

Social Choice Theory

Christian List

Social choice theory is the study of collective decision procedures. It is not a single theory, but a cluster of models and results concerning the aggregation of individual inputs (e.g., votes, preferences, judgments, welfare) into collective outputs (e.g., collective decisions, preferences, judgments, welfare). Central questions are: How can a group of individuals choose a winning outcome (e.g., policy, electoral candidate) from a given set of options? What are the properties of different voting systems? When is a voting system democratic? How can a collective (e.g., electorate, legislature, collegial court, expert panel, or committee) arrive at coherent collective preferences or judgments on some issues, on the basis of its members' individual preferences or judgments? How can we rank different social alternatives in an order of social welfare? Social choice theorists study these questions not just by looking at examples, but by developing general models and proving theorems.

Pioneered in the 18th century by Nicolas de Condorcet and Jean-Charles de Borda and in the 19th century by Charles Dodgson (also known as Lewis Carroll), social choice theory took off in the 20th century with the works of Kenneth Arrow, Amartya Sen, and Duncan Black. Its influence extends across economics, political science, philosophy, mathematics, and recently computer science and biology. Apart from contributing to our understanding of collective decision procedures, social choice theory has applications in the areas of institutional design, welfare economics, and social epistemology.

1. History of social choice theory

1.1 Condorcet

The two scholars most often associated with the development of social choice theory are the Frenchman Nicolas de Condorcet (1743-1794) and the American Kenneth Arrow (born 1921). Condorcet was a liberal thinker in the era of the French Revolution who was pursued by the revolutionary authorities for criticizing them. After a period of hiding, he was eventually arrested (though apparently not immediately identified), and died in prison (for more details on Condorcet, see McLean and Hewitt 1994). In his *Essay on the Application of Analysis to the Probability of Majority Decisions* (1785), he advocated a particular voting system, pairwise majority voting, and presented his two most prominent insights. The first, known as *Condorcet's jury theorem*, is that if each member of a jury has an equal and independent chance better than random, but worse than perfect, of making a correct judgment on whether a defendant is guilty (or on some other factual proposition), the majority of jurors is more likely to be correct than each individual juror, and the probability of a correct majority judgment approaches 1 as the jury size increases. Thus, under certain conditions, majority rule is good at 'tracking the truth' (e.g., Grofman, Owen, and Feld 1983; List and Goodin 2001).

Condorcet's second insight, often called *Condorcet's paradox*, is the observation that majority preferences can be 'irrational' (specifically, intransitive) even when individual preferences are 'rational' (specifically, transitive). Suppose, for example, that one third of a group prefers alternative x to y to z , a second third prefers y to z to

x , and a final third prefers z to x to y . Then there are majorities (of two thirds) for x against y , for y against z , and for z against x : a ‘cycle’, which violates transitivity. Furthermore, no alternative is a *Condorcet winner*, an alternative that beats, or at least ties with, every other alternative in pairwise majority contests.

Condorcet anticipated a key theme of modern social choice theory: majority rule is at once a plausible method of collective decision making and yet subject to some surprising problems. Resolving or bypassing these problems remains one of social choice theory’s core concerns.

1.2 Arrow and his influence

While Condorcet had investigated a particular voting method (majority voting), Arrow, who won the Nobel Memorial Prize in Economics in 1972, introduced a general approach to the study of preference aggregation, partly inspired by his teacher of logic Alfred Tarski (1901-1983) from whom he had learnt relation theory as an undergraduate at the City College of New York (Suppes 2005). Arrow considered a class of *possible* aggregation methods, which he called *social welfare functions*, and asked which of them satisfy certain axioms or desiderata. He proved that, surprisingly, there exists no method for aggregating the preferences of two or more individuals over three or more alternatives into collective preferences, where this method satisfies five seemingly plausible axioms, discussed below.

This result, known as *Arrow’s impossibility theorem*, prompted much work and many debates in social choice theory and welfare economics. William Riker (1920-1993), who inspired the Rochester school in political science, interpreted it as a mathematical proof of the impossibility of populist democracy (e.g., Riker 1982). Others, most prominently Amartya Sen (born 1933), who won the 1998 Nobel Memorial Prize, took it to show that ordinal preferences are insufficient for making satisfactory social choices. Commentators also questioned whether Arrow’s desiderata on an aggregation method are as innocuous as claimed or whether they should be relaxed.

The lessons from Arrow’s theorem depend, in part, on how we interpret an Arrowian social welfare function. The use of ordinal preferences as the ‘aggregenda’ may be easier to justify if we interpret the aggregation rule as a *voting method* than if we interpret it as a *welfare evaluation method*. Sen argued that, in the latter case (where a social planner seeks to rank different social alternatives in an order of social welfare), it may be justifiable to use additional information over and above ordinal preferences, such as interpersonally comparable welfare measurements (e.g., Sen 1982).

Arrow himself held the view ‘that interpersonal comparison of utilities has no meaning and ... that there is no meaning relevant to welfare comparisons in the measurability of individual utility’ (1951/1963, p. 9). This view was influenced by neoclassical economics, associated with scholars such as Vilfredo Pareto (1848-1923), Lionel Robbins (1898-1984), John Hicks (1904-1989), co-winner of the Economics Nobel Prize with Arrow, and Paul Samuelson (1915-2009), another Nobel Laureate. Arrow’s theorem demonstrates the implications of the ‘ordinalist’ assumptions of neoclassical thought.

Nowadays most social choice theorists have moved beyond the early negative interpretations of Arrow's theorem and are interested in the trade-offs involved in finding satisfactory decision procedures. Sen has promoted this 'possibilist' interpretation of social choice theory (e.g., in his 1998 Nobel lecture).

Within this approach, Arrow's axiomatic method is perhaps even more influential than his impossibility theorem (on the axiomatic method, see Thomson 2000). The paradigmatic kind of result in contemporary axiomatic work is the 'characterization theorem'. Here the aim is to identify a set of plausible necessary and sufficient conditions that uniquely characterize a particular solution (or class of solutions) to a given type of collective decision problem. An early example is Kenneth May's (1952) characterization of majority rule, discussed below.

1.3 Borda, Carroll, Black, and others

Condorcet and Arrow are not the only founding figures of social choice theory. Condorcet's contemporary and co-national Jean-Charles de Borda (1733-1799) defended a voting system that is often seen as a prominent alternative to majority voting. The *Borda count*, formally defined later, avoids Condorcet's paradox but violates one of Arrow's conditions, the *independence of irrelevant alternatives*. Thus the debate between Condorcet and Borda is a precursor to some modern debates on how to respond to Arrow's theorem.

The origins of this debate precede Condorcet and Borda. In the Middle Ages, Ramon Llull (c1235-1315) proposed the aggregation method of pairwise majority voting, while Nicolas Cusanus (1401-1464) proposed a variant of the Borda count (McLean 1990). In 1672, the German statesman and scholar Samuel von Pufendorf (1632-1694) compared simple majority, qualified majority, and unanimity rules and offered an analysis of the structure of preferences that can be seen as a precursor to later discoveries (e.g., on *single-peakedness*, discussed below) (Gaertner 2005).

In the 19th century, the British mathematician and clergyman Charles Dodgson (1832-1898), better known as Lewis Carroll, independently rediscovered many of Condorcet's and Borda's insights and also developed a theory of proportional representation. It was largely thanks to the Scottish economist Duncan Black (1908-1991) that Condorcet's, Borda's, and Dodgson's social-choice-theoretic ideas were drawn to the attention of the modern research community (McLean, McMillan, and Monroe 1995). Black also made several discoveries related to majority voting, some of which are discussed below.

In France, George-Théodule Guilbaud ([1952] 1966) wrote an important but often overlooked paper, revisiting Condorcet's theory of voting from a logical perspective and sparking a French literature on the *Condorcet effect*, the logical problem underlying Condorcet's paradox, which has only recently received more attention in Anglophone social choice theory (Monjardet 2005). For further contributions on the history of social choice, see McLean, McMillan, and Monroe (1996), McLean and Urken (1995), McLean and Hewitt (1994), and a special issue of *Social Choice and Welfare*, edited by Salles (2005).

2. Three formal arguments for majority rule

To introduce social choice theory formally, it helps to consider a simple decision problem: a collective choice between two alternatives.

2.1 The concept of an aggregation rule

Let $N = \{1, 2, \dots, n\}$ be a set of individuals, where $n \geq 2$. The individuals have to choose between two alternatives (candidates, policies etc.). Each individual $i \in N$ casts a vote, denoted v_i , where

$v_i = 1$ represents a vote for the first alternative,
 $v_i = -1$ represents a vote for the second alternative, and optionally
 $v_i = 0$ represents an abstention (for simplicity, we set this possibility aside).

A combination of votes across the individuals, $\langle v_1, v_2, \dots, v_n \rangle$, is called a *profile*. For any profile, the group seeks to arrive at a social decision v , where

$v = 1$ represents a decision for the first alternative,
 $v = -1$ represents a decision for the second alternative, and
 $v = 0$ represents a tie.

An *aggregation rule* is a function f that assigns to each profile $\langle v_1, v_2, \dots, v_n \rangle$ (in some domain of admissible profiles) a social decision $v = f(v_1, v_2, \dots, v_n)$. Examples are:

Majority rule: For each profile $\langle v_1, v_2, \dots, v_n \rangle$,

$$f(v_1, v_2, \dots, v_n) = \begin{cases} 1 & \text{if } v_1 + v_2 + \dots + v_n > 0 \text{ ('there are more 1s than -1s');} \\ 0 & \text{if } v_1 + v_2 + \dots + v_n = 0 \text{ ('there are as many 1s as -1s');} \\ -1 & \text{if } v_1 + v_2 + \dots + v_n < 0 \text{ ('there are more -1s than 1s').} \end{cases}$$

Dictatorship: For each profile $\langle v_1, v_2, \dots, v_n \rangle$,

$$f(v_1, v_2, \dots, v_n) = v_i,$$

where $i \in N$ is an antecedently fixed individual (the 'dictator').

Weighted majority rule: For each profile $\langle v_1, v_2, \dots, v_n \rangle$,

$$f(v_1, v_2, \dots, v_n) = \begin{cases} 1 & \text{if } w_1v_1 + w_2v_2 + \dots + w_nv_n > 0, \\ 0 & \text{if } w_1v_1 + w_2v_2 + \dots + w_nv_n = 0, \\ -1 & \text{if } w_1v_1 + w_2v_2 + \dots + w_nv_n < 0, \end{cases}$$

where w_1, w_2, \dots, w_n are real numbers, interpreted as the 'voting weights' of the n individuals.

Two points about the concept of an aggregation rule are worth noting. First, under the standard definition, an aggregation rule is defined *extensionally*, not *intensionally*: it is a *mapping (functional relationship)* between individual inputs and collective

outputs, not a *set of explicit instructions* (a *rule* in the ordinary-language sense) that could be extended to inputs outside the function's formal domain. Secondly, an aggregation rule is defined for a fixed set of individuals N and a fixed decision problem, so that majority rule in a group of two individuals is a different mathematical object from majority rule in a group of three.

To illustrate, Tables 1 and 2 show majority rule for these two group sizes as extensional objects. The rows of each table correspond to the different possible profiles of votes; the final column displays the resulting social decisions.

Table 1: Majority rule among two individuals

Individual 1's vote	Individual 2's vote	Collective decision
1	1	1
1	-1	0
-1	1	0
-1	-1	-1

Table 2: Majority rule among three individuals

Individual 1's vote	Individual 2's vote	Individual 3's vote	Collective decision
1	1	1	1
1	1	-1	1
1	-1	1	1
1	-1	-1	-1
-1	1	1	1
-1	1	-1	-1
-1	-1	1	-1
-1	-1	-1	-1

The present way of representing an aggregation rule helps us see how many possible aggregation rules there are (e.g., List 2011). Suppose there are k profiles in the domain of admissible inputs (in the present example, $k = 2^n$, since each of the n individuals has two choices, with abstention disallowed). Suppose, further, there are l possible social decisions for each profile (in the example, $l = 3$, allowing ties). Then there are l^k possible aggregation rules: the relevant table has k rows, and in each row, there are l possible ways of specifying the final entry (the collective decision). Thus the number of possible aggregation rules grows exponentially with the number of admissible profiles and the number of possible decision outcomes.

To select an aggregation rule non-arbitrarily from this large class of possible ones, some constraints are needed. I now consider three formal arguments for majority rule.

2.2 A procedural argument for majority rule

The first involves imposing some 'procedural' requirements on the relationship between individual votes and social decisions and showing that majority rule is the only aggregation rule satisfying them. May (1952) introduced four such requirements:

Universal domain: The domain of admissible inputs of the aggregation rule consists of all logically possible profiles of votes $\langle v_1, v_2, \dots, v_n \rangle$, where each $v_i \in \{-1, 1\}$.

Anonymity: For any admissible profiles $\langle v_1, v_2, \dots, v_n \rangle$ and $\langle w_1, w_2, \dots, w_n \rangle$ that are permutations of each other (i.e., one can be obtained from the other by reordering the entries), the social decision is the same, i.e., $f(v_1, v_2, \dots, v_n) = f(w_1, w_2, \dots, w_n)$.

Neutrality: For any admissible profile $\langle v_1, v_2, \dots, v_n \rangle$, if the votes for the two alternatives are reversed, the social decision is reversed too, i.e., $f(-v_1, -v_2, \dots, -v_n) = \neg f(v_1, v_2, \dots, v_n)$.

Positive responsiveness: For any admissible profile $\langle v_1, v_2, \dots, v_n \rangle$, if some voters change their votes in favour of one alternative (say the first) and all other votes remain the same, the social decision does not change in the opposite direction; if the social decision was a tie prior to the change, the tie is broken in the direction of the change, i.e., if $[w_i > v_i \text{ for some } i \text{ and } w_j = v_j \text{ for all other } j]$ and $f(v_1, v_2, \dots, v_n) = 0$ or 1 , then $f(w_1, w_2, \dots, w_n) = 1$.

Universal domain requires the aggregation rule to cope with any level of ‘pluralism’ in its inputs; anonymity requires it to treat all voters equally; neutrality requires it to treat all alternatives equally; and positive responsiveness requires the social decision to be a positive function of the way people vote. May proved the following:

Theorem (May 1952): An aggregation rule satisfies universal domain, anonymity, neutrality, and positive responsiveness if and only if it is majority rule.

Apart from providing an argument for majority rule based on four plausible procedural desiderata, the theorem helps us characterize other aggregation rules in terms of which desiderata they violate. Dictatorships and weighted majority rules with unequal individual weights violate anonymity. Asymmetrical supermajority rules (under which a supermajority of the votes, such as two thirds or three quarters, is required for a decision in favour of one of the alternatives, while the other alternative is the default choice) violate neutrality. This may sometimes be justifiable, for instance when there is a presumption in favour of one alternative, such as a presumption of innocence in a jury decision. Symmetrical supermajority rules (under which neither alternative is chosen unless it is supported by a sufficiently large supermajority) violate positive responsiveness. A more far-fetched example of an aggregation rule violating positive responsiveness is the inverse majority rule (here the alternative *rejected* by a majority wins).

2.3 An epistemic argument for majority rule

Condorcet’s jury theorem provides a consequentialist argument for majority rule. The argument is ‘epistemic’, insofar as the aggregation rule is interpreted as a truth-tracking device (e.g., Grofman, Owen and Feld 1983; List and Goodin 2001).

Suppose the aim is to make a judgment on some procedure-independent fact or state of the world, denoted X . In a jury decision, the defendant is either guilty ($X = 1$) or innocent ($X = -1$). In an expert-panel decision on the safety of some technology, the technology may be either safe ($X = 1$) or not ($X = -1$). Each individual’s vote expresses a judgment on that fact or state, and the social decision represents the collective judgment. The goal is to reach a factually correct collective judgment.

Which aggregation rule performs best at ‘truth tracking’ depends on the relationship between the individual votes and the relevant fact or state of the world.

Condorcet assumed that each individual is better than random at making a correct judgment (the competence assumption) and that different individuals’ judgments are stochastically independent, given the state of the world (the independence assumption). Formally, let V_1, V_2, \dots, V_n (capital letters) denote the random variables generating the specific individual votes v_1, v_2, \dots, v_n (small letters), and let $V = f(V_1, V_2, \dots, V_n)$ denote the resulting random variable representing the social decision $v = f(v_1, v_2, \dots, v_n)$ under a given aggregation rule f , such as majority rule. Condorcet’s assumptions can be stated as follows:

Competence: For each individual $i \in N$ and each state of the world $x \in \{-1, 1\}$, $\Pr(V_i = x \mid X = x) = p > 1/2$, where p is the same across individuals and states.

Independence: The votes of different individuals V_1, V_2, \dots, V_n are independent of each other, conditional on each value $x \in \{-1, 1\}$ of X .

Under these assumptions, majority voting is a good truth-tracker:

Theorem (Condorcet’s jury theorem): For each state of the world $x \in \{-1, 1\}$, the probability of a correct majority decision, $\Pr(V = x \mid X = x)$, is greater than each individual’s probability of a correct vote, $\Pr(V_i = x \mid X = x)$, and converges to 1, as the number of individuals n increases.¹

The first conjunct (‘is greater than each individual’s probability’) is the *non-asymptotic conclusion*, the second (‘converges to 1’) the *asymptotic conclusion*. One can further show that, if the two states of the world have an equal prior probability (i.e., $\Pr(X = 1) = \Pr(X = -1) = 1/2$), majority rule is the most reliable of all aggregation rules, maximizing $\Pr(V = X)$ (e.g., Ben-Yashar and Nitzan 1997).

Although the jury theorem is often invoked to establish the epistemic merits of democracy, its assumptions are highly idealistic. The competence assumption is not a conceptual claim but an empirical one and depends on any given decision problem. Although an average (not necessarily equal) individual competence above $1/2$ may be sufficient for Condorcet’s conclusion (e.g., Grofman, Owen, and Feld 1983, Boland 1989, Kanazawa 1998),² the theorem ceases to hold if individuals are randomizers (no better and no worse than a coin toss) or if they are worse than random ($p < 1/2$). In the latter case, the probability of a correct majority decision is *less* than each individual’s probability of a correct vote and converges to 0, as the jury size increases. The theorem’s conclusion can also be undermined in less extreme cases (Berend and Paroush 1998), for instance when each individual’s reliability, though above $1/2$, is an exponentially decreasing function approaching $1/2$ with increasing jury size (List 2003a).

¹ When n is even, the first part of the theorem only holds for group sizes n above a certain lower bound (which depends on p), due to the possibility of majority ties. When n is odd, it holds for any $n > 1$.

² If different individuals have different *known* levels of reliability, weighted majority voting outperforms simple majority voting at maximizing the probability of a correct decision, with each individual’s voting weight proportional to $\log(p/(1-p))$, where p is the individual’s reliability as defined above (Shapley and Grofman 1984; Grofman, Owen, and Feld 1983; Ben-Yashar and Nitzan 1997).

Similarly, whether the independence assumption is true depends on the decision problem in question. Although Condorcet's conclusion is robust to the presence of some interdependencies between individual votes, the structure of these interdependencies matters (e.g., Boland 1989; Ladha 1992; Estlund 1994; Dietrich and List 2004; Berend and Sapir 2007; Dietrich and Spiekermann 2013). If all individuals' votes are perfectly correlated with one another or mimic a small number of opinion leaders, the collective judgment is no more reliable than the judgment among a small number of independent individuals.

Bayesian networks, as employed in Pearl's work on causation (2000), have been used to model the effects of voter dependencies on the jury theorem and to distinguish between stronger and weaker variants of conditional independence (Dietrich and List 2004, Dietrich and Spiekermann 2013). Dietrich (2008) has argued that Condorcet's two assumptions are never simultaneously justified, in the sense that, even when they are both true, one cannot obtain evidence to support both at once.

Finally, game-theoretic work challenges an implicit assumption of the jury theorem, namely that voters will always reveal their judgments truthfully. Even if all voters prefer a correct to an incorrect collective judgment, they may still have incentives to misrepresent their individual judgments. This can happen when, conditional on the event of being pivotal for the outcome, a voter expects a higher chance of bringing about a correct collective judgment by voting *against* his or her own private judgment than in line with it (Austin-Smith and Banks 1996; Feddersen and Pesendorfer 1998).

2.4 A utilitarian argument for majority rule

Another consequentialist argument for majority rule is utilitarian rather than epistemic. It does not require the existence of an independent fact or state of the world that the collective decision is supposed to track. Suppose each voter gets some utility from the collective decision, which depends on whether the decision matches his or her vote (preference): specifically, each voter gets a utility of 1 from a match between his or her vote and the collective outcome and a utility of 0 from a mismatch.³ The Rae-Taylor theorem then states that if each individual has an equal prior probability of preferring each of the two alternatives, majority rule maximizes each individual's expected utility (see, e.g., Mueller 2003).

Relatedly, majority rule minimizes the number of frustrated voters (defined as voters on the losing side) and maximizes total utility across voters. Brighthouse and Fleurbaey (2010) generalize this result. Define voter i 's *stake* in the decision, δ_i , as the utility difference between his or her preferred outcome and his or her dispreferred outcome. The Rae-Taylor theorem rests on an implicit equal-stakes assumption, i.e., $\delta_i = 1$ for every $i \in N$. Brighthouse and Fleurbaey show that when stakes are allowed to vary across voters, total utility is maximized not by majority rule, but by a weighted majority rule, where each individual i 's voting weight w_i is proportional to his or her stake δ_i .

³ Optionally, one can stipulate that the utility from a tie is 1/2.

3. Preference aggregation

At the heart of social choice theory is the analysis of preference aggregation, understood as the aggregation of several individuals' preference rankings of two or more social alternatives into a single, collective preference ranking (or choice) over these alternatives.

The basic model is as follows. Again, consider a set $N = \{1, 2, \dots, n\}$ of individuals ($n \geq 2$). Let $X = \{x, y, z, \dots\}$ be a set of social alternatives, for example possible worlds, policy platforms, election candidates, or allocations of goods. Each individual $i \in N$ has a *preference ordering* R_i over these alternatives: a complete and transitive binary relation on X .⁴ For any $x, y \in X$, xR_iy means that individual i weakly prefers x to y . We write xP_iy if xR_iy and not yR_ix ('individual i strictly prefers x to y '), and xI_iy if xR_iy and yR_ix ('individual i is indifferent between x and y ').

A combination of preference orderings across the individuals, $\langle R_1, R_2, \dots, R_n \rangle$, is called a *profile*. A *preference aggregation rule*, F , is a function that assigns to each profile $\langle R_1, R_2, \dots, R_n \rangle$ (in some domain of admissible profiles) a social preference relation $R = F(R_1, R_2, \dots, R_n)$ on X . When F is clear from the context, we simply write R for the social preference relation corresponding to $\langle R_1, R_2, \dots, R_n \rangle$.

For any $x, y \in X$, xRy means that x is socially weakly preferred to y . We also write xPy if xRy and not yRx (' x is strictly socially preferred to y '), and xIy if xRy and yRx (' x and y are socially tied'). For generality, the requirement that R be complete and transitive is not built into the definition of a preference aggregation rule.

The paradigmatic example of a preference aggregation rule is *pairwise majority voting*, as discussed by Condorcet. Here, for any profile $\langle R_1, R_2, \dots, R_n \rangle$ and any $x, y \in X$, xRy if and only if at least as many individuals have xR_iy as have yR_ix , formally $|\{i \in N : xR_iy\}| \geq |\{i \in N : yR_ix\}|$. As we have seen, this does not guarantee transitive social preferences.⁵

How frequent are intransitive majority preferences? It can be shown that the proportion of preference profiles (among all possible ones) that lead to cyclical majority preferences increases with the number of individuals (n) and the number of alternatives ($|X|$). If all possible preference profiles are equally likely to occur (the so-called 'impartial culture' scenario), majority cycles should therefore be probable in large electorates (Gehrlein 1983). (Technical work further distinguishes between 'top-cycles' and cycles below a possible Condorcet-winning alternative.) However, the probability of cycles can be significantly lower under certain systematic, even small, deviations from an impartial culture (List and Goodin 2001, Appendix 3; Tsetlin, Regenwetter, and Grofman 2003; Regenwetter et al. 2006).

⁴ Completeness requires that, for any $x, y \in X$, xR_iy or yR_ix , and transitivity requires that, for any $x, y, z \in X$, if xR_iy and yR_iz , then xR_iz .

⁵ In the classic example, there are three individuals with preference orderings xP_1yP_1z , yP_2zP_2x , and zP_3xP_3y over three alternatives x, y , and z . The resulting majority preferences are cyclical: we have xPy , yRz , and yet zPx .

3.1 Arrow's theorem

Abstracting from pairwise majority voting, Arrow introduced the following conditions on a preference aggregation rule, F .

Universal domain: The domain of F is the set of all logically possible profiles of complete and transitive individual preference orderings.

Ordering: For any profile $\langle R_1, R_2, \dots, R_n \rangle$ in the domain of F , the social preference relation R is complete and transitive.

Weak Pareto principle: For any profile $\langle R_1, R_2, \dots, R_n \rangle$ in the domain of F , if for all $i \in N$ $xP_i y$, then xPy .

Independence of irrelevant alternatives: For any two profiles $\langle R_1, R_2, \dots, R_n \rangle$ and $\langle R^*_1, R^*_2, \dots, R^*_n \rangle$ in the domain of F and any $x, y \in X$, if for all $i \in N$ $xR_i y$ if and only if $xR^*_i y$, then xRy if and only if xR^*y .

Non-dictatorship: There does not exist an individual $i \in N$ such that, for all $\langle R_1, R_2, \dots, R_n \rangle$ in the domain of F and all $x, y \in X$, $xP_i y$ implies xPy .

Universal domain requires the aggregation rule to cope with any level of 'pluralism' in its inputs. Ordering requires it to produce 'rational' social preferences, avoiding Condorcet cycles. The weak Pareto principle requires that when all individuals strictly prefer alternative x to alternative y , so does society. Independence of irrelevant alternatives requires that the social preference between any two alternatives x and y depend only on the individual preferences between x and y , not on individuals' preferences over other alternatives. Non-dictatorship requires that there be no 'dictator', who always determines the social preference, regardless of other individuals' preferences. (Note that pairwise majority voting satisfies all of these conditions except ordering.)

Theorem (Arrow 1951/1963): If $|X| > 2$, there exists no preference aggregation rule satisfying universal domain, ordering, the weak Pareto principle, independence of irrelevant alternatives, and non-dictatorship.

It is evident that this result carries over to the aggregation of other kinds of orderings, as distinct from preference orderings, such as (i) belief orderings over several hypotheses (ordinal credences), (ii) multiple criteria that a single decision maker may use to generate an all-things-considered ordering of several decision options, and (iii) conflicting value rankings to be reconciled.

Examples of other such aggregation problems to which Arrow's theorem has been applied include: intrapersonal aggregation problems (e.g., May 1954; Hurley 1985), constraint aggregation in optimality theory in linguistics (e.g., Harbour and List 2000), theory choice (e.g., Okasha 2011; cf. Morreau forthcoming), evidence amalgamation (e.g., Stegenga 2011), and the aggregation of multiple similarity orderings into an all-things-considered similarity ordering (e.g., Morreau 2010, Kroedel and Huber 2012). In each case, the plausibility of Arrow's theorem depends on the case-specific plausibility of Arrow's ordinalist framework and the theorem's

conditions.

Generally, if we consider Arrow's framework appropriate and his conditions indispensable, Arrow's theorem raises a serious challenge. To avoid it, we must relax at least one of the five conditions or give up the restriction of the aggregation rule's inputs to orderings and defend the use of richer inputs, as discussed in Section 4.

3.2 Non-dictatorial preference aggregation rules

3.2.1 Relaxing universal domain

One way to avoid Arrow's theorem is to relax universal domain. If the aggregation rule is required to accept as input only preference profiles that satisfy certain 'cohesion' conditions, then aggregation rules such as pairwise majority voting will produce complete and transitive social preferences. The best-known cohesion condition is *single-peakedness* (Black 1948).

A profile $\langle R_1, R_2, \dots, R_n \rangle$ is *single-peaked* if the alternatives can be aligned from 'left' to 'right' (e.g., on some cognitive or ideological dimension) such that each individual has a most preferred position on that alignment with decreasing preference as alternatives get more distant (in either direction) from the most preferred position. Formally, this requires the existence of a linear ordering Ω on X such that, for every triple of alternatives $x, y, z \in X$, if y lies between x and z with respect to Ω , it is not the case that $xR_i y$ and $zR_i x$ (this rules out a 'cave' between x and z , at y). Single-peakedness is plausible in some democratic contexts. If the alternatives in X are different tax rates, for example, each individual may have a most preferred tax rate (which will be lower for a libertarian individual than for a socialist) and prefer other tax rates less as they get more distant from the ideal.

Black (1948) proved that if the domain of the aggregation rule is restricted to the set of all profiles of individual preference orderings satisfying single-peakedness, majority cycles cannot occur, and the most preferred alternative of the median individual relative to the relevant left-right alignment is a Condorcet winner (assuming n is odd). Pairwise majority voting then satisfies the rest of Arrow's conditions.

Other domain-restriction conditions with similar implications include *single-cavedness*, a geometrical mirror image of single-peakedness (Inada 1964), *separability into two groups* (ibid.), and *latin-squarelessness* (Ward 1965), the latter two more complicated combinatorial conditions. (For a review, see Gaertner 2001.) Sen (1966) showed that all these conditions imply a weaker condition, *triple-wise value-restriction*. It requires that, for every triple of alternatives $x, y, z \in X$, there exists one alternative in $\{x, y, z\}$ and one rank $r \in \{1, 2, 3\}$ such that no individual ranks that alternative in r^{th} place among x, y , and z . For instance, all individuals may agree that y is not bottom ($r = 3$) among x, y , and z . Triple-wise value-restriction suffices for transitive majority preferences.

There has been much discussion on whether, and under what conditions, real-world preferences fall into such a restricted domain. It has been suggested, for example, that group deliberation can induce single-peaked preferences, by leading participants to

focus on a shared cognitive or ideological dimension (Miller 1992; Knight and Johnson 1995; Dryzek and List 2003). Experimental evidence from deliberative opinion polls is consistent with this hypothesis (List, Luskin, Fishkin, and McLean 2013), though further empirical work is needed.

3.2.2 Relaxing ordering

Preference aggregation rules are normally expected to produce orderings as their outputs, but sometimes we may only require partial orderings or not fully transitive binary relations. An aggregation rule that produces transitive but often incomplete social preferences is the *Pareto dominance procedure*: here, for any profile $\langle R_1, R_2, \dots, R_n \rangle$ and any $x, y \in X$, xRy if and only if, for all $i \in N$, xP_iy . An aggregation rule that produces complete but often intransitive social preferences is the *Pareto extension procedure*: here, for any profile $\langle R_1, R_2, \dots, R_n \rangle$ and any $x, y \in X$, xRy if and only if it is not the case that, for all $i \in N$, yP_ix . Both rules have a unanimitarian spirit, giving each individual veto power either against the *presence* of a weak social preference for x over y or against its *absence*.

Gibbard (1969) proved that even if we replace the requirement of transitivity with what he called *quasi-transitivity*, the resulting possibilities of aggregation are still very limited. Call a preference relation R *quasi-transitive* if the induced strict relation P is transitive (while the indifference relation I need not be transitive). Call an aggregation rule *oligarchic* if there is a subset $M \subseteq N$ (the ‘oligarchs’) such that (i) if, for all $i \in M$, xP_iy , then xPy , and (ii) if, for some $i \in M$, xP_iy , then xRy . The Pareto extension procedure is an example of an oligarchic aggregation rule with $M = N$. In an oligarchy, the oligarchs are jointly decisive and have individual veto power. Gibbard proved the following:

Theorem (Gibbard 1969): If $|X| > 2$, there exists no preference aggregation rule satisfying universal domain, quasi-transitivity and completeness of social preferences, the weak Pareto principle, independence of irrelevant alternatives, and non-oligarchy.

3.2.3 Relaxing the weak Pareto principle

The weak Pareto principle is arguably hard to give up. One case in which we may lift it is that of spurious unanimity, where a unanimous preference for x over y is based on mutually inconsistent reasons (e.g., Mongin 1997; Gilboa, Samet, and Schmeidler 2004). Two men may each prefer to fight a duel (alternative x) to not fighting it (alternative y) because each over-estimates his chances of winning. There may exist no mutually agreeable probability assignment over possible outcomes of the duel (i.e., who would win) that would ‘rationalize’ the unanimous preference for x over y . In this case, the unanimous preference is a bad indicator of social preferability. This example, however, depends on the fact that the alternatives of fighting and not fighting are not fully specified outcomes but uncertain prospects. Arguably, the weak Pareto principle is more plausible in cases without uncertainty.

An aggregation rule that becomes possible when the weak Pareto principle is dropped is an *imposed procedure*, where, for any profile $\langle R_1, R_2, \dots, R_n \rangle$, the social preference relation R is an antecedently fixed (‘imposed’) ordering $R_{imposed}$ of the alternatives. Though completely unresponsive to individual preferences, this aggregation rule

satisfies the rest of Arrow's conditions.

Sen (1970a) offered another critique of the weak Pareto principle, showing that it conflicts with a 'liberal' principle. Here we interpret the aggregation rule as a method a social planner can use to rank social alternatives in an order of social welfare. Suppose each individual in society is given some basic rights, to the effect that his or her preference is sometimes socially decisive (i.e., cannot be overridden by others' preferences). Each of Lewd and Prude, for example, should be decisive over whether he himself reads a particular book, *Lady Chatterley's Lover*.

Minimal liberalism: There are *at least* two distinct individuals $i, j \in N$ who are each decisive on at least one pair of alternatives; i.e., there is at least one pair of alternatives $x, y \in X$ such that, for every profile $\langle R_1, R_2, \dots, R_n \rangle$, $xP_i y$ implies xPy , and $yP_j x$ implies yPx , and at least one pair of alternatives $x^*, y^* \in X$ such that, for every profile $\langle R_1, R_2, \dots, R_n \rangle$, $x^*P_j y^*$ implies x^*Py^* , and $y^*P_i x^*$ implies y^*Px^* .

Sen asked us to imagine that Lewd most prefers that Prude read the book (alternative x), second-most prefers that he read the book himself (alternative y), and least prefers that neither read the book (z). Prude most prefers that neither read the book (z), second-most prefers that he read the book himself (x), and least prefers that Lewd read the book (y). Assuming Lewd is decisive over the pair y and z , society should prefer y to z . Assuming Prude is decisive over the pair x and z , society should prefer z to x . But since Lewd and Prude both prefer x to y , the weak Pareto principle (applied to $N = \{\text{Lewd, Prude}\}$) implies that society should prefer x to y , and so we are faced with a social preference cycle. Sen called this problem the 'liberal paradox' and generalized it as follows.

Theorem (Sen 1970a): There exists no preference aggregation rule satisfying universal domain, acyclicity of social preferences, the weak Pareto principle, and minimal liberalism.

The result suggests that if we wish to respect individual rights, we may sometimes have to sacrifice Paretian efficiency. An alternative conclusion is that the weak Pareto principle can be rendered compatible with minimal liberalism only when the domain of admissible preference profiles is suitably restricted, for instance to preferences that are 'tolerant' or not 'meddlesome' (Blau 1975, Craven 1982, Gigliotti 1986, Sen 1983). Lewd's and Prude's preferences in Sen's example are 'meddlesome'.

Several authors have challenged the relevance of Sen's result, however, by criticizing his formalization of rights (e.g., Gaertner, Pattanaik, and Suzumura 1992, Dowding and van Hees 2003).

3.2.3 Relaxing independence of irrelevant alternatives

A common way to obtain possible preference aggregation rules is to give up independence of irrelevant alternatives. Almost all familiar voting methods over three or more alternatives that involve some form of preferential voting (with voters being asked to express full or partial preference orderings) violate this condition.

A standard example is *plurality rule*: here, for any profile $\langle R_1, R_2, \dots, R_n \rangle$ and any x ,

$y \in X$, xRy if and only if $|\{i \in N : \text{for all } z \neq x, xP_i z\}| \geq |\{i \in N : \text{for all } z \neq y, yP_i z\}|$. Informally, alternatives are socially ranked in the order of how many individuals most prefer each of them. Plurality rule avoids Condorcet's paradox, but runs into other problems. Most notably, an alternative that is majority-dispreferred to every other alternative may win under plurality rule: if 34% of the voters rank x above y above z , 33% rank y above z above x , and 33% rank z above y above x , plurality rule ranks x above each of y and z , while pairwise majority voting would rank y above z above x (y is the Condorcet winner). By disregarding individuals' lower-ranked alternatives, plurality rule also violates the weak Pareto principle. However, plurality rule may be plausible in 'restricted informational environments', where the balloting procedure collects information only about voters' top preferences, not about their full preference rankings. Here plurality rule satisfies generalized variants of May's four conditions introduced above (Goodin and List 2006).

A second example of a preference aggregation rule that violates independence of irrelevant alternatives is the *Borda count* (e.g., Saari 1990). Here, for any profile $\langle R_1, R_2, \dots, R_n \rangle$ and any $x, y \in X$, xRy if and only if $|\sum_{i \in N} \{z \in X : xR_i z\}| \geq |\sum_{i \in N} \{z \in X : yR_i z\}|$. Informally, each voter assigns a score to each alternative, which depends on its rank in his or her preference ranking. The most-preferred alternative gets a score of k (where $k = |X|$), the second-most-preferred alternative a score of $k-1$, the third-most-preferred alternative a score of $k-2$, and so on. Alternatives are then socially ordered in terms of the sums of their scores across voters: the alternative with the largest sum-total is top, the alternative with the second-largest sum-total next, and so on.

To see how this violates independence of irrelevant alternatives, consider the two profiles of individual preference orderings over four alternatives (x, y, z, w) in Tables 3 and 4.

Table 3: A profile of individual preference orderings

	Individual 1	Individuals 2 to 7	Individuals 8 to 15
1 st preference	y	x	z
2 nd preference	x	z	x
3 rd preference	z	w	y
4 th preference	w	y	w

Table 4: A slightly modified profile of individual preference orderings

	Individual 1	Individuals 2 to 7	Individuals 8 to 15
1 st preference	x	x	z
2 nd preference	y	z	x
3 rd preference	w	w	y
4 th preference	z	y	w

In Table 3, the Borda scores of the four alternatives are:

$$\begin{aligned} x: & 9 \cdot 3 + 6 \cdot 4 = 51, \\ y: & 1 \cdot 4 + 6 \cdot 1 + 8 \cdot 2 = 26, \\ z: & 1 \cdot 2 + 6 \cdot 3 + 8 \cdot 4 = 52, \\ w: & 1 \cdot 1 + 6 \cdot 2 + 8 \cdot 1 = 21, \end{aligned}$$

leading to a social preference for z over x over y over w . In Table 4 the Borda scores

are:

$$\begin{aligned} x: & 7*4 + 8*3 = 52, \\ y: & 1*3 + 6*1 + 8*2 = 25, \\ z: & 1*1 + 6*3 + 8*4 = 51, \\ w: & 7*2 + 8*1 = 22, \end{aligned}$$

leading to a social preference for x over z over y over w . The only difference between the two profiles lies in Individual 1's preference ordering, and even here there is no change in the relative ranking of x and z . Despite identical individual preferences between x and z in Tables 3 and 4, the social preference between x and z is reversed, a violation of independence of irrelevant alternatives.

Such violations are common in real-world voting rules, and they make preference aggregation potentially vulnerable to strategic voting and/or strategic agenda setting. I illustrate this in the case of strategic voting.

3.3 The Gibbard-Satterthwaite theorem

So far we have discussed *preference aggregation rules*, which map profiles of individual preference orderings to social preference relations. We now consider *social choice rules*, whose output, instead, is one or several winning alternatives. Formally, a *social choice rule*, f , is a function that assigns to each profile $\langle R_1, R_2, \dots, R_n \rangle$ (in some domain of admissible profiles) a social choice set $f(R_1, R_2, \dots, R_n) \subseteq X$. A social choice rule f can be derived from a preference aggregation rule F , by defining $f(R_1, R_2, \dots, R_n) = \{x \in X : \text{for all } y \in X, xRy\}$ where $R = F(R_1, R_2, \dots, R_n)$; the reverse does not generally hold. We call the set of sometimes-chosen alternatives the *range* of f .⁶

The *Condorcet winner criterion* defines a social choice rule, where, for each profile $\langle R_1, R_2, \dots, R_n \rangle$, $f(R_1, R_2, \dots, R_n)$ contains every alternative in X that wins or at least ties with every other alternative in pairwise majority voting. As shown by Condorcet's paradox, this may produce an empty choice set. By contrast, plurality rule and the Borda count induce social choice rules that always produce non-empty choice sets. They also satisfy the following basic conditions (the last for $|X| \geq 3$):

Universal Domain: The domain of f is the set of all logically possible profiles of complete and transitive individual preference orderings.

Non-dictatorship: There does not exist an individual $i \in N$ such that, for all $\langle R_1, R_2, \dots, R_n \rangle$ in the domain of f and all x in the range of f , $yR_i x$ where $y \in f(R_1, R_2, \dots, R_n)$.⁷

The range constraint: The range of f contains at least three distinct alternatives (and ideally all alternatives in X).

When supplemented with an appropriate tie-breaking criterion, the plurality and Borda rules can further be made 'resolute':

⁶ Formally, $\{x \in X : x \in f(R_1, R_2, \dots, R_n) \text{ for some } \langle R_1, R_2, \dots, R_n \rangle \text{ in the domain of } f\}$.

⁷ For present purposes, one can stipulate that the last clause (for all x in the range of f , $yR_i x$ where $y \in f(R_1, R_2, \dots, R_n)$) is violated if $f(R_1, R_2, \dots, R_n)$ is empty.

Resoluteness: The social choice rule f always produces a unique winning alternative (a singleton choice set). (We then write $x = f(R_1, R_2, \dots, R_n)$ to denote the winning alternative for the profile $\langle R_1, R_2, \dots, R_n \rangle$.)

Surprisingly, this list of conditions conflicts with the following further requirement.

Strategy-proofness: There does not exist a profile $\langle R_1, R_2, \dots, R_n \rangle$ in the domain of f at which f is manipulable by some individual $i \in N$, where *manipulability* means the following: if i submits a *false* preference ordering R'_i ($\neq R_i$), the winner is an alternative y' that i strictly prefers (according to R_i) to the alternative y that would win if i submitted the *true* preference ordering R_i .⁸

Theorem (Gibbard 1973, Satterthwaite 1975): There exists no social choice rule satisfying universal domain, non-dictatorship, the range constraint, resoluteness, and strategy-proofness.

This result raises important questions about the trade-offs between different requirements on a social choice rule. A dictatorship, which always chooses the dictator's most preferred alternative, is trivially strategy-proof. The dictator obviously has no incentive to vote strategically, and no-one else does so either, since the outcome depends only on the dictator.

To see that the Borda count violates strategy-proofness, recall the example of Tables 3 and 4 above. If Individual 1 in Table 3 truthfully submits the preference ordering $yP_1xP_1zP_1w$, the Borda winner is z , as we have seen. If Individual 1 falsely submits the preference ordering $xP_1yP_1wP_1z$, as in Table 4, the Borda winner is x . But Individual 1 prefers x to z according to his or her true preference ordering (in Table 3), and so he or she has an incentive to vote strategically.

Moulin (1980) has shown that when the domain of the social choice rule is restricted to single-peaked preference profiles, pairwise majority voting and other so-called 'median voting' schemes can satisfy the rest of the conditions of the Gibbard-Satterthwaite theorem. Similarly, when collective decisions are restricted to binary choices alone, which amounts to dropping the range constraint, majority voting satisfies the rest of the conditions. Other possible escape routes from the theorem open up if resoluteness is dropped. In the limiting case in which all alternatives are always chosen, the other conditions are vacuously satisfied.

The requirement of strategy-proofness has been challenged too. One line of argument is that, even when there exist strategic incentives in the technical sense of the Gibbard-Satterthwaite theorem, individuals will not necessarily act on them. They would require detailed information about others' preferences and enough computational power to figure out what the optimal strategically modified preferences would be. Neither demand is generally met. Bartholdi, Tovey, and Trick (1989) showed that, due to computational complexity, some social choice rules are resistant

⁸ Formally, $y'Py$, where $y' = f(R_1, \dots, R'_i, \dots, R_n)$ and $y = f(R_1, \dots, R_i, \dots, R_n)$, assuming that $\langle R_1, \dots, R'_i, \dots, R_n \rangle$ is in the domain of f . The definition presupposes that the social choice sets for the profiles $\langle R_1, \dots, R_i, \dots, R_n \rangle$ and $\langle R_1, \dots, R'_i, \dots, R_n \rangle$ are singleton.

to strategic manipulation: it may be an NP-hard problem for a voter to determine how to vote strategically. Harrison and McDaniel (2008) provide experimental evidence suggesting that the ‘Kemeny rule’, an extension of pairwise majority voting designed to avoid Condorcet cycles, is ‘behaviourally incentive-compatible’: i.e., strategic manipulation is computationally hard.

Dowding and van Hees (2008) have argued that not all forms of strategic voting are normatively problematic. They distinguish between ‘sincere’ and ‘insincere’ forms of manipulation and argue that only the latter but not the former are normatively troublesome. Sincere manipulation involves (i) voting for a compromise alternative whose chances of winning are thereby increased, where (ii) one prefers the compromise alternative to the alternative that would otherwise win. Supporters of Ralph Nader, a third-party US presidential candidate in 2000 with little chance of winning, who voted in favour of Al Gore to increase his chances of beating George W. Bush engaged in sincere manipulation in the sense of (i) and (ii). Plurality rule is susceptible to sincere manipulation, but not vulnerable to insincere manipulation.

4. Welfare aggregation

An implicit assumption so far has been that preferences are ordinal and not interpersonally comparable: preference orderings contain no information about each individual’s strength of preference or about how to compare different individuals’ preferences with one another. Statements such as ‘Individual 1 prefers alternative x more than Individual 2 prefers alternative y ’ or ‘Individual 1 prefers a switch from x to y more than Individual 2 prefers a switch from x^* to y^* ’ are considered meaningless.

In voting contexts, this assumption may be plausible, but in welfare-evaluation contexts – when a social planner seeks to rank different social alternatives in an order of social welfare – the use of richer information may be justified. Sen (1970b) generalized Arrow’s model to incorporate such richer information.

As before, consider a set $N = \{1, 2, \dots, n\}$ of individuals ($n \geq 2$) and a set $X = \{x, y, z, \dots\}$ of social alternatives. Now each individual $i \in N$ has a *welfare function* W_i over these alternatives, which assigns a real number $W_i(x)$ to each alternative $x \in X$, interpreted as a measure of i ’s welfare under alternative x . Any welfare function on X induces an ordering on X , but the converse is not true: welfare functions encode more information. A combination of welfare functions across the individuals, $\langle W_1, W_2, \dots, W_n \rangle$, is called a *profile*.

A *social welfare functional (SWFL)*, also denoted F , is a function that assigns to each profile $\langle W_1, W_2, \dots, W_n \rangle$ (in some domain of admissible profiles) a social preference relation $R = F(W_1, W_2, \dots, W_n)$ on X , with the familiar interpretation. Again, when F is clear from the context, we write R for the social preference relation corresponding to $\langle W_1, W_2, \dots, W_n \rangle$. The *output* of a SWFL is similar to that of a preference aggregation rule (again, we do not build the completeness or transitivity of R into the definition⁹), but its *input* is richer.

⁹ Sen, like Arrow in his definition of social welfare functions (as opposed to functionals), required R to be an ordering by definition.

What we gain from this depends on how much of the enriched informational input we allow ourselves to use in determining society's preferences: technically, it depends on our assumption about measurability and interpersonal comparability of welfare.

4.1 Measurability and interpersonal comparability of welfare

Numerically, welfare profiles contain much information over and above the profiles of orderings on X they induce. But we may not consider all this information meaningful. Some of it could be an artifact of the numerical representation. For example, the difference between the profile $\langle W_1, W_2, \dots, W_n \rangle$ and its scaled-up version $\langle 10*W_1, 10*W_2, \dots, 10*W_n \rangle$, where everything is the same in proportional terms, could be like the difference between length measurements in centimeters and in inches.

Some examples of statements about individual welfare that may, or may not, be meaningful are the following (List 2003b; see also Bossert and Weymark 1996, Section 5):

A level comparison: Individual i 's welfare under alternative x is at least as great as individual j 's welfare under alternative y , formally $W_i(x) \geq W_j(y)$.

The comparison is *intrapersonal* if $i = j$, and *interpersonal* if $i \neq j$.

A unit comparison: The ratio of [individual i 's welfare gain or loss if we switch from alternative y_1 to alternative x_1] to [individual j 's welfare gain or loss if we switch from alternative y_2 to alternative x_2] is λ , where λ is some real number, formally $(x_1 - y_1) / (x_2 - y_2) = \lambda$.

Again, the comparison is *intrapersonal* if $i = j$, and *interpersonal* if $i \neq j$.

A zero comparison: Individual i 's welfare under alternative x is greater than / equal to / less than zero, formally $sign(W_i(x)) = \lambda$, where $\lambda \in \{-1, 0, 1\}$ and $sign$ is a real-valued function that maps strictly negative numbers to -1 , zero to 0 , and strictly positive numbers to $+1$.

Arrow's view, as noted, is that only intrapersonal level comparisons are meaningful, while all other kinds of comparisons are not. Sen (1970b) formalized various assumptions about measurability and interpersonal comparability of welfare by (i) defining an equivalence relation on welfare profiles that specifies when two profiles count as 'containing the same information', and (ii) requiring any profiles in the same equivalence class to generate the same social preference ordering. Of the kinds of comparisons introduced above, the meaningful ones are those that are invariant in each equivalence class. Arrow's ordinalist assumption can be expressed as follows:

Ordinal measurability with no interpersonal comparability (ONC): Two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ contain the same information whenever, for each $i \in N$, $W^*_i = \phi_i(W_i)$, where ϕ_i is some positive monotonic transformation, possibly different for different individuals.

Thus the individual welfare functions in any profile can be arbitrarily monotonically

transformed ('stretched or squeezed') without informational loss, thereby ruling out any interpersonal comparisons or even intrapersonal unit comparisons.

If welfare is cardinally measurable but still interpersonally non-comparable, we have:

Cardinal measurability with no interpersonal comparability (CNC): Two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ contain the same information whenever, for each $i \in N$, $W^*_i = a_i W_i + b_i$, where the a_i s and b_i s are real numbers (with $a_i > 0$), possibly different for different individuals.

Here, each individual's welfare function is unique up to positive affine transformations ('scaling and shifting'), but there is still no common scale across individuals. This secures intrapersonal level and unit comparisons, but rules out interpersonal comparisons and zero comparisons.

Interpersonal level comparability is achieved under the following enriched variant of ordinal measurability:

Ordinal measurability with interpersonal level comparability (OLC): Two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ contain the same information whenever, for each $i \in N$, $W^*_i = \phi(W_i)$, where ϕ is the same positive monotonic transformation for all individuals.

Here, a profile of individual welfare functions can be arbitrarily monotonically transformed ('stretched or squeezed') without informational loss, but the same transformation must be used for all individuals, thereby securing interpersonal level comparability.

Interpersonal unit comparability is achieved under the following enriched variant of cardinal measurability:

Cardinal measurability with interpersonal unit comparability (CUC): Two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ contain the same information whenever, for each $i \in N$, $W^*_i = aW_i + b_i$, where a is the same real number for all individuals ($a > 0$) and the b_i s are real numbers.

Here, the welfare functions in each profile can be re-scaled and shifted without informational loss, but the same scalar multiple (though not necessarily the same shifting constant) must be used for all individuals, thereby rendering interpersonal unit comparisons invariant in each equivalence class.

Zero comparisons, finally, become meaningful under the following enriched variant of ordinal measurability (List 2001):

Ordinal measurability with zero comparability (ONC+0): Two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ contain the same information whenever, for each $i \in N$, $W^*_i = \phi_i(W_i)$, where ϕ_i is some positive monotonic and zero-preserving transformation, possibly different for different individuals. (Here *zero-preserving* means that $\phi_i(0) = 0$.)

This allows arbitrary stretching and squeezing of individual welfare functions without informational loss, *provided the welfare level of zero remains fixed*, thereby ensuring zero comparability.

Several other measurability and interpersonal comparability assumptions have been discussed in the literature. The following ensures interpersonal comparability of both levels and units:

Cardinal measurability with full interpersonal comparability (CFC): Two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ contain the same information whenever, for each $i \in N$, $W^*_i = aW_i + b$, where a, b are the same real numbers for all individuals ($a > 0$).

Lastly, intra- and interpersonal comparisons of all *three* kinds (level, unit, and zero) are meaningful if we accept the following:

Ratio-scale measurability with full interpersonal comparability (RFC): Two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ contain the same information whenever, for each $i \in N$, $W^*_i = aW_i$, where a is the same real number for all individuals ($a > 0$).

Which assumption is warranted depends on how welfare is interpreted. If welfare is *hedonic utility*, which can be experienced only from a first-person perspective, interpersonal comparisons are harder to justify than if welfare is the *objective satisfaction of subjective preferences or desires* (the *desire-satisfaction* view) or an *objective good or state* (an *objective-list* view) (e.g., Hausman 1995, List 2003b). The desire-satisfaction view may render interpersonal comparisons empirically meaningful (by relating the interpersonally significant maximal and minimal levels of welfare for each individual to the attainment of his or her most and least preferred alternatives), but at the expense of running into problems of expensive tastes or adaptive preferences (Hausman 1995). Resource-based, functioning-based, or primary-goods-based currencies of welfare, by contrast, may allow empirically meaningful and less morally problematic interpersonal comparisons.

4.2 The possibility of welfare aggregation

Once we introduce interpersonal comparisons of welfare levels or units or zero comparisons, there exist possible SWFLs satisfying the analogues of Arrow's conditions as well as stronger desiderata. In a welfare-aggregation context, Arrow's impossibility can therefore be traced to a lack of interpersonal comparability.

As noted, a SWFL respects a given assumption about measurability and interpersonal comparability if, for any two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ that are deemed to contain the same information, $F(W_1, W_2, \dots, W_n) = F(W^*_1, W^*_2, \dots, W^*_n)$. Arrow's conditions and theorem can be restated as follows:

Universal domain: The domain of F is the set of all logically possible profiles of individual welfare functions.

Ordering: For any profile $\langle W_1, W_2, \dots, W_n \rangle$ in the domain of F , the social preference relation R is complete and transitive.

Weak Pareto principle: For any profile $\langle W_1, W_2, \dots, W_n \rangle$ in the domain of F , if for all $i \in N$ $W_i(x) > W_i(y)$, then xPy .

Independence of irrelevant alternatives: For any two profiles $\langle W_1, W_2, \dots, W_n \rangle$ and $\langle W^*_1, W^*_2, \dots, W^*_n \rangle$ in the domain of F and any $x, y \in X$, if for all $i \in N$ $W_i(x) = W^*_i(x)$ and $W_i(y) = W^*_i(y)$, then xRy if and only if xR^*y .

Non-dictatorship: There does not exist an individual $i \in N$ such that, for all $\langle W_1, W_2, \dots, W_n \rangle$ in the domain of F and all $x, y \in X$, $W_i(x) > W_i(y)$ implies xPy .

Theorem: Under ONC (or CNC, as Sen 1970b has shown), if $|X| > 2$, there exists no SWFL satisfying universal domain, ordering, the weak Pareto principle, independence of irrelevant alternatives, and non-dictatorship.

Crucially, however, each of OLC, CUC, and ONC+0 is sufficient for the existence of SWFLs satisfying all other conditions:

Theorem: Under each of OLC, CUC, and ONC+0, there exist SWFLs satisfying universal domain, ordering, the weak Pareto principle, independence of irrelevant alternatives, and non-dictatorship (as well as stronger conditions).

Some examples of such SWFLs come from political philosophy and welfare economics. A possible SWFL under OLC is a version of Rawls's difference principle (1971).

Maximin: For any profile $\langle W_1, W_2, \dots, W_n \rangle$ and any $x, y \in X$, xRy if and only if $\min_{i \in N}(W_i(x)) \geq \min_{i \in N}(W_i(y))$.

While maximin rank-orders social alternatives in terms of the welfare level of the worst-off individual alone, its lexicographic extension (*leximin*), which was endorsed by Rawls himself, uses the welfare level of the second-worst-off individual as a tie-breaker when there is tie at the level of the worst off, the welfare level of the third-worst-off individual as a tie-breaker when there is a tie at the second stage, and so on. (Note, however, that Rawls focused on primary goods, rather than welfare, as the relevant 'currency'.) This satisfies the *strong* (not just weak) *Pareto principle*, requiring that if for all $i \in N$ $W_i(x) \geq W_i(y)$, then xRy , and if in addition for some $i \in N$ $W_i(x) > W_i(y)$, then xPy .

An example of a possible SWFL under CUC is classical utilitarianism.

Utilitarianism: For any profile $\langle W_1, W_2, \dots, W_n \rangle$ and any $x, y \in X$, xRy if and only if $W_1(x) + W_2(x) + \dots + W_n(x) \geq W_1(y) + W_2(y) + \dots + W_n(y)$.

Finally, an example of a possible SWFL under ONC+0 is a variant of a frequently used, though simplistic poverty measure.

A head-count rule: For any profile $\langle W_1, W_2, \dots, W_n \rangle$ and any $x, y \in X$, xRy if and only if $|\{i \in N : W_i(x) < 0\}| < |\{i \in N : W_i(y) < 0\}|$ or $[|\{i \in N : W_i(x) < 0\}| = |\{i \in N : W_i(y) < 0\}|$

$< 0\}$] and xR_jy], where $j \in N$ is some antecedently fixed tie-breaking individual.

While substantively less compelling than maximin or utilitarian rules, head-count rules require only zero-comparability of welfare (List 2001).

An important conclusion, therefore, is that Rawls's difference principle, the classical utilitarian principle, and even the head-count method of poverty measurement can all be seen as solutions to Arrow's aggregation problem that become possible once we go beyond Arrow's framework of ordinal, interpersonally non-comparable preferences.

Under CFC, one can provide a simultaneous characterization of Rawlsian maximin and utilitarianism (Deschamps and Gevers 1978). It uses two additional axioms. One, *minimal equity*, requires (in the words of Sen 1977, p. 1548) 'that a person who is going to be best off anyway does not *always* strictly have his way', and another, *separability*, requires that two welfare profiles that coincide for some subset $M \subseteq N$ while everyone in $N \setminus M$ is indifferent between all alternatives in X lead to the same social ordering.

Theorem (Deschamps and Gevers 1978): Under CFC, any SWFL satisfying universal domain, ordering, the strong Pareto principle, independence of irrelevant alternatives, anonymity (as in May's theorem), minimal equity, and separability is either leximin or of a utilitarian type (meaning that, except possibly when there are ties in sum-total welfare, it coincides with the utilitarian SWFL defined above).

Finally, the additional information available under RFC makes 'prioritarian' SWFLs possible.¹⁰ Like utilitarian SWFLs, they rank-order social alternatives on the basis of welfare sums across the individuals in N , but rather than summing up welfare directly, they sum up concavely transformed welfare, giving greater marginal weight to lower levels of welfare.

Prioritarianism: For any profile $\langle W_1, W_2, \dots, W_n \rangle$ and any $x, y \in X$, xRy if and only if $W_1^r(x) + W_2^r(x) + \dots + W_n^r(x) \geq W_1^r(y) + W_2^r(y) + \dots + W_n^r(y)$, where $0 < r < 1$.

Prioritarianism requires RFC and not merely CFC because, by design, the prioritarian social ordering for any welfare profile is not invariant under changes in welfare levels (shifting).

4.3 Other topics

The present framework has been applied to several further areas. It has been generalized to variable-population-size choice problems, so as to formalize population ethics in the tradition of Parfit (1984). Here, we must rank-order social alternatives (e.g., possible worlds) in which different individuals exist. For example, the set of individuals existing under alternative x , $N(x)$, could differ from the set of individuals existing under y , $N(y)$. This raises questions such as whether a world with a smaller number of better-off individuals is better than, equally good as, or worse than a world with a larger number of worse-off individuals. (The focus here is on axiological questions about the relative goodness of such worlds, not normative questions about

¹⁰ Technically, this requires a domain restriction to positive welfare profiles.

the rightness or wrongness of bringing them about.)

Parfit (1984) and others argued that classical utilitarianism is subject to the *repugnant conclusion*: a world with a very large number of individuals whose welfare levels are barely above zero could have a larger sum-total of welfare, and therefore count as better, than a world with a smaller number of very well-off individuals.

Blackorby, Donaldson, and Bossert (e.g., 2005) have axiomatically characterized different variable-population-size welfare aggregation methods that avoid the repugnant conclusion and satisfy some other desiderata. One solution is the following:

Critical-level utilitarianism: For any profile $\langle W_1, W_2, \dots, W_n \rangle$ and any $x, y \in X$, xRy if and only if $\sum_{i \in N(x)} [W_i(x) - c] \geq \sum_{i \in N(y)} [W_i(y) - c]$, where $c \geq 0$ is some ‘critical level’ of welfare above which the quality of life counts as ‘decent/good’.

Critical-level utilitarianism avoids the repugnant conclusion when the parameter c is set sufficiently large. It requires stronger measurability of welfare than classical utilitarianism, since it generates a social ordering R that is not generally invariant under re-scaling of welfare units or shifts in welfare levels. Even the rich framework of RFC would force the critical level c to be zero, thereby collapsing critical-level utilitarianism into classical utilitarianism and making it vulnerable to the repugnant conclusion again. As Blackorby, Bossert, and Donaldson (1999, p. 420) note, ‘[s]ome information environments that are ethically adequate in fixed-population settings have ethically unattractive consequences in variable-population environments.’ Thus, in the variable-population-size case, a more significant departure from the informational framework of Arrow’s original model is needed to avoid impossibility results.

The SWFL approach has been generalized to the case in which each individual has multiple welfare functions (e.g., a k -tuple of them), capturing (i) multiple opinions about each individual’s welfare (e.g., Roberts 1995, Ooghe and Lauwers 2005) or (ii) multiple dimensions of welfare (e.g., List 2004a). In this case, we are faced not only with issues of measurability and interpersonal comparability, but also with issues of inter-opinion or inter-dimensional comparability. To obtain compelling possibility results, comparability across both individuals and dimensions/opinions is needed. A related literature addresses multidimensional inequality measurement (for an introductory review, see Weymark 2006).

Finally, in the philosophy of biology, the one-dimensional and multi-dimensional SWFL frameworks have been used (by Okasha 2009 and Bossert, Qi, and Weymark 2013) to analyse the notion of group fitness, defined as a function of individual fitness indicators.

5. Judgment aggregation

A more recent branch of social choice theory is the theory of judgment aggregation. It can be motivated by observing that votes for a most preferred alternative or orderings (or welfare functions) over multiple alternatives are not the only objects we may wish to aggregate from an individual to a collective level. Many decision-making bodies, such as legislatures, collegial courts, expert panels, and other committees, are faced with more complex ‘aggreganda’. In particular, they may have to aggregate individual

sets of judgments on multiple, logically connected propositions into collective sets of judgments.

A court may have to judge whether a defendant is liable for breach of contract on the basis of whether there was a valid contract in place and whether there was a breach. An expert panel may have to judge whether atmospheric greenhouse-gas concentrations will exceed a particular threshold by 2050, whether there is a causal chain from greater greenhouse-gas concentrations to temperature increases, and whether the temperature will increase. Legislators may have to judge whether a particular end is socially desirable, whether a proposed policy is the best means for achieving that end, and whether to pursue that policy.

These problems cannot be formalized in standard preference-aggregation models, since the aggreganda are not orderings but sets of judgments on multiple propositions. The theory of judgment aggregation represents these aggreganda in propositional logic (or another suitable logic). The field was inspired by the ‘doctrinal paradox’ in jurisprudence, with which we begin.

5.1 The ‘doctrinal paradox’ and the ‘discursive dilemma’

Kornhauser and Sager (1986) described the following problem. (A structurally similar problem was discovered by Vacca 1921 and, as Elster 2013 points out, by Poisson 1837.) A three-judge court has to make judgments on the following propositions:

- p : The defendant was contractually obliged not to do action X.
- q : The defendant did action X.
- r : The defendant is liable for breach of contract.

According to legal doctrine, the *premises* p and q are jointly necessary and sufficient for the *conclusion* r . Suppose the individual judges hold the views shown in Table 5.

Table 5: An example of the ‘doctrinal paradox’

	p (obligation)	q (action)	r (liability)
Judge 1	True	True	True
Judge 2	False	True	False
Judge 3	True	False	False
Majority	True	True	False

Although each individual judge respects the relevant legal doctrine, there is a majority for p , a majority for q , and yet a majority against r – in breach of legal doctrine. The court faces a dilemma: it can *either* go with the majority judgments on the premises (p and q) and reach a ‘liable’ verdict by logical inference (the *issue-by-issue* or *premise-based approach*); *or* go with the majority judgment on the conclusion (r) and reach a ‘not liable’ verdict, ignoring the majority judgments on the premises (the *case-by-case* or *conclusion-based approach*). Kornhauser and Sager’s ‘doctrinal paradox’ consists in the fact that these two approaches may lead to opposite outcomes.

We can learn another lesson from this example. Relative to the legal doctrine, the majority judgments are logically inconsistent. Formally expressed, the set of majority-accepted propositions, $\{p, q, \text{not } r\}$, is inconsistent relative to the constraint r if and

only if (p and q). This observation was the starting point of the more recent, logic-based literature on judgment aggregation (beginning with List and Pettit 2002).

The possibility of inconsistent majority judgments is not tied to the presence of a legal doctrine or other explicit side constraint (as pointed out by Pettit 2001, who called this phenomenon the ‘discursive dilemma’). Suppose, for example, an expert panel has to make judgments on three propositions (and their negations):

p : Atmospheric CO₂ will exceed 600ppm by 2050.

if p then q: If atmospheric CO₂ exceeds this level by 2050, there will be a temperature increase of more than 3.5° by 2010.

q : There will be a temperature increase of more than 3.5° by 2010.

If individual judgments are as shown in Table 6, the majority judgments are inconsistent: despite individually consistent judgments, the set of majority-accepted propositions, $\{p, \textit{if p then q}, \textit{not q}\}$, is logically inconsistent.

Table 6: A majoritarian inconsistency

	p	<i>if p then q</i>	q
Expert 1	True	True	True
Expert 2	False	True	False
Expert 3	True	False	False
Majority	True	True	False

5.2 The conditions for inconsistent majority judgments

The patterns of judgments in Tables 5 and 6 are structurally equivalent to the pattern of preferences leading to Condorcet’s paradox, when we reinterpret those preferences as judgments on propositions of the form ‘ x is preferable to y ’, ‘ y is preferable to z ’, and so on, as shown in Table 7 (List and Pettit 2004; an earlier interpretation of preferences along these lines can be found in Guilbaud [1952] 1966). Here, the set of majority-accepted propositions is inconsistent relative to the constraint of transitivity.

Table 7: Condorcet’s paradox, propositionally reinterpreted

	‘ x is preferable to y ’	‘ y is preferable to z ’	‘ x is preferable to z ’
Individual 1 (prefers x to y to z)	True	True	True
Individual 2 (prefers y to z to x)	False	True	False
Individual 3 (prefers z to x to y)	True	False	False
Majority (prefers x to y to z to x , a ‘cycle’)	True	True	False

A general combinatorial result subsumes all these phenomena. Call a set of propositions *minimally inconsistent* if it is a logically inconsistent set, but all its proper subsets are consistent.

Proposition (Dietrich and List 2007a, Nehring and Puppe 2007): Propositionwise majority voting may generate inconsistent collective judgments if and only if the set

of propositions (and their negations) on which judgments are to be made has a minimally inconsistent subset of three or more propositions.

In the examples of Tables 6, 5, and 7, the relevant minimally inconsistent sets of size (at least) three are: $\{p, \text{if } p \text{ then } q, \text{not } q\}$; $\{p, q, \text{not } r\}$, which is minimally inconsistent relative to the side constraint $r \text{ if and only if } (p \text{ and } q)$; and $\{‘x \text{ is preferable to } y’, ‘y \text{ is preferable to } z’, ‘z \text{ is preferable to } x’\}$, which is minimally inconsistent relative to a transitivity constraint on preferability.

5.3 A general model and a simple impossibility result

The basic model of judgment aggregation can be defined as follows (List and Pettit 2002). Let $N = \{1, 2, \dots, n\}$ be a set of individuals ($n \geq 2$). The propositions on which judgments are to be made are represented by sentences from propositional logic (or some other, expressively richer logic, such as a predicate, modal, or conditional logic; see Dietrich 2007). We define the *agenda*, X , as a finite set of propositions, closed under single negation.¹¹ For example, X could be $\{p, \neg p, p \rightarrow q, \neg(p \rightarrow q), q, \neg q\}$, as in the expert-panel case.

Each individual $i \in N$ has a *judgment set* J_i , defined as a subset $J_i \subseteq X$ and interpreted as the set of propositions that individual i accepts. A judgment set is *consistent* if it is a logically consistent set of propositions¹² and *complete* (relative to X) if it contains a member of every proposition-negation pair $p, \neg p \in X$.

A combination of judgment sets across the individuals, $\langle J_1, J_2, \dots, J_n \rangle$, is called a *profile*. A *judgment aggregation rule*, F , is a function that assigns to each profile $\langle J_1, J_2, \dots, J_n \rangle$ (in some domain of admissible profiles) a collective judgment set $J = F(J_1, J_2, \dots, J_n) \subseteq X$, interpreted as the set of propositions accepted by the group as a whole. As before, when F is clear from the context, we write J for the collective judgment set corresponding to $\langle J_1, J_2, \dots, J_n \rangle$. Again, for generality, we build no rationality requirement on J (such as consistency or completeness) into the definition of a judgment aggregation rule.

The simplest example of a judgment aggregation rule is *propositionwise majority voting*. Here, for any profile $\langle J_1, J_2, \dots, J_n \rangle$, $J = \{p \in X : |\{i \in N : p \in J_i\}| > n/2\}$. As we have seen, this may produce inconsistent collective judgments.

Consider the following conditions on an aggregation rule:

Universal domain: The domain of F is the set of all logically possible profiles of consistent and complete individual judgment sets.

Collective rationality: For any profile $\langle J_1, J_2, \dots, J_n \rangle$ in the domain of F , the collective judgment set J is consistent and complete.

¹¹ Formally, $X = \{p, \neg p : p \in X^+\}$, where X^+ is a set of un-negated propositions. To avoid technicalities, we assume that X contains no contradictory or tautological propositions.

¹² In principle, consistency can be defined relative to some side constraint such as the legal doctrine in the ‘doctrinal paradox’ example.

Anonymity: For any two profiles $\langle J_1, J_2, \dots, J_n \rangle$ and $\langle J_1^*, J_2^*, \dots, J_n^* \rangle$ that are permutations of each other, $F(J_1, J_2, \dots, J_n) = F(J_1^*, J_2^*, \dots, J_n^*)$.

Systematicity: For any two profiles $\langle J_1, J_2, \dots, J_n \rangle$ and $\langle J_1^*, J_2^*, \dots, J_n^* \rangle$ in the domain of F and any $p, q \in X$, if for all $i \in N$ $p \in J_i$ if and only if $q \in J_i^*$, then $p \in J$ if and only if $q \in J^*$.

The first three conditions are analogous to universal domain, ordering, and anonymity in preference aggregation. The last is the counterpart of independence of irrelevant alternatives, though stronger: it requires that (i) the collective judgment on any proposition $p \in X$ (of which a binary ranking proposition such as ‘ x is preferable to y ’ is a special case) depend only on individual propositions on p (the *independence* part), and (ii) the pattern of dependence between individual and collective judgments be the same across all propositions in X (the *neutrality* part). Formally, *independence* is the special case with quantification restricted to $p = q$. Propositionwise majority voting satisfies all these conditions, except the consistency part of collective rationality.

Theorem (List and Pettit 2002): If $\{p, q, p \wedge q\} \subseteq X$ (where p and q are mutually independent propositions and ‘ \wedge ’ can also be replaced by ‘ \vee ’ or ‘ \rightarrow ’), there exists no judgment aggregation rule satisfying universal domain, collective rationality, anonymity, and systematicity.

Like other impossibility theorems, this result is best interpreted as describing the trade-offs between different conditions on an aggregation rule. The result has been generalized and strengthened in various ways, beginning with Pauly and van Hees’s (2006) proof that the impossibility persists if anonymity is weakened to non-dictatorship (for other generalizations, see Dietrich 2006 and Mongin 2008).

5.4 More general impossibility and possibility theorems

As we have seen, in preference aggregation, the ‘boundary’ between possibility and impossibility results is easy to draw: when there are only two decision alternatives, all of the desiderata on a preference aggregation rule reviewed above can be satisfied (and majority rule does the job); when there are three or more alternatives, there are impossibility results. In judgment aggregation, by contrast, the picture is more complicated. What matters is not the number of propositions in X but the nature of the logical interconnections between them.

Impossibility results in judgment aggregation have the following generic form: for a given class of agendas, the aggregation rules satisfying a particular set of conditions (usually, a domain condition, a rationality condition, and some responsiveness conditions) are non-existent or degenerate (e.g., dictatorial). Different kinds of agendas trigger different instances of this scheme, with stronger or weaker conditions imposed on the aggregation rule depending on the properties of those agendas (for a more detailed review, see List 2012). The significance of combinatorial properties of the agenda was first discovered by Nehring and Puppe (2002) in a mathematically related but interpretationally distinct framework (strategy-proof social choice over so-called property spaces). Three kinds of agenda stand out:

A non-simple agenda: X has a minimally inconsistent subset of three or more

propositions.

A pair-negatable agenda: X has a minimally inconsistent subset Y that can be rendered consistent by negating a pair of propositions in it. (Equivalently, X is not isomorphic to a set of propositions whose only connectives are \neg and \leftrightarrow ; see Dokow and Holzman 2010a.)

A path-connected agenda (or *totally blocked*, in Nehring and Puppe 2002): For any $p, q \in X$, there is a sequence $p_1, p_2, \dots, p_k \in X$ with $p_1 = p$ and $p_k = q$ such that p_1 conditionally entails p_2 , p_2 conditionally entails p_3 , ..., and p_{k-1} conditionally entails p_k . (Here, p_i conditionally entails p_j if $p_i \cup Y$ entails p_j for some $Y \subseteq X$ consistent with each of p_i and $\neg p_j$.)

Some agendas have two or more of these properties. The agendas in our ‘doctrinal paradox’ and ‘discursive dilemma’ examples are both non-simple and pair-negatable. The *preference agenda*, $X = \{‘x$ is preferable to $y’$, ‘ y is preferable to $x’$, ‘ x is preferable to $z’$, ‘ z is preferable to $x’$, ...}, is non-simple, pair-negatable, and path-connected (assuming preferability is transitive and complete). The following result holds:

Theorem (Dietrich and List 2007b, Dokow and Holzman 2010a, building on Nehring and Puppe 2002): If X is non-simple, pair-negatable, and path-connected, there exists no judgment aggregation rule satisfying universal domain, collective rationality, independence, unanimity preservation (requiring that, for any unanimous profile $\langle J, J, \dots, J \rangle$, $F(J, J, \dots, J) = J$), and non-dictatorship.¹³

Applied to the preference agenda, this result yields Arrow’s theorem (for strict preference orderings) as a corollary (predecessors of this result can be found in List and Pettit 2004 and Nehring 2003).¹⁴ Thus Arrovian preference aggregation can be reinterpreted as a special case of judgment aggregation.

The literature contains several variants of this theorem. One variant drops the agenda property of path-connectedness and strengthens independence to systematicity. A second variant drops the agenda property of pair-negatability and imposes a monotonicity condition on the aggregation rule (requiring that additional support never hurt an accepted proposition) (Nehring and Puppe 2010; the latter result was first proved in the above-mentioned mathematically related framework by Nehring and Puppe 2002). A final variant drops both path-connectedness and pair-negatability while imposing both systematicity and monotonicity (ibid.).

In each case, the agenda properties are not only sufficient but also necessary for the result (Nehring and Puppe 2002, 2010; Dokow and Holzman 2010a). Note also that path-connectedness implies non-simplicity. Therefore, non-simplicity need not be listed among the theorem’s conditions, though it is needed in the variants dropping path-connectedness.

¹³ See also the remark on the relationship between path-connectedness and non-simplicity at the end of this subsection.

¹⁴ An earlier mathematically related, though interpretationally distinct contribution is Wilson’s work on abstract aggregation (1975).

5.5 Non-dictatorial judgment aggregation rules

5.5.1 Relaxing universal domain

As in preference aggregation, one way to avoid the present impossibility results is to relax universal domain. If the domain of admissible profiles of individual judgment sets is restricted to those satisfying specific ‘cohesion’ conditions, propositionwise majority voting produces consistent collective judgments.

The simplest cohesion condition is *unidimensional alignment* (List 2003c). A profile $\langle J_1, J_2, \dots, J_n \rangle$ is *unidimensionally aligned* if the individuals in N can be ordered from left to right (e.g., on some cognitive or ideological dimension) such that, for every proposition $p \in X$, 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 p (i.e., those with $p \notin J_i$), as illustrated in Table 8. For any such profile, the majority judgments are consistent: the judgment set of the median individual relative to the left-right ordering will prevail (where n is odd). This judgment set will inherit its consistency from the median individual, assuming individual judgments are consistent. By implication, on unidimensionally aligned domains, propositionwise majority voting will satisfy the rest of the conditions on judgment aggregation rules reviewed above.

Table 8: Unidimensional alignment

	Individual 1	Individual 2	Individual 3	Individual 4	Individual 5
p	True	True	False	False	False
q	True	True	True	True	False
r	False	False	False	True	True
$p \wedge q \wedge r$	False	False	False	False	False

In analogy with the case of single-peakedness in preference aggregation, several less restrictive conditions already suffice for consistent majority judgments. One such condition (introduced in Dietrich and List 2010a, where a survey is provided) generalizes Sen’s triple-wise value-restriction. A profile $\langle J_1, J_2, \dots, J_n \rangle$ is *value-restricted* if every minimally inconsistent subset $Y \subseteq X$ has a pair of elements p, q such that no individual $i \in N$ has $\{p, q\} \subseteq J_i$. Value-restriction prevents any minimally inconsistent subset of X from becoming majority-accepted, and hence ensures consistent majority judgments. Applied to the preference agenda, value-restriction reduces to Sen’s equally named condition.

5.5.2 Relaxing collective rationality

While the requirement that collective judgments be consistent is widely accepted, the requirement that collective judgments be complete (in X) is more contentious. In support of completeness, one might say that a given proposition would not be included in X unless it is supposed to be collectively adjudicated. Against completeness, one might say that there are circumstances in which the level of disagreement on a particular proposition (or set of propositions) is so great that forming a collective view on it is undesirable or counterproductive. Several papers offer possibility or impossibility results on completeness relaxations (e.g., List and Pettit 2002; Gärdenfors 2006; Dietrich and List 2007a, 2008; Dokow and Holzman

2010b).

Judgment aggregation rules violating collective completeness while satisfying (all or most of) the other conditions introduced above include: *unanimity rule*, where, for any profile $\langle J_1, J_2, \dots, J_n \rangle$, $J = \{p \in X : p \in J_i \text{ for all } i \in N\}$; *supermajority rules*, where, for any profile $\langle J_1, J_2, \dots, J_n \rangle$, $J = \{p \in X : |\{i \in N : p \in J_i\}| > qn\}$ for a suitable acceptance quota $q \in [0.5, 1)$; and *conclusion-based rules*, where a subset $Y \subseteq X$ of logically independent propositions (and their negations) is designated as a set of conclusions and $J = \{p \in Y : |\{i \in N : p \in J_i\}| > n/2\}$. In the multi-member court example of Table 5, the set of conclusions is simply $Y = \{r, \neg r\}$.

Given consistent individual judgment sets, unanimity rule guarantees consistent collective judgment sets, because the intersection of several consistent sets of propositions is always consistent. Supermajority rules guarantee consistent collective judgment sets too, provided the quota q is chosen to be at least $(k-1)/k$, where k is the size of the largest minimally inconsistent subset of X . The reason is combinatorial: any k distinct supermajorities of the relevant size will always have at least one individual in common. So, for any minimally inconsistent set of propositions (which is at most of size k) to be majority-accepted, at least one individual would have to accept all the propositions in the set, contradicting this individual's consistency (Dietrich and List 2007a, List and Pettit 2002). Conclusion-based rules, finally, produce consistent collective judgment sets by construction, but always leave non-conclusions undecided.

Gärdenfors (2006) and more generally Dietrich and List (2008) and Dokow and Holzman (2010b) have shown that if – while relaxing completeness – we require collective judgment sets to be *deductively closed* (i.e., for any $p \in X$ entailed by J , it must be that $p \in J$), we face an impossibility result again. For the same agendas that lead to the impossibility result reviewed in Section 5.4, there exists no judgment aggregation rule satisfying universal domain, collective consistency and deductive closure, independence, unanimity preservation, and non-oligarchy. An aggregation rule is called *oligarchic* if there is an antecedently fixed subset $M \subseteq N$ (the ‘oligarchs’) such that, for any profile $\langle J_1, J_2, \dots, J_n \rangle$, $J = \{p \in X : p \in J_i \text{ for all } i \in M\}$. Unanimity rule and dictatorships are special cases with $M = N$ and $M = \{i\}$ for some $i \in N$, respectively.

The downside of oligarchic aggregation rules is that they either lapse into dictatorship or lead to stalemate, with the slightest disagreements between oligarchs resulting in indecision (since every oligarch has veto power on every proposition).

5.5.3 Relaxing systematicity/independence

A variety of judgment aggregation rules become possible when we relax systematicity/independence. Recall that systematicity combines an independence and a neutrality requirement. Relaxing only neutrality does not get us very far, since for many agendas there are impossibility results with independence alone, as illustrated in Section 5.4.

One much-discussed class of aggregation rules violating independence is given by the

premise-based rules. Here, a subset $Y \subseteq X$ of logically independent propositions (and their negations) is designated as a set of *premises*, as in the court example. For any profile $\langle J_1, J_2, \dots, J_n \rangle$, $J = \{p \in X : J_Y \text{ entails } p\}$ where J_Y are the majority-accepted propositions among the premises, formally $\{p \in Y : |\{i \in N : p \in J_i\}| > n/2\}$. Informally, majority votes are taken on the premises, and the collective judgments on all other propositions are determined by logical implication. If the premises constitute a logical basis for the entire agenda, a premise-based rule guarantees consistent and (absent ties) complete collective judgment sets. (The present definition follows List and Pettit 2002; for generalizations, see Dietrich and Mongin 2010. The procedural and epistemic properties of premise-based rules are discussed in Pettit 2001, Chapman 2002, Bovens and Rabinowicz 2006, Dietrich 2006, and List 2006.)

A generalization is given by the *sequential priority rules* (List 2004b, Dietrich and List 2007a). Here, for each profile $\langle J_1, J_2, \dots, J_n \rangle$, the propositions in X are collectively adjudicated in a fixed order of priority, for instance, a temporal or epistemic one. The collective judgment on each proposition $p \in X$ is made as follows. If the majority judgment on p is consistent with collective judgments on prior propositions, this majority judgment prevails; otherwise the collective judgment on p is determined by the implications of prior judgments. By construction, this guarantees consistent and (absent ties) complete collective judgments. However, it is *path-dependent*: the order in which propositions are considered may affect the outcome, specifically when the underlying majority judgments are inconsistent. For example, when this aggregation rule is applied to the profiles in Tables 5, 6, and 7 (but not 8), the collective judgments depend on the order in which the propositions are considered. Thus sequential priority rules are vulnerable to *agenda manipulation*. Similar phenomena occur in sequential pairwise majority voting in preference aggregation (e.g., Riker 1982).

Another prominent class of aggregation rules violating independence is given by the *distance-based rules* (Pigozzi 2006, building on Konieczny and Pino Pérez 2002; see also Miller and Osherson 2009). A distance-based rule is defined in terms of some distance metric between judgment sets, for example the *Hamming distance*, where, for any two judgment sets $J, J' \subseteq X$, $d(J, J') = |\{p \in X : \text{not } [p \in J \Leftrightarrow p \in J']\}|$. Each profile $\langle J_1, J_2, \dots, J_n \rangle$ is mapped to a consistent and complete judgment set J that minimizes the sum-total distance from each of the J_i s. Distance-based rules can be interpreted as capturing the idea of identifying compromise judgments. Unlike premise-based or sequential priority rules, they do not require a distinction between premises and conclusions or any other order of priority among the propositions.

As in preference aggregation, the cost of relaxing independence is the loss of strategy-proofness. The conjunction of independence and monotonicity is necessary and sufficient for the non-manipulability of a judgment aggregation rule by strategic voting (Dietrich and List 2007c; for related results, see Nehring and Puppe 2002). Thus we cannot generally achieve strategy-proofness without relaxing either universal domain, or collective rationality, or unanimity preservation, or non-dictatorship. In practice, we must therefore look for ways of rendering opportunities for strategic manipulation less of a threat.

6. Other topics

As should be evident, social choice theory is a vast field. Areas not covered in this entry, or mentioned only in passing, include: theories of fair division (how to divide one or several divisible or indivisible goods, such as cakes or houses, between several claimants; e.g., Brams and Taylor 1996 and Moulin 2004); behavioural social choice theory (analyzing empirical evidence of voting behaviour under various aggregation rules; e.g., Regenwetter et al. 2006, List, Luskin, Fishkin, and McLean 2013); empirical social choice theory (analysing surveys and experiments on people's intuitions about distributive justice; e.g., Gaertner and Schokkaert 2012); computational social choice theory (analysing computational properties of aggregation rules, including their computational complexity; e.g., Bartholdi, Tovey, and Trick 1989, Brandt, Conitzer, and Endriss 2013); theories of probability aggregation (studying the aggregation of probability or credence functions; e.g., Lehrer and Wagner 1981, McConway 1981, Genest and Zidek 1986, Mongin 1995, Dietrich and List 2007d); theories of general attitude aggregation (generalizing two-valued judgment aggregation, probability/credence aggregation, and preference aggregation; e.g., Dietrich and List 2010b, Dokow and Holzman 2010c); the study of collective decision-making in non-human animals (studying group decisions in a variety of animal species from social insects to primates, as surveyed in Conradt and List 2009 and the special issue it introduces); and applications to social epistemology (the analysis of group doxastic states and their relationship to individual doxastic states; e.g., Goldman 2004, 2010).

Acknowledgments

I am grateful to the Editors, their Reviewers, and Iain McLean for comments.

References

- Arrow, K. (1951/1963). *Social Choice and Individual Values*. New York (Wiley).
- Austen-Smith, D., and J. S. Banks (1996). "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review* 90: 34-45.
- Bartholdi, J. J., C. A. Tovey, and M. A. Trick (1989). "The computational difficulty of manipulating an election." *Social Choice and Welfare* 6: 227-241.
- Ben-Yashar, R., and S. Nitzan (1997). "The optimal decision rule for fixed-size committees in dichotomous choice situations: the general result." *International Economic Review* 38: 175-186.
- Berend, D., and J. Paroush (1998). "When is Condorcet's Jury Theorem valid?" *Social Choice and Welfare* 15: 481-488.
- Berend, D., and L. Sapir. (2007). "Monotonicity in Condorcet's Jury Theorem with dependent voters." *Social Choice and Welfare* 28: 507-528.
- Black, D. (1948). "On the Rationale of Group Decision-Making." *Journal of Political Economy* 56: 23-34.
- Blackorby, C., W. Bossert, and D. Donaldson (1999). "Information Invariance in Variable-Population Social-Choice Problems." *International Economic Review* 40: 403-422.
- Blackorby, C., W. Bossert, and D. Donaldson (2005). *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge (Cambridge University Press).
- Blau, J. H. (1975). "Liberal Values and Independence." *Review of Economic Studies* 42: 395-401.

- Boland, P. J. (1989). "Majority systems and the Condorcet jury theorem." *Statistician* 38: 181-189.
- Bossert, W., and J. A. Weymark (1996). "Utility in social choice." *Handbook of Utility Theory, Volume 2*. S. Barberà, P. J. Hammond, and C. Seidel (eds.). Boston (Kluwer).
- Bossert, W., C. X. Qi, and J. A. Weymark (2013). "Extensive social choice and the measurement of group fitness in biological hierarchies." *Biology and Philosophy* 28: 75-98.
- Bovens, L., and W. Rabinowicz (2006). "Democratic Answers to Complex Questions: An Epistemic Perspective." *Synthese* 150: 131-153.
- Brams, S. J., and A. D. Taylor (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge (Cambridge University Press).
- Brandt, F., V. Conitzer, and U. Endriss (2013). "Computational Social Choice." *Multiagent Systems*. G. Weiss (ed.). Cambridge, MA (MIT Press), pp. 213-283.
- Brighthouse, H., and M. Fleurbaey (2010). "Democracy and Proportionality." *Journal of Political Philosophy* 18: 137-155.
- Chapman, B. (2002). "Rational Aggregation." *Politics, Philosophy and Economics* 1: 337-354.
- Condorcet, Nicolas de (1785). *Essay sur l'Application de l'Analyse à la Probabilité des Décisions Rendue à la Pluralité des Voix*. Paris.
- Conradt, L., and C. List (2009). "Group decisions in humans and animals: a survey." *Philosophical Transactions of the Royal Society B* 364: 719-742.
- Craven, J. (1982). "Liberalism and Individual Preferences." *Theory and Decision* 14: 351-360.
- Deschamps, R., and L. Gevers (1978). "Leximin and utilitarian rules: A joint characterization." *Journal of Economic Theory* 17: 143-163.
- Dietrich, F. (2006). "Judgment Aggregation: (Im)Possibility Theorems." *Journal of Economic Theory* 126: 286-298.
- Dietrich, F. (2007). "A generalised model of judgment aggregation." *Social Choice and Welfare* 28: 529-565.
- Dietrich, F. (2008). "The premises of Condorcet's jury theorem are not simultaneously justified." *Episteme* 5: 56-73.
- Dietrich, F., and C. List (2004). "A Model of Jury Decisions Where All Jurors Have the Same Evidence." *Synthese* 142: 175-202.
- Dietrich, F., and C. List (2007a). "Judgment aggregation by quota rules: majority voting generalized." *Journal of Theoretical Politics* 19: 391-424.
- Dietrich, F., and C. List (2007b). "Arrow's theorem in judgment aggregation." *Social Choice and Welfare* 29: 19-33.
- Dietrich, F., and C. List (2007c). "Strategy-proof judgment aggregation." *Economics and Philosophy* 23: 269-300.
- Dietrich, F., and C. List (2007d). "Opinion pooling on general agendas." Online paper, available at: <http://personal.lse.ac.uk/list/pdf-files/opinionpooling.pdf>
- Dietrich, F., and C. List (2008). "Judgment aggregation without full rationality." *Social Choice and Welfare* 31: 15-39.
- Dietrich, F., and C. List (2010a). "Majority voting on restricted domains." *Journal of Economic Theory* 145: 512-543.
- Dietrich, F., and C. List (2010b). "The aggregation of propositional attitudes: towards a general theory." *Oxford Studies in Epistemology* 3: 215-234.
- Dietrich, F., and P. Mongin (2010). "The premise-based approach to judgment

- aggregation.” *Journal of Economic Theory* 145: 562-582.
- Dietrich, F., and K. Spiekermann (2013). “Epistemic Democracy with Defensible Premises.” *Economics and Philosophy* 29 (in press).
- Dokow, E., and R. Holzman (2010a). “Aggregation of binary evaluations.” *Journal of Economic Theory* 145: 495-511.
- Dokow, E., and R. Holzman (2010b). “Aggregation of binary evaluations with abstentions.” *Journal of Economic Theory* 145: 544-561.
- Dokow, E., and R. Holzman (2010c). “Aggregation of non-binary evaluations.” *Advances in Applied Mathematics* 45: 487-504.
- Dowding, K., and M. van Hees (2003). “The Construction of Rights.” *American Political Science Review* 97: 281-293.
- Dowding, K., and M. van Hees (2008). “In Praise of Manipulation.” *British Journal of Political Science* 38: 1-15.
- Dryzek, J., and C. List (2003). “Social Choice Theory and Deliberative Democracy: A Reconciliation.” *British Journal of Political Science* 33: 1-28.
- Elster, J. (2013). “Excessive Ambitions (II).” *Capitalism and Society* 8, Issue 1, Article 1.
- Estlund, D. (1994). “Opinion Leaders, Independence, and Condorcet’s Jury Theorem.” *Theory and Decision* 36: 131-162.
- Feddersen, T. J., and W. Pesendorfer (1998). “Convicting the Innocent.” *American Political Science Review* 92: 23-35.
- Gaertner, W. (2001). *Domain Conditions in Social Choice Theory*. Cambridge (Cambridge University Press).
- Gaertner, W. (2005). “De jure naturae et gentium: Samuel von Pufendorf’s contribution to social choice theory and economics.” *Social Choice and Welfare* 25: 231-241.
- Gaertner, W., P. K. Pattanaik, and K. Suzumura (1992). “Individual Rights Revisited.” *Economica* 59: 161-177.
- Gaertner, W., and E. Schokkaert (2012). *Empirical Social Choice: Questionnaire-Experimental Studies on Distributive Justice*. Cambridge (Cambridge University Press).
- Gärdenfors, P. (2006). “An Arrow-like theorem for voting with logical consequences.” *Economics and Philosophy* 22: 181-190.
- Gehrlein, W. V. (1983). “Condorcet’s Paradox.” *Theory and Decision* 15: 161-197.
- Genest, C., and J. V. Zidek (1986). “Combining Probability Distributions: A Critique and Annotated Bibliography.” *Statistical Science* 1: 113-135.
- Gibbard, A. (1969). “Social Choice and the Arrow Conditions.” Unpublished manuscript, available at: <http://www-personal.umich.edu/~gibbard/Gibbard1968-Social-Choice-Arrow-Conditions.pdf>
- Gibbard, A. (1973). “Manipulation of voting schemes: a general result.” *Econometrica* 41: 587-601.
- Gigliotti, G. A. (1986). “Comment on Craven.” *Theory and Decision* 21: 89-95.
- Gilboa, I., D. Samet, and D. Schmeidler (2004). “Utilitarian Aggregation of Beliefs and Tastes.” *Journal of Political Economy* 112: 932-938.
- Goldman, A. (2004). “Group Knowledge versus Group Rationality: Two Approaches to Social Epistemology.” *Episteme, A Journal of Social Epistemology* 1: 11-22.

- Goldman, A. (2010). "Why Social Epistemology Is Real Epistemology." *Social Epistemology*. A. Haddock, A. Millar, and D. Pritchard (eds.). Oxford (Oxford University Press).
- Goodin, R. E., and C. List (2006). "A Conditional Defense of Plurality Rule: Generalizing May's Theorem in a Restricted Informational Environment." *American Journal of Political Science* 50: 940-949.
- Grofman, B., G. Owen, and S. L. Feld (1983). "Thirteen theorems in search of the truth." *Theory and Decision* 15: 261-278.
- Guilbaud, G. T. ([1952] 1966). "Theories of the General Interest, and the Logical Problem of Aggregation." *Readings in Mathematical Social Science*. P. F. Lazarsfeld and N. W. Henry (eds.). Cambridge/MA (MIT Press), pp. 262-307.
- Harbour, D., and C. List (2000). "Optimality Theory and the problem of constraint aggregation." *MIT Working Papers in Linguistics and Philosophy 1: The Linguistics/Philosophy Interface*. R. Bhatt, P. Hawley, M. Hackl and I. Maitra (eds.). Cambridge, MA (MITWIPL), pp. 175-213.
- Harrison, G. W., and T. McDaniel (2008). "Voting Games and Computational Complexity." *Oxford Economic Papers* 60: 546-565.
- Hausman, D. (1995). "The Impossibility of Interpersonal Utility Comparisons." *Mind* 104: 473-490.
- Hurley, S. (1985). "Supervenience and the Possibility of Coherence." *Mind* 94: 501-525.
- Inada, K.-I. (1964). "A Note on the Simple Majority Decision Rule." *Econometrica* 32: 525-531.
- Kanazawa, S. (1998). "A brief note on a further refinement of the Condorcet Jury Theorem for heterogeneous groups." *Mathematical Social Sciences* 35: 69-73.
- Knight, J., and J. Johnson (1994). "Aggregation and Deliberation: On the Possibility of Democratic Legitimacy." *Political Theory* 22: 277-296.
- Konieczny, S., and R. Pino Pérez (2002). "Merging Information Under Constraints: A Logical Framework." *Journal of Logic and Computation* 12: 773-808.
- Kornhauser, L. A., and L. G. Sager (1986). "Unpacking the Court." *Yale Law Journal* 96: 82-117.
- Kroedel, T., and F. Huber (2012). "Counterfactual Dependence and Arrow." *Nous* (in press).
- Ladha, K. (1992). "The Condorcet Jury Theorem, Free Speech and Correlated Votes." *American Journal of Political Science* 36: 617-634.
- Lehrer, K., and C. Wagner (1981). *Rational Consensus in Science and Society*. Dordrecht/Boston (Reidel).
- List, C. (2001). "A Note on Introducing a 'Zero-Line' of Welfare as an Escape-Route from Arrow's Theorem." *Pacific Economic Review* 6, special section in honour of Amartya Sen, 223-238.
- List, C. (2003a). "The epistemology of special majority voting." Working paper, London School of Economics, available at: <http://personal.lse.ac.uk/list/PDF-files/proportion3107.pdf>
- List, C. (2003b). "Are Interpersonal Comparisons of Utility Indeterminate?" *Erkenntnis* 58: 229-260.
- List, C. (2003c). "A Possibility Theorem on Aggregation over Multiple Interconnected Propositions." *Mathematical Social Sciences* 45: 1-13 (with Corrigendum in *Mathematical Social Sciences* 52: 109-110).
- List, C. (2004a). "Multidimensional Welfare Aggregation." *Public Choice* 119: 119-142.

- List, C. (2004b). "A Model of Path-Dependence in Decisions over Multiple Propositions." *American Political Science Review* 98: 495-513.
- List, C. (2006). "The Discursive Dilemma and Public Reason." *Ethics* 116: 362-402.
- List, C. (2011). "The Logical Space of Democracy." *Philosophy and Public Affairs* 39: 262-297.
- List, C. (2012). "The Theory of Judgment Aggregation: An Introductory Review." *Synthese* 187: 179-207.
- List, C., and R. E. Goodin (2001). "Epistemic Democracy: Generalizing the Condorcet Jury Theorem." *Journal of Political Philosophy* 9: 277-306.
- List, C., R. C. Luskin, J. S. Fishkin, and I. McLean (2013). "Deliberation, Single-Peakedness, and the Possibility of Meaningful Democracy: Evidence from Deliberative Polls." *Journal of Politics* 75: 80-95.
- List, C., and P. Pettit (2002). "Aggregating Sets of Judgments: An Impossibility Result." *Economics and Philosophy* 18(1): 89-110.
- List, C., and P. Pettit (2004). "Aggregating Sets of Judgments: Two Impossibility Results Compared." *Synthese* 140: 207-235.
- May, K. O. (1952). "A set of independent, necessary and sufficient conditions for simple majority decision." *Econometrica* 20: 680-684.
- May, K. O. (1954). "Intransitivity, Utility, and the Aggregation of Preference Patterns." *Econometrica* 22: 1-13.
- McConway, K. J. (1981). "Marginalization and Linear Opinion Pools." *Journal of the American Statistical Association* 76: 410-414.
- McLean, I. (1990). "The Borda and Condorcet principles: Three medieval applications." *Social Choice and Welfare* 7: 99-108.
- McLean, I., and F. Hewitt (eds.) (1994). *Condorcet: Foundations of Social Choice and Political Theory*. Cheltenham (Edward Elgar Publishing).
- McLean, I. S., A. McMillan, and B. L. Monroe (1995). "Duncan Black and Lewis Carroll." *Journal of Theoretical Politics* 7: 107-123.
- McLean, I. S., A. McMillan, and B. L. Monroe (eds.) (1996). *A Mathematical Approach to Proportional Representation: Duncan Black on Lewis Carroll*. Dordrecht (Kluwer).
- McLean, I., and A. B. Urken (eds.) (1995). *Classics of Social Choice*. Ann Arbor (University of Michigan Press).
- Miller, D. (1992). "Deliberative Democracy and Social Choice." *Political Studies* 40 (special issue): 54-67.
- Miller, M. K., and D. Osherson (2009). "Methods for distance-based judgment aggregation." *Social Choice and Welfare* 32: 575-601.
- Monjardet, B. (2005). "Social choice theory and the 'Centre de Mathématique Sociale': some historical notes." *Social Choice and Welfare* 25: 433-456.
- Mongin, P. (1995). "Consistent Bayesian aggregation." *Journal of Economic Theory* 66: 313-351.
- Mongin, P. (1997). "Spurious Unanimity and the Pareto Principle." Paper presented at the Conference on Utilitarianism, New Orleans, March 1997, available at: [http://eprints.lse.ac.uk/23942/1/Spurious Unanimity and the Pareto Principle \(LSERO\).pdf](http://eprints.lse.ac.uk/23942/1/Spurious_Unanimity_and_the_Pareto_Principle_(LSERO).pdf)
- Mongin, P. (2008). "Factoring Out the Impossibility of Logical Aggregation." *Journal of Economic Theory* 141: 100-113.
- Morreau, M. (2010). "It simply does not add up: Trouble with overall similarity." *Journal of Philosophy* 107: 469-490.

- Morreau, M. (forthcoming). "Theory Choice and Social Choice: Kuhn Vindicated." *Mind*.
- Moulin, H. (1980). "On Strategy-Proofness and Single Peakedness." *Public Choice* 35: 437-455.
- Moulin, H. (2004). *Fair Division And Collective Welfare*. Cambridge, MA (MIT Press).
- Mueller, D. C. (2003). *Public Choice III*. Cambridge (Cambridge University Press).
- Nehring, K. (2003). "Arrow's theorem as a corollary." *Economics Letters* 80: 379-382.
- Nehring, K., and C. Puppe (2002). "Strategyproof Social Choice on Single-Peaked Domains: Possibility, Impossibility and the Space Between." Unpublished manuscript, University of California at Davis.
- Nehring, K., and C. Puppe (2007). "The structure of strategy-proof social choice – Part I: General characterization and possibility results on median spaces." *Journal of Economic Theory* 135: 269-305.
- Nehring, K., and C. Puppe (2010). "Abstract Arrovian Aggregation." *Journal of Economic Theory* 145: 467-494.
- Okasha, S. (2009). "Individuals, groups, fitness and utility: multi-level selection meets social choice theory." *Biology and Philosophy* 24: 561-584.
- Okasha, S. (2011). "Theory Choice and Social Choice: Kuhn versus Arrow." *Mind* 120: 83-115.
- Ooghe, E., and L. Lauwers (2005). "Non-dictatorial extensive social choice." *Economic Theory* 25: 721-743.
- Parfit, D. (1984). *Reasons and Persons*. Oxford (Oxford University Press).
- Pauly, M., and M. van Hees (2006). "Logical Constraints on Judgment Aggregation." *Journal of Philosophical Logic* 35: 569-585.
- Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge (Cambridge University Press).
- Pettit, P. (2001). "Deliberative Democracy and the Discursive Dilemma." *Philosophical Issues* 11: 268-299.
- Pigozzi, G. (2006). "Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation." *Synthese* 152: 285-298.
- Poisson, S. D. (1837). *Recherches sur la probabilité des jugements en matière criminelle et en matière civile: précédées des règles générales du calcul des probabilités*.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge/MA (Harvard University Press).
- Regenwetter, M., B. Grofman, A. A. J. Marley, and I. Tsetlin (2006). *Behavioral Social Choice*. Cambridge (Cambridge University Press).
- Riker, W. (1982). *Liberalism against Populism*. San Francisco (W.H. Freeman and Co).
- Roberts, K. (1995). "Valued Opinions or Opinionated Values: The Double Aggregation Problem." *Choice, Welfare and Development: A Festschrift in Honour of Amartya Sen*. K. Basu, P. K. Pattanaik, and K. Suzumura (eds.). Oxford (Oxford University Press), pp. 141-165.
- Saari, D. G. (1990). "The Borda dictionary." *Social Choice and Welfare* 7: 279-317.
- Salles, M. (edited with introduction) (2005). "The history of Social Choice." Special issue, *Social Choice and Welfare* 25: 229- 564.
- Satterthwaite, M. (1975). "Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions." *Journal of Economic Theory* 10: 187-217.

- Sen, A. K. (1966). "A Possibility Theorem on Majority Decisions." *Econometrica* 34: 491-499.
- Sen, A. K. (1970a). "The Impossibility of a Paretian Liberal." *Journal of Political Economy* 78: 152-157.
- Sen, A. K. (1970b). *Collective Choice and Social Welfare*. San Francisco (Holden-Day).
- Sen, A. K. (1977). "On weights and measures: informational constraints in social welfare analysis." *Econometrica* 45: 1539-1572.
- Sen, A. K. (1982). *Choice, Welfare and Measurement*. Oxford (Blackwell).
- Sen, A. (1983). "Liberty and social choice." *Journal of Philosophy* 80: 5-28.
- Sen, A. K. (1998). "The Possibility of Social Choice." Nobel lecture, December 8, 1998, Stockholm, available at: http://www.nobelprize.org/nobel_prizes/economics/laureates/1998/sen-lecture.pdf
- Shapley, L., and B. Grofman (1984). "Optimizing group judgment accuracy in the presence of interdependencies." *Public Choice* 43: 329-343.
- Stegenga, J. (2011). "An impossibility theorem for amalgamating evidence." *Synthese* (in press).
- Suppes, P. (2005). "The pre-history of Kenneth Arrow's social choice and individual values." *Social Choice and Welfare* 25: 319-226.
- Thomson, W. (2000). "On the axiomatic method and its recent applications to game theory and resource allocation." *Social Choice and Welfare* 18: 327-386.
- Tsetlin, I., M. Regenwetter, and B. Grofman (2003). "The Impartial Culture Maximizes the Probability of Majority Cycles." *Social Choice and Welfare* 21: 387-398.
- Vacca, R. (1921). "Opinioni Individuali e Deliberazioni Collettive." *Rivista Internazionale di Filosofia del Diritto*: 52-59.
- Ward, B. (1965). "Majority Voting and Alternative Forms of Public Enterprises." *The Public Economy of Urban Communities*. J. Margolis (ed.). Baltimore (Johns Hopkins Press).
- Weymark, J. (2006). "The Normative Approach to the Measurement of Multidimensional Inequality." *Inequality and Economic Integration*. F. Farina and E. Savaglio (eds.). London (Routledge), pp. 303-328.
- Wilson, R. (1975). "On the Theory of Aggregation." *Journal of Economic Theory* 10: 89-99.