

$$Z = \{(1, 1, 1), (0, 0, 0), (0, 1, 0), (1, 0, 1), (1, 0, 0), (0, 1, 1)\}.$$

In the first problem, $(1, 1, 0)$ and $(0, 0, 1)$ represent cyclic preferences and hence are ruled out. In the second, they represent internally inconsistent choices and hence are ruled out. The details of the setting, including the nature of the inconsistencies, are abstracted away.

The advantage of one framework over the other depends on the setting (e.g., the nature of the constraints) and the question at hand (e.g., the degree to which the interpretation of the propositions and their interconnections matter).

6 Concluding remarks

Arguably, the initial phase of the recent research effort on judgment aggregation, namely the quest to extend some classic results on preference aggregation to general judgment-aggregation problems, is now largely complete. In addition to the results discussed above, particularly Arrow's theorem and various responses to it, some other classic results have been generalized, notably the Gibbard-Satterthwaite theorem (Nehring and Puppe [49, 47] and relatedly Dietrich and List [11]), and Sen's liberal paradox (Dietrich and List [14] and relatedly Nehring [45, 46]).¹⁹

There are at least two challenges for the next phase of work in this area. The move to general judgment-aggregation settings has introduced additional complexities not present in preference aggregation. One important challenge will be to explore those additional complexities further, especially in the case of aggregation rules that give up propositionwise independence. As noted, Dietrich and Mongin's symposium contribution on premise-based approaches [20] and Dietrich's recent work on general relevance relations [7] are steps in this direction, but more work remains to be done on characterizations of compelling non-independent aggregation rules.

Another challenge will be to extend the theory of judgment aggregation from the case of binary judgments to that of general propositional attitudes, which need not be binary and which may be governed by very different rationality or feasibility constraints. Recent contributions along these lines include some works on judgment aggregation in many-valued logics (Pauly and van Hees [52], van Hees [60]), some works on the relationship between judgment aggregation and probability aggregation (Osherson and Vardi [51], Dietrich and List [12], Nehring [46]), and some works on the aggregation of abstract, non-binary evaluations (Rubinstein and Fishburn [56], Dokow and Holzman [22]) or of general propositional attitudes (Dietrich and List [17], Duddy and Piggins [25]). The non-binary case, like the non-independent one, permits far more possibilities than the more constrained binary and independent case, and the difficulty will be to arrive at elegant characterizations of aggregation rules given the additional degrees of freedom. We hope that the present symposium will provide a useful entrance point to this vibrant area of research and stimulate further work.

References

- [1] Arrow, K. (1951/1963) *Social Choice and Individual Values*. New York (Wiley).
- [2] Bovens, L., Rabinowicz, W. (2006) Democratic Answers to Complex Questions – An Epistemic Perspective. *Synthese* 150(1): 131-153.
- [3] Brennan, G. (2001) Collective coherence? *International Review of Law and Economics* 21(2): 197-211.
- [4] Campbell, D. E., Kelly, J. S. (2002) Impossibility Theorems in the Arrovian Framework. In Arrow, K. J., Sen, A. K., Suzumura, K. (eds.), *Handbook of Social Choice and Welfare*, Vol. 1. Amsterdam (North-Holland): 35-94.

¹⁹Variants of the Condorcet jury theorem and results about the probability of Condorcet-style paradoxes have also been obtained in judgment-aggregation settings (e.g., Bovens and Rabinowicz [2], List [34]), but there is scope for further generalization here.

- [5] Dietrich, F. (2006) Judgment Aggregation: (Im)Possibility Theorems. *Journal of Economic Theory* 126(1): 286-298.
- [6] Dietrich, F. (2007) A generalised model of judgment aggregation. *Social Choice and Welfare* 28(4): 529-565.
- [7] Dietrich, F. (2007) Aggregation theory and the relevance of some issues to others. Working paper, University of Maastricht.
- [8] Dietrich, F. (2010) The possibility of judgment aggregation on agendas with subjunctive implications. *Journal of Economic Theory*, this issue.
- [9] Dietrich, F., List, C. (2007) Arrow's theorem in judgment aggregation. *Social Choice and Welfare* 29(1): 19-33.
- [10] Dietrich, F., List, C. (2007) Judgment aggregation by quota rules: majority voting generalized. *Journal of Theoretical Politics* 19(4): 391-424.
- [11] Dietrich, F., List, C. (2007) Strategy-proof judgment aggregation. *Economics and Philosophy* 23(3).
- [12] Dietrich, F., List, C. (2007) Opinion pooling on general agendas. Working paper, London School of Economics.
- [13] Dietrich, F., List, C. (2008) Judgment aggregation without full rationality. *Social Choice and Welfare* 31(1): 15-39.
- [14] Dietrich, F., List, C. (2008) A liberal paradox for judgment aggregation. *Social Choice and Welfare* 31(1): 59-78.
- [15] Dietrich, F., List, C. (2009) Propositionwise judgment aggregation: the general case. Working paper, London School of Economics.
- [16] Dietrich, F., List, C. (2010) Majority voting on restricted domains. *Journal of Economic Theory*, this issue.
- [17] Dietrich, F., List, C. (2010) The aggregation of propositional attitudes: towards a general theory. *Oxford Studies in Epistemology*.
- [18] Dietrich, F., List, C. (forthcoming) The impossibility of unbiased judgment aggregation. *Theory and Decision*.
- [19] Dietrich, F., List, C. (forthcoming) Judgment aggregation with consistency alone. *Social Choice and Welfare*.
- [20] Dietrich, F., Mongin, P. (2010) The Premiss-Based Approach to Judgment Aggregation. *Journal of Economic Theory*, this issue.
- [21] Dokow, E., Holzman, R. (2009) Aggregation of binary evaluations for truth-functional agendas. *Social Choice and Welfare* 32: 221-241.
- [22] Dokow, E., Holzman, R. (2009) Aggregation of non-binary evaluations. Working paper, Technion-Israel Institute of Technology.

- [23] Dokow, E., Holzman, R. (2010) Aggregation of binary evaluations. *Journal of Economic Theory*, this issue.
- [24] Dokow, E., Holzman, R. (2010) Aggregation of binary evaluations with abstentions. *Journal of Economic Theory*, this issue.
- [25] Duddy, C., Piggins, A. (2009) Many-valued judgment aggregation: characterizing the possibility/impossibility boundary for an important class of agendas. Working paper, National University of Ireland, Galway.
- [26] Gärdenfors, P. (2006) An Arrow-like theorem for voting with logical consequences. *Economics and Philosophy* 22(2): 181-190.
- [27] Guilbaud, G. Th. (1966) Theories of the General Interest, and the Logical Problem of Aggregation. In P. F. Lazarsfeld and N. W. Henry (eds.), *Readings in Mathematical Social Science*. Cambridge/MA (MIT Press): 262-307.
- [28] Konieczny, S., Pino Pérez, R. (2002) Merging Information Under Constraints: A Logical Framework. *Journal of Logic and Computation* 12(5): 773-808.
- [29] Kornhauser, L. A. (1992) Modeling Collegial Courts. II. Legal Doctrine. *Journal of Law, Economics and Organization* 8: 441-470.
- [30] Kornhauser, L. A., Sager, L. G. (1986) Unpacking the Court. *Yale Law Journal* 96(1): 82-117.
- [31] Kornhauser, L. A., Sager, L. G. (1993) The One and the Many: Adjudication in Collegial Courts. *California Law Review* 81: 1-59.
- [32] List, C. (2003) A Possibility Theorem on Aggregation over Multiple Interconnected Propositions. *Mathematical Social Sciences* 45(1): 1-13 (Corrigendum in *Mathematical Social Sciences* 52: 109-110).
- [33] List, C. (2004) A Model of Path-Dependence in Decisions over Multiple Propositions. *American Political Science Review* 98(3): 495-513.
- [34] List, C. (2005) The Probability of Inconsistencies in Complex Collective Decisions. *Social Choice and Welfare* 24(1): 3-32.
- [35] List, C. (forthcoming) The theory of judgment aggregation: An introductory review. *Synthese*.
- [36] List, C., Pettit, P. (2002) Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy* 18: 89-110.
- [37] List, C., Pettit, P. (2004) Aggregating Sets of Judgments: Two Impossibility Results Compared. *Synthese* 140(1-2): 207-235.
- [38] List, C., Puppe, C. (2009) Judgment aggregation: a survey. In Anand, P., Puppe, C., Pattanaik, P. (eds.), *Oxford Handbook of Rational and Social Choice*. Oxford (Oxford University Press).
- [39] Mas-Colell, A., Whinston, M. D., Green, J. R. (1995) *Microeconomic Theory*. New York (Oxford University Press).

- [40] McLean, I. (1990) The Borda and Condorcet Principles: Three Medieval Applications. *Social Choice and Welfare* 7(2): 99-108.
- [41] Miller, M., Osherson, D. (2009) Methods for distance-based judgment aggregation. *Social Choice and Welfare* 32(4): 575-601.
- [42] Mongin, P. (2008) Factoring Out the Impossibility of Logical Aggregation. *Journal of Economic Theory* 141(1): 100-113.
- [43] Nehring, K. (2003) Arrow's theorem as a corollary. *Economics Letters* 80(3): 379-382.
- [44] Nehring, K. (2010) Abstract Arrowian Aggregation: The Possibility-Impossibility Frontier. Working paper, University of California at Davis.
- [45] Nehring, K. (2006) The impossibility of a Paretian rational. Working paper, University of California at Davis.
- [46] Nehring, K. (2007) The impossibility of a Paretian rational: A Bayesian perspective. *Economics Letters* 96(1): 45-50.
- [47] Nehring, K., Puppe, C. (2002) Strategy-Proof Social Choice on Single-Peaked Domains: Possibility, Impossibility and the Space Between. Working paper, University of California at Davis.
- [48] Nehring, K., Puppe, C. (2008) Consistent judgement aggregation: the truth-functional case. *Social Choice and Welfare* 31(1): 41-57.
- [49] Nehring, K., Puppe, C. (2010) Abstract Arrowian Aggregation. *Journal of Economic Theory*, this issue.
- [50] Nehring, K., Puppe, C. (2010) Reasonable Group Choice. *Journal of Economic Theory*, this issue.
- [51] Osherson, D., Vardi, M. Y. (2006) Aggregating disparate estimates of chance. *Games and Economic Behavior* 56(1): 148-173.
- [52] Pauly, M., van Hees, M. (2006) Logical Constraints on Judgment Aggregation. *Journal of Philosophical Logic* 35: 569-585.
- [53] Pettit, P. (2001) Deliberative Democracy and the Discursive Dilemma. *Philosophical Issues* 11: 268-299.
- [54] Pigozzi, G. (2006) Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation. *Synthese* 152(2): 285-298.
- [55] Rothstein, P. (1990) Order Restricted Preferences and Majority Rule. *Social Choice and Welfare* 7(4): 331-342.
- [56] Rubinstein, A., Fishburn, P. (1986) Algebraic Aggregation Theory. *Journal of Economic Theory* 38: 63-77.
- [57] Sen, A. K. (1966) A Possibility Theorem on Majority Decisions. *Econometrica* 34(2): 491-499.

- [58] Sen, A. K. (1970/1979) *Collective Choice and Social Welfare*. San Francisco (Holden-Day).
- [59] Sen, A. K. (1995) Rationality and Social Choice. *American Economic Review* 85(1): 1-24.
- [60] van Hees, M. (2007) The limits of epistemic democracy. *Social Choice and Welfare* 28(4): 649-666.
- [61] Wilson, R. (1975) On the Theory of Aggregation. *Journal of Economic Theory* 10: 89-99.