

Political Science and Political Economy Working Paper

Department of Government

London School of Economics

No. 4/2007

Group Deliberation and the Transformation of Judgments: An Impossibility Result

Christian List (LSE)



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■

Group deliberation and the transformation of judgments: an impossibility result

Christian List

London School of Economics

March 8, 2007

Abstract

While a large social-choice-theoretic literature discusses the aggregation of individual judgments into collective ones, there is relatively little formal work on the transformation of individual judgments in group deliberation. I develop a model of judgment transformation and prove a baseline impossibility result: Any judgment transformation function satisfying some initially plausible conditions is the identity function, under which no opinion change occurs. I identify escape routes from this impossibility result and argue that successful group deliberation must be ‘holistic’: individuals cannot generally revise their judgments on a proposition based on judgments on that proposition alone but must take other propositions into account too. I discuss the significance of these findings for democratic theory.

1 Introduction

Aggregation and deliberation are often contrasted as two very different approaches to collective decision-making. While aggregation is the merging of

conflicting individual opinions into a social outcome, deliberation involves the discussion of these opinions and their possible revision by the individuals deliberating.¹ Jon Elster (1986, p. 112) summarizes the difference between the two approaches as follows: ‘The core of the [deliberative approach] ... is that rather than aggregating or filtering preferences, the political system should be set up with a view to changing them by public debate and confrontation.’ He continues, ‘there would [then] not be any need for an aggregation mechanism, since a rational discussion would tend to produce unanimous preferences.’

The contrast between the two approaches is probably overstated. More plausibly, they are complementary rather than contradictory, with deliberation often preceding aggregation (Miller 1992, Knight and Johnson 1994, Dryzek and List 2003). Nonetheless, social choice theory, arguably our best formal theory of collective decision-making, has focused mostly on aggregation and said relatively little about deliberation. Game theorists have recently given more attention to deliberation, investigating for example the incentives for and against truth-telling in deliberative settings (e.g., Austen-Smith and Feddersen 2006, Calvert 2006, Landa and Meirowitz 2006, Hafer and Landa forthcoming), but we still lack a social-choice-theoretic model of the transformation of opinion in deliberation. The aim of this paper is to contribute to filling this gap in the literature.

I model opinions as judgments – acceptance or rejection – on certain propositions, drawing on the emerging theory of judgment aggregation.² The

¹On deliberative democracy, see, for example, Cohen (1989), Dryzek (1990, 2000), Fishkin (1991), Gutman and Thompson (1996), Bohman and Rehg (1997).

²Inspired by the ‘doctrinal’ and ‘discursive’ paradoxes (Kornhauser and Sager 1986, Pettit 2001), judgment aggregation was formalized by List and Pettit (2002, 2004), combining Arrow’s (1951/1963) axiomatic approach to social choice theory with a logical representation of propositions. Further results and model extensions were provided by List (2003, 2004), Pauly and van Hees (2006), Dietrich (2006, forthcoming), Nehring and

propositions may be logically interconnected, so that the judgments on some propositions constrain those that can rationally be held on others. This way of modelling opinions is very general. As illustrated below, it can represent not only beliefs but also preferences. While the theory of judgment aggregation focuses on *judgment aggregation functions*, I here analyse what I call *judgment transformation functions*. Such a function maps each admissible profile of individual sets of judgments on the given propositions not to a collective set of judgments on them (as an aggregation function does), but to another, possibly revised profile of individual sets of judgments. The input profile represents the individuals' pre-deliberation judgments, the output profile their post-deliberation judgments. The process may or may not lead to consensus.

To develop an axiomatic approach to studying judgment transformation, I present a baseline impossibility result. It shows that, when the propositions under consideration are non-trivially interconnected, any judgment transformation function satisfying some initially plausible conditions must be maximally conservative: it must be the identity function, under which nobody ever changes his or her judgment on anything. The conditions, informally stated, are the following. First, any profile of rational individual judgment sets is admissible as input to the deliberation (*universal domain*). Second, the output of the deliberation is also a profile of rational individual judgment sets (*rational co-domain*). Third, if there is unanimity on *every* proposition before deliberation (not just on a single one), this is preserved after deliberation (*unanimity preservation*); fourth, the individuals do not always ignore their pre-deliberation judgments in forming their post-deliberation judgments

Puppe (2005), van Hees (forthcoming), Dietrich and List (forthcoming a, b), Dokow and Holzman (2005) and Pigozzi (2006). Judgment aggregation theory is closely related to abstract aggregation theory, e.g., Wilson (1975), Rubinstein and Fishburn (1986), Nehring and Puppe (2002), and to the theory of belief merging (Konieczny and Pino Pérez 2002).

(*minimal relevance*); and fifth, the deliberation focuses on one proposition at a time (*independence/systematicity*).³

Since only a degenerate deliberative process in which there is no opinion change satisfies these five conditions together, I consider relaxing some of them so as to avoid this negative result. I argue that, except in special cases, the first four conditions are hard to give up, but that the fifth – independence/systematicity – is a plausible candidate for relaxation. Thus the impossibility result shows that successful group deliberation as envisaged by deliberative democrats requires some kind of holism: individuals cannot generally revise their judgments on each proposition based on judgments on that proposition alone. The revision must take judgments on other propositions into account too. I give some examples of feasible judgment transformation functions, distinguishing between those that generate consensus and those that generate something less than consensus.

Throughout the paper, I relate my approach to other works on the intersection between social choice theory and deliberative democracy. Other important related contributions include the Alchourrón-Gärdenfors-Makinson theory of belief revision (1985), whose focus, however, is on a single individual's belief change in response to new information rather than belief revision in groups, and the theory of conciliation and consensus in belief merging (Konieczny 2004; Gauwin, Konieczny and Marquis 2005), whose key concept – a conciliation operator – is formally related to the present concept of a judgment transformation function. I conclude with a brief discussion of how the present work is related to game-theoretic works on deliberation.

³Systematicity adds to this a neutrality requirement, as formally defined below.

2 Opinions as judgments on propositions

How can the opinions held by a group of individuals at a given time be modelled? In this section, I explain how to model them as judgments on propositions expressed in logic (following List and Pettit 2002 and Dietrich forthcoming). I turn to their aggregation and transformation subsequently.

The ingredients of the model are the following. There is a group of individuals, labelled $1, 2, \dots, n$. The set of propositions considered by them is called the *agenda*. Propositions are represented by sentences, denoted ‘ p ’, ‘ q ’, ‘ r ’, ..., from propositional logic or a more general language.⁴ Propositional logic can express atomic propositions, without logical connectives, and compound propositions, with the logical connectives ‘not’, ‘and’, ‘or’, ‘if-then’ and ‘if and only if’. Atomic propositions are denoted ‘ a ’, ‘ b ’, ‘ c ’, ...; examples of compound propositions are ‘ a and b ’, ‘if a or b , then not c ’. As is standard in logic, one can distinguish between consistent and inconsistent sets of propositions.⁵

Each individual’s opinions (at a given time) are represented by a *judgment set*: the set of all those propositions in the agenda that the individual accepts.⁶ On the standard interpretation, to accept proposition ‘ p ’ means to believe ‘ p ’; thus judgments are binary cognitive attitudes. (Alternatively, to accept ‘ p ’ could mean to desire ‘ p ’; judgments would then be binary emotive

⁴Formally, the *agenda* is a subset X of the logic, where (i) X is closed under negation (if ‘ p ’ is in X , then so is ‘not p ’), (ii) ‘not not p ’ is identified with ‘ p ’, and (iii) X contains no tautological or contradictory propositions. Instead of propositional logic, any logic with some minimal properties can be used, including expressively richer logics such as predicate, modal, deontic and conditional logics (Dietrich forthcoming).

⁵In propositional logic, a set of propositions is *consistent* if all its members can be simultaneously true, and *inconsistent* otherwise. For example, {‘ a ’, ‘ a or b ’} is a consistent set, whereas {‘ a ’, ‘not a ’} and {‘ a ’, ‘if a then b ’, ‘not b ’} are not. More generally, consistency is definable in terms of a more basic notion of *logical entailment* (Dietrich forthcoming).

⁶Formally, individual i ’s *judgment set* is a subset J_i of the agenda X .

attitudes.) A judgment set is *consistent* if it is a consistent set of propositions and *complete* if it contains a member of each proposition-negation pair in the agenda. A combination of judgment sets across all the individuals in the group is called a *profile*.⁷

Let me give some examples of agendas of propositions on which groups of individuals may make judgments.

Example 1: Climate change. A panel of experts deliberates about climate change. The agenda on which the experts make judgments contains the following propositions and their negations: ‘Global CO₂ emissions are above 6500 million metric tons of carbon per annum’ (*a*); ‘If global CO₂ emissions are above this threshold, then the global temperature will increase by at least 1.2°C by 2030’ (‘if *a* then *b*’); ‘The global temperature will increase by at least 1.2°C by 2030’ (*b*).⁸

Example 2: A tenure case. A university committee deliberates about whether to grant tenure to a junior academic. The agenda on which the committee members make judgments contains the following propositions and their negations: ‘The candidate is excellent at teaching’ (*a*); ‘The candidate is excellent at research’ (*b*); ‘Excellence at both teaching and research is necessary and sufficient for tenure’ (*c* if and only if (*a* and *b*)); ‘The candidate should be given tenure’ (*c*).⁹

Example 3: Ranking candidates or policy options. A political decision-making body (e.g., a legislature, committee or electorate) deliberates about how to rank three or more candidates or policy options in an order of social preference. The agenda on which the individuals make judgments contains all propositions of the form ‘*x* is preferable to *y*’ and their negations,

⁷Formally, a *profile* is an n -tuple (J_1, J_2, \dots, J_n) .

⁸Variants of this example appear across the literature on judgment aggregation.

⁹This example is due to Bovens and Rabinowicz (2006).

where x and y are distinct candidates or options from some set of available ones and ‘is preferable to’ is a binary relation, with the rationality constraints on preferences built into the (predicate) logic.¹⁰

Example 4: Group membership. A club, society or association deliberates about which candidates from a list of three or more available ones should be granted membership, subject to the constraint that some, but not all, candidates should be granted membership. The agenda on which the individuals make judgments contains all propositions of the form ‘candidate j should be granted membership’ and their negations, where j is any available candidate and where the mentioned constraint is built into the logic.¹¹

Each of these agendas exhibits certain logical connections between propositions. By contrast, ‘trivial’ agendas such as those containing only a single proposition-negation pair are not typical in deliberative settings. To set them aside, I assume throughout the paper that the agenda is at least *minimally complex* in a sense that is satisfied in all the examples but whose technical details are not central for the exposition.¹²

¹⁰For details, see Dietrich and List (forthcoming-a), drawing on List and Pettit (2004).

¹¹The conjunction of the propositions in quotes is stipulated to be false and their disjunction to be true. The example is due to Kasher and Rubinstein (1997).

¹²Formally, I assume that (i) the agenda has an inconsistent subset of three or more propositions that becomes consistent upon removing one or more propositions, and (ii) it is not (nor isomorphic to) a set of propositions whose only logical connectives are ‘not’ and ‘if and only if’ (Dokow and Holzman 2005, Dietrich and List forthcoming-a). Properties (i) and (ii) are met in examples 1 to 4. For instance, the agenda containing ‘ a ’, ‘if a then b ’, ‘ b ’ and negations (example 1) satisfies (i) because its three-member inconsistent subset {‘ a ’, ‘if a then b ’, ‘not b ’} becomes consistent if one proposition is removed; it obviously satisfies (ii). In examples 2 to 4, a further property is met, which I assume only where explicitly stated: (iii) any proposition in the agenda can be deduced from any other proposition in it via a sequence of pairwise conditional entailments (Nehring and Puppe 2002).

3 The aggregation of judgments

Before I formally define the problem of judgment transformation, it is useful to recapitulate the problem of judgment aggregation: How can each profile of individual judgment sets on a given agenda be aggregated into a collective judgment set? This problem arises, for example, in referenda involving multiple propositions, in legislatures or committees deciding what factual and normative propositions to accept in legislation, in multi-member courts resolving cases on the basis of several premises, and in expert panels seeking to merge several scientific viewpoints into a collective viewpoint.

As illustrated in Figure 1, an *aggregation function* is a function that maps each profile of individual judgment sets in some domain to a collective judgment set.¹³

profile of individual judgment sets




Figure 1: Judgment aggregation

Examples of aggregation functions are majority voting, where each proposition is collectively accepted if and only if it is accepted by a majority of individuals; supermajority or unanimity rules, where each proposition is col-

¹³While a judgment aggregation function, as defined in List and Pettit (2002), goes back to Arrowian social choice theory, a related concept is that of a merging operator in belief merging (Konieczny and Pino Pérez 2002). Parallels are discussed in Pigozzi (2006).

lectively accepted if and only if it is accepted by a certain qualified majority of individuals, for example, two thirds, three quarters, or all of them; and dictatorships, where the collective judgment set is always the individual judgment set of the same antecedently fixed individual. Many other aggregation functions have been proposed in the literature.

Although the possibilities seem abundant, it is surprisingly difficult to find an aggregation function that guarantees consistent collective judgment sets. Notoriously, majority voting can produce inconsistent collective judgment sets even when all individual judgment sets are consistent – an insight which has become known as the ‘discursive paradox’ (Pettit 2001, extending Kornhauser and Sager 1986) and which generalizes Condorcet’s well-known paradox of majority voting (List and Pettit 2004).

To illustrate, consider the climate change example above (example 1), and suppose there are three experts on the panel, with opinions as shown in Table 1. The first expert judges that ‘ a ’, ‘if a then b ’ and ‘ b ’; the second judges that ‘ a ’, but ‘not (if a then b)’ and ‘not b ’; and the third judges that ‘if a then b ’, but ‘not a ’ and ‘not b ’. Clearly, each expert holds an individually consistent judgment set. Yet, the majority judgments are inconsistent: a majority accepts ‘ a ’, ‘if a then b ’ and ‘not b ’, an inconsistent set of propositions in the standard sense of logic.

	‘ a ’	‘if a then b ’	‘ b ’
Individual 1	True	True	True
Individual 2	True	False	False
Individual 3	False	True	False
Majority	True	True	False

Table 1: A profile of individual judgment sets

Although this observation shows that majority voting on interconnected propositions may generate inconsistent collective judgments, it leaves open

the question of whether other aggregation functions are immune to this problem and, if so, which ones. The recent literature on judgment aggregation has explored this question in great generality. One of the generic findings is that there exist no democratically appealing aggregation functions satisfying the following conditions:

Universal domain. The aggregation function accepts as admissible input any possible profile of consistent and complete individual judgment sets.

Collective rationality. The aggregation function produces as output consistent and complete collective judgment sets.

Unanimity preservation. If all individuals hold the same judgment set, this is also the collective judgment set.

Independence/systematicity. The collective judgment on any proposition p on the agenda depends only on individual judgments on p [and the pattern of dependence is the same across propositions]. (Independence omits, and systematicity includes, the clause in square brackets.)

Theorem 1 (*Dietrich and List forthcoming-a, Dokow and Holzman 2005*)
Any aggregation function satisfying universal domain, collective rationality, unanimity preservation and independence/systematicity is a dictatorship of one individual. (Whether the result requires independence or systematicity depends on how the minimal complexity of the agenda is defined.¹⁴)

This theorem extends several earlier results (List and Pettit 2002, Nehring and Puppe 2002, Pauly and van Hees 2006, Dietrich 2006). As it also applies

¹⁴If the agenda meets only properties (i) and (ii), systematicity is needed for the result; if it also meets property (iii), independence is enough.

to the example of ranking candidates or options (example 3 above), Arrow's impossibility theorem (1951/1963) on preference aggregation can be shown to be a corollary.¹⁵

Obviously, a lot could be said about how to interpret the result. To avoid the dictatorship conclusion, we must relax one of universal domain, collective rationality, unanimity preservation or independence/systematicity. All these routes have been discussed extensively in the literature. Arguably, the relaxation of universal domain and of independence/systematicity are the most promising routes. Since the focus of this paper is on deliberation rather than aggregation, however, I set these issues aside here and return to analogous issues in the discussion of judgment transformation. Theorem 1 suffices to illustrate that, under conditions of pluralism, the aggregation of individual judgments into collective ones raises a significant, albeit not insurmountable, challenge.

4 The transformation of judgments

To model the transformation (as opposed to aggregation) of judgments, only one further concept is needed: that of a judgment transformation function. As illustrated in Figure 2, a *transformation function* is a function that maps each profile of individual judgment sets in some domain to a profile of individual judgment sets in some co-domain (possibly the same as the domain).¹⁶ The input profile represents the individuals' judgments before deliberation,

¹⁵Formal constructions are given in Dietrich and List (forthcoming-a) and Dokow and Holzman (2005). For earlier derivations of Arrow-like results from judgment and abstract aggregation results, see Wilson (1975), List and Pettit (2004) and Nehring (2003).

¹⁶Just as a judgment aggregation function is related to a belief merging operator, so a judgment transformation function is related to a belief conciliation operator, defined in Gauwin, Konieczny and Marquis (2006). Pigozzi's (2006) insights on the parallels between judgment aggregation and belief merging apply, *mutatis mutandis*, to transformation too.

the output profile their judgments after deliberation. The output judgments may or may not differ from the input judgments, and the transformation may or may not lead to consensus.




Figure 2: Judgment transformation

A simple example of a transformation function is deference to the majority, where, after deliberation, each individual accepts all those propositions that a majority accepts before deliberation. But just as majority voting as an aggregation function fails to guarantee consistent collective judgments, so deference to the majority as a transformation function fails to guarantee consistent post-deliberation judgments (Pettit 2006). If each expert in the climate change example were to defer to the majority judgments in Table 1, for instance, the resulting post-deliberation judgments would be inconsistent.

An alternative to deference to the majority is deference to a supermajority or unanimity: here each individual accepts all those propositions after deliberation that a certain qualified majority (perhaps everyone) accepts before deliberation. If the supermajority threshold is sufficiently large, such a transformation function performs better than deference to a majority at securing consistency (List 2006b), but it has problems of its own. First, the individuals' post-deliberation judgments will be incomplete on all those

issues on which there is no supermajority consensus; and second, they may violate deductive closure: an individual may come to accept ‘ a ’ and ‘if a then b ’, because each receives the required supermajority support, and yet fail to accept ‘ b ’, because there is no supermajority consensus on ‘ b ’. Moreover, it is hard to solve these two problems together. Only a unanimity threshold can generally prevent violations of deductive closure (Dietrich and List forthcoming-b), but it also amplifies the incompleteness problem, because it permits the acceptance of only those propositions on which there is perfect consensus.

Other examples of transformation functions are opinion leader functions, where each individual adopts as his or her output judgment set the input judgment set of an antecedently fixed individual, the individual’s opinion leader. The opinion leader may differ for different individuals or be the same across individuals. In the latter case, the opinion leader function is the deliberative analogue of a dictatorial aggregation function. Finally, a degenerate transformation function is the identity function, where the output profile is always the same as the input profile: nobody ever changes his or her judgments. Below I introduce some other transformation functions.

Just as the theory of judgment aggregation seeks to characterize the logical space of possible aggregation functions satisfying various conditions, so I now want to explore the logical space of possible transformation functions satisfying certain conditions. This exercise is interesting from two perspectives. From a normative perspective, deliberative democrats have proposed a number of desiderata that a good deliberative process should meet. By formalizing these desiderata as conditions on a transformation function, we may ask whether they can be met together and what a transformation function looks like that meets them all. From a positive perspective, several effects of group deliberation on individual opinions are empirically known, ranging from deliberation-induced ‘meta-agreement’ (List, Luskin, Fishkin and

McLean 2000/2006) to deliberation-induced ‘group polarization’ (Sunstein 2002). By formalizing such effects as conditions satisfied by the underlying transformation function, we may investigate what transformation functions explain those empirically observed effects. While my model is consistent with either of these interpretations, the results provable in it must obviously be viewed differently depending on whether the conditions on a transformation function are interpreted normatively or positively.

5 An impossibility result

Let me introduce some conditions on a transformation function. Although each condition can be made plausible, I do not suggest that they are all equally compelling; indeed, I relax some of them below. However, they are useful for analysing the logical space of possible transformation functions.

Universal domain. The transformation function accepts as admissible input any possible profile of consistent and complete individual judgment sets.

Universal domain requires the transformation function to cope with conditions of pluralism on the input side, subject to the constraint of full individual rationality. A more demanding input condition would require the transformation function to cope even with less than fully rational individual judgments; but the main result below shows that even the present requirement is far from undemanding.

Rational co-domain. The transformation function produces as output a profile of consistent and complete individual judgment sets.

Rational co-domain requires the transformation function to produce output profiles that also meet the constraint of full individual rationality. Below

I also consider a weaker variant of this requirement.

Unanimity preservation. The transformation function maps any unanimous profile to itself.

Unanimity preservation requires that, if the individuals agree on *all* propositions on the agenda before deliberation, this all-way consensus be preserved after deliberation. This is distinct from either of the following, arguably less plausible conditions. The first is *unanimity generation*, the requirement that the transformation function map every profile to a unanimous profile, as captured by Elster’s quote in the introduction to this paper. Although endorsed by many deliberative democrats, especially those of a Habermasian orientation, this requirement seems unduly demanding and empirically unrealistic. The second condition from which unanimity preservation is distinct is *propositionwise unanimity preservation*, the requirement that if all individuals agree on a particular proposition ‘*p*’ before deliberation, without necessarily agreeing on other propositions, this consensus on ‘*p*’ be preserved after deliberation. This requires that even an incompletely theorized agreement on ‘*p*’ be preserved in deliberation, even if different individuals agree on ‘*p*’ for incompatible reasons. Such a requirement is neither normatively compelling nor empirically realistic.¹⁷ For example, upon noticing that you and I support ‘*p*’ for incompatible reasons, we may each decide to give up our belief in ‘*p*’. By contrast, unanimity preservation is the much milder requirement that an all-way consensus on everything – in those rare cases in which it occurs – be stable under deliberation.

To state the next condition, call two profiles *variants for a given individual* if they coincide for all individuals except the given one.

¹⁷On incompletely theorized agreements, see Sunstein (1994) and, in the context of judgment aggregation, List (2006a). For critiques of propositionwise unanimity preservation, see Bradley (2007), Mongin (2005) and Nehring (2005).

Minimal relevance. For each individual, there exists at least one admissible pair of variant profiles for which the individual’s output judgment sets differ.

Minimal relevance is the mild requirement that individuals do not always ignore their pre-deliberation judgments. It only rules out that an individual’s pre-deliberation judgments *never* make any difference to his or her post-deliberation judgments. It does not require those pre-deliberation judgments to make a difference more than once, nor does it say anything about *how* they should make a difference. Consistently with minimal relevance, the individual’s post-deliberation judgments could even respond negatively to his or her pre-deliberation judgments.

Independence/systematicity. Each individual’s output judgment on any proposition on the agenda depends only on the individuals’ input judgments on that proposition [and the pattern of dependence is the same across propositions]. (Again, independence omits, and systematicity includes, the clause in square brackets.)

Independence is a requirement of ‘local’ (as opposed to ‘holistic’) deliberation: the post-deliberation judgments on any proposition should be determined by pre-deliberation judgments on that proposition alone and should not depend on pre-deliberation judgments on other propositions. (Systematicity adds to this a neutrality requirement across propositions.) In the climate change example, independence requires, for instance, that individuals’ post-deliberation judgments on whether emissions above the relevant threshold would lead to the specified temperature increase (‘if *a* then *b*’) depend only on pre-deliberation judgments on this proposition and not on pre-deliberation judgments on, say, whether emissions are in fact above the threshold (‘*a*’). In the group membership example, to give another illustration, independence

requires that post-deliberation judgments on whether a particular candidate should be granted membership depend only on pre-deliberation judgments regarding this candidate, not on pre-deliberation judgments regarding other candidates.

Whether one considers independence compelling depends, from a normative perspective, on whether a focus on one proposition at a time is deemed desirable in deliberation and, from a positive perspective, on whether real-world deliberative processes display such a focus. Regardless of its intrinsic appeal, independence can be justified instrumentally by appealing to a *strategy-proofness* requirement, that is, the requirement that, when individuals choose whether or not to reveal their judgments truthfully in deliberation, truthfulness be a weakly dominant strategy. Under sufficiently permissive assumptions about individual incentives, strategy-proofness is met if and only if the transformation function satisfies independence and a further condition called *monotonicity*.¹⁸ This can be shown by adapting related results on aggregation (Dietrich and List forthcoming-c, Nehring and Puppe 2002). If one considers strategy-proofness desirable, as many deliberative democrats do, one may therefore have to endorse independence too.

Although these conditions may seem plausible baseline requirements on deliberation, it turns out that only a maximally conservative deliberative process can meet them all. Recall that the identity function is the transformation function that maps every profile to itself.

Theorem 2 *Any transformation function satisfying universal domain, rational co-domain, unanimity preservation, minimal relevance and independence/systematicity is the identity function. (As before, whether the result*

¹⁸Monotonicity requires that any individual's post-deliberation acceptance of a given proposition should not be reversed if, in the pre-deliberation profile, one additional individual supports the proposition, with other individuals' judgments remaining fixed.

*requires independence or systematicity depends on how the minimal complexity of the agenda is defined.*¹⁹)

Proof. Consider any transformation function satisfying the conditions of Theorem 2. Notice that the transformation function can be decomposed into n separate functions, where the i -th such function maps each profile of individual judgment sets in the domain of the transformation function to individual i 's output judgment set. Formally, each of these n functions – being a mapping from profiles of judgment sets to single judgment sets – is an aggregation function; interpretational issues can be set aside here. Since the underlying transformation function satisfies universal domain, rational codomain, unanimity preservation and independence/systematicity (minimal relevance is not yet used), each induced aggregation function satisfies universal domain, collective rationality (here meaning rationality of the output judgment sets), unanimity preservation and independence/systematicity. By Theorem 1 above, it is therefore a dictatorship of one individual.²⁰ This already shows that the underlying transformation function must be an opinion leader function, where each individual adopts as his or her output judgment set the input judgment set of some antecedently fixed individual, his or her opinion leader (or ‘dictator’ in the terminology of the induced aggregation function). Could any individual’s opinion leader be distinct from the individual him- or herself? If this were the case, minimal relevance would be violated contrary to the proof’s assumption, because the individual’s output judgment set would be invariant under any changes of his or her input judgment set. Each individual must therefore be his or her own opinion leader. Consequently, the transformation function is the identity function.

¹⁹Again, if the agenda meets only properties (i) and (ii), systematicity is needed for the result; if it also meets property (iii), independence is enough.

²⁰The qualifications regarding independence and systematicity in the footnote to Theorem 1 apply here too and thus carry over to Theorem 2.

This completes the proof.

Theorem 2 is clearly an impossibility result, showing that ‘successful’ group deliberation is impossible under the given conditions. The result casts doubt on these conditions. First, from a normative perspective, one does not want to impose conditions on deliberation that are so restrictive as to be met only by a degenerate deliberative process in which nobody ever changes his or her judgments. This would obviously be against the spirit of the normative literature on deliberative democracy. Further, Theorem 2 immediately implies that the five introduced conditions are inconsistent with the further condition of unanimity generation discussed above and implicit in many writings on deliberative democracy, as illustrated by Elster’s opening quote. Therefore, if one did want deliberation to produce unanimity, one could not also expect it to meet the five introduced conditions.

Secondly, from an empirical perspective, it is simply not true that group deliberation exhibits the extreme conservatism implied by the theorem. There is plenty of empirical evidence that opinions do change in deliberation; more on this later. Let me therefore go through the conditions one by one and consider relaxing them.

6 Escape routes

6.1 Relaxing universal domain

Universal domain requires the transformation function to cope with any level of pluralism in its input, subject only to the constraint of individual rationality. What happens if this is weakened to the requirement that it should cope only with those input profiles that exhibit a certain amount of cohesion among the individuals? Then there exist transformation functions other than the identity function that satisfy all the other conditions. An example

is deference to the majority, which guarantees consistent post-deliberation judgments provided only profiles are deemed admissible in which there are no distinct majorities supporting mutually inconsistent propositions. Could pre-deliberation judgments exhibit this amount of cohesion?

Suppose, for example, that even before deliberation the individuals agree on some cognitive or ideological dimension in terms of which to think about the propositions on the agenda – a ‘meta-agreement’ (List 2002, 2003) – and that, in consequence, they can be aligned from left to right on that dimension such that, for each proposition on the agenda, the individuals accepting the proposition lie either all to the left, or all to the right, of those rejecting it. Deference to the majority is then guaranteed to yield consistent and (absent ties) complete post-deliberation judgments. Consider, for example, the individual judgments over the agenda containing ‘ a ’, ‘if a then b ’ and ‘ b ’ as shown in Table 2, where the required left-right alignment of the individuals (here from 1 to 5) holds.

	Ind. 1	Ind. 2	Ind. 3	Ind. 4	Ind. 5
‘ a ’	True	False	False	False	False
‘if a then b ’	False	True	True	True	True
‘ b ’	False	False	False	True	True

Table 2: Unidimensionally aligned judgments

Notice that the majority judgments in Table 2 coincide with the judgments of the median individual relative to the left-right alignment, here individual 3. Generally, given any profile of the form described, no proposition can be supported by a majority unless it is also supported by the median individual. So, by deferring to the majority, individuals inherit the consistent judgments of the median individual.²¹

²¹Assuming full rationality of that individual, in accordance with universal domain.




Figure 3: Single-peaked ranking judgments

Another type of cohesion that is sufficient for consistent majority judgments applies to ranking judgments, as in example 3 above. Note that an individual's set of ranking judgments can be viewed as expressing a ranking of the given candidates or options from most to least preferable. Let some left-right ordering of these candidates/options be given; this could order them from most socialist to most capitalist, from most secular to most religious, from most urban to most rural, or in any other way. An individual's set of ranking judgments is called *single-peaked relative to that left-right ordering* if the individual has a most highly ranked candidate/option somewhere on the ordering with a decreasing ranking as candidates/options get more distant from it in either direction. This is illustrated by the two rankings in Figure 3 of the options x, y, z, v, w from most (1st) to least (5th) preferable. A profile (across individuals) is called *single-peaked* if there exists a left-right ordering of the candidates/options relative to which *all* individuals' ranking judgment sets are single-peaked. A classic result by Black (1948) shows that, for any single-peaked profile of ranking judgments, the resulting majority judgments are consistent. Consequently, if pre-deliberation ranking judgments are single-peaked, individuals can form consistent post-deliberation judgments by deferring to the majority.

Does such pre-deliberation cohesion provide a plausible escape route from the impossibility result on judgment transformation? Although pre-deliberation profiles may sometimes exhibit the required amount of cohesion, this cannot generally be assumed to be the case. Often the aim of deliberation is precisely to deal with pluralism. While consensus, or some other form of cohesion, may ideally be the *output* of deliberation, requiring it as its *input* appears to miss the point of deliberation. Nonetheless, one possible interpretation of the impossibility result is that, *if* rational co-domain, unanimity preservation, minimal relevance and independence/systematicity are required, *then* non-degenerate judgment transformation is possible only if individuals enter the process with sufficient initial cohesion.

6.2 Relaxing rational co-domain

Rational co-domain requires the individuals' output judgment sets to be both consistent and complete. Suppose this is weakened to the requirement that output judgment sets be merely consistent and deductively closed, where deductive closure means that individuals accept the implications of other accepted propositions, at least when they are also contained in the agenda. Deductive closure is much less demanding than completeness, as it is satisfied, for example, even by an empty judgment set. Requiring deductive closure in a deliberative setting is plausible (at least when confined to propositions on the agenda) because a frequently stated aim of deliberation is not just to lead people to form considered judgments on various propositions but also to make them aware of the implications of their judgments.

What happens if rational co-domain is relaxed in this way? Unfortunately, it does not open up a compelling escape route from the impossibility result. Any transformation function satisfying the weakened co-domain condition together with the other conditions (universal domain, unanimity

preservation, minimal relevance, independence/systematicity) is of the following form. For each individual, there exists a fixed subset of individuals in which he or she is included – his or her ‘peer group’ (in the limiting case, this could be the singleton set of the individual him- or herself) – such that the individual’s output judgment set is always the intersection of the input judgment sets among the individual’s peers. This is a corollary of a recent result on judgment aggregation without full rationality (Dietrich and List 2006, generalizing Gärdenfors 2006).²² Arguably, such a transformation function is no better, and possibly worse, than the identity function: it has the property that each individual’s output judgment set is always a subset of the individual’s input judgment set. At best an individual’s judgment set remains unchanged after deliberation, at worst it shrinks. How much it shrinks depends on the size of the individual’s peer group and the amount of disagreement among the peers. Such a transformation function perhaps instantiates the combination of a conservative and a sceptical attitude: an individual never comes to accept a proposition he or she did not accept in the first place and never continues to accept a proposition unless everyone in his or her peer group agrees with it.

6.3 Relaxing unanimity preservation

Unanimity preservation is the requirement that the transformation function map any unanimous profile to itself. Relaxing this requirement is not a very promising route. First, the requirement is already very mild, as argued above. But, secondly, even if one were prepared to drop it, this would not lead very far: under slightly stronger assumptions about how the propositions

²²The result still holds if the transformation function admits as input any possible profile of consistent and deductively closed judgment sets (not requiring completeness); here even a weakened independence/systematicity condition suffices for the result (Dietrich and List 2006). For an equivalent result in abstract aggregation, see Dokow and Holzman (2006).

on the agenda are interconnected, Theorem 2 continues to hold even without unanimity preservation.²³ I therefore set this route aside here.

6.4 Relaxing minimal relevance

Minimal relevance rules out that an individual's pre-deliberation judgment set never makes any difference to his or her post-deliberation judgment set. As in the case of unanimity preservation, relaxing minimal relevance does not lead to a strong escape route from the impossibility theorem. Not only is minimal relevance a mild requirement, but, as shown in the proof of Theorem 2, its relaxation makes possible only a very restrictive class of transformation functions, namely that of opinion leader functions. Under an opinion leader function, each individual adopts as his or her output judgment set the input judgment set of an antecedently fixed individual, the individual's opinion leader. As noted, an opinion leader function is analogous to a dictatorial aggregation function except that different individuals may defer to different opinion leaders. Obviously, such a transformation function is plausible at most in special circumstances, for example when individuals have reasons to think that their opinion leaders have a special expertise on the agenda of propositions under consideration.

6.5 Relaxing independence/systematicity

Independence requires the transformation function to determine the output judgment on any proposition solely on the basis of the individuals' input judgments on that proposition, with systematicity requiring in addition that

²³If systematicity is required, the relevant agenda assumption is the conjunction of (i), (ii) and a property called *asymmetry*; this follows from a result by Dietrich (forthcoming). If only independence is required, the relevant agenda assumption is *atomic closure* or *atomicity*; this follows from results by Pauly and van Hees (2006) and Dietrich (2006).

the pattern of dependence be the same across propositions. As noted above, this can be viewed as a constraint of ‘local’ deliberation, requiring the consideration of one proposition at a time. Given the limited appeal of the previous escape routes from the impossibility result, it seems natural to relax this requirement.

If we give it up, one possibility is for each individual to designate some propositions on the agenda as ‘premises’ and others as ‘conclusions’ and to generate his or her post-deliberation judgments by deferring to the pre-deliberation majority judgment on each premise and then deriving the judgments on other propositions from these majority judgments on the premises. If the premises are chosen as a ‘logical basis’ for the entire agenda – that is, they are mutually independent and any assignment of truth-values to them settles the truth-values of all other propositions – the resulting transformation function guarantees consistent and complete post-deliberation judgments and satisfies all the other conditions introduced above. The choice of premises and conclusions need not be the same across individuals.

While the present class of transformation functions is the judgment-transformation analogue of the ‘premise-based procedures’ in the context of aggregation (e.g., Pettit 2001), a more general class of transformation functions draws on the ‘sequential priority procedures’ of aggregation (e.g., List 2004). Here each individual determines a particular order of priority among the propositions on the agenda, interpreting earlier propositions in that order as epistemically (or otherwise) prior to later ones. The individual then considers the propositions one-by-one in the chosen order and forms his or her post-deliberation judgment on each proposition as follows. If the pre-deliberation majority judgment on the proposition is consistent with the judgments the individual has made on propositions considered earlier, then he or she defers to that pre-deliberation majority judgment; but if it is inconsistent with those earlier judgments, then he or she accepts the implications

of those earlier judgments. In the case of Table 1, for example, an individual may consider the propositions in the order ‘ a ’, ‘if a then b ’, ‘ b ’ (with negations interspersed) and then accept ‘ a ’ and ‘if a then b ’ by deferring to the pre-deliberation majority judgments while accepting ‘ b ’ by logical inference. The output profile under such a transformation function is sensitive to each individual’s chosen order of priority among the propositions. This property of the transformation function can be seen as a virtue or as a vice, depending on one’s perspective. On the one hand, it takes into account the fact that different propositions may have a different status (Pettit 2001, Chapman 2002). But on the other hand, it makes individuals manipulable by a cunning Rikerian ‘herethetician’ who leads them to consider propositions in a strategically adjusted order (Riker 1986, List 2004). In the next section, I discuss a third class of transformation functions that becomes possible once independence/systematicity is dropped.

What is the cost of violating independence? As already noted, a transformation function violating it may provide incentives for strategic misrepresentation of pre-deliberation judgments. To illustrate, consider the climate change example with individual judgments as shown in Table 1, and suppose the experts form their post-deliberation judgments in the sequential manner just explained, considering the propositions in the order ‘ a ’, ‘if a then b ’, ‘ b ’. Suppose, further, that the second expert, who does not accept that there will be the specified temperature increase (proposition ‘ b ’), does not want his or her colleagues to accept that proposition either. Under the sequential transformation function, he or she may lead them to accept ‘not b ’ by misrepresenting his or her judgment on ‘ a ’, strategically expressing the view that ‘not a ’. It can be shown that such possibilities of strategic manipulation arise as soon as the transformation function violates independence (adapting Dietrich and List forthcoming-c). The impossibility theorem presented can therefore also be seen as describing a dilemma between two problems of a

deliberative process: maximal conservatism on the one hand and strategic manipulability on the other.

7 Away from consensus

While it is sometimes held that group deliberation should bring about consensus, there is very little empirical evidence of this effect. It is also unclear whether achieving consensus is always normatively desirable. Moreover, if we impose unanimity generation as a condition on a transformation function, then the problem of judgment transformation formally collapses into that of judgment aggregation, and the only transformation functions satisfying universal domain, rational co-domain, unanimity preservation and independence/systematicity – dropping minimal relevance – are those opinion leader functions in which all individuals defer to the same opinion leader, the equivalent of an Arrowian dictator. This is particularly ironic in so far as the possibility of a deliberation-induced consensus is often proposed as a solution to, not a variant of, the notorious problem of aggregation.

Could deliberation bring about something less than consensus that is still helpful for democratic decision-making – for example, by facilitating the consistent aggregation of post-deliberation judgments?

A recent literature suggests that deliberation may have this effect, at least under favourable conditions (Miller 1992, Knight and Johnson 1994, Dryzek and List 2003). Recall the earlier discussion of the possibility that individuals agree on some cognitive or ideological dimension in terms of which to think about the relevant propositions or, in the case of ranking judgments, the candidates or policy options. While such ‘meta-agreement’ may be too demanding as a pre-condition for deliberation, it can more plausibly be expected as the outcome of deliberation. In an empirical study using data from Fishkin’s deliberative polls (List, Luskin, Fishkin and McLean 2000/2006),

such an effect has been identified. Groups of between 150 and 350 randomly sampled participants were interviewed on their ranking judgments over multiple policy options both before and after a period of group deliberation. Deliberation increased the proximity of these ranking judgments to single-peakedness, as defined above – a deliberation-induced movement towards ‘meta-agreement’.

Can we find an empirically plausible class of transformation functions to explain this effect? Let me define the class of constrained minimal revision functions. Under such a function, the transformation of judgments takes place in two stages. An input profile is given. In the first stage, the group identifies a particular set of judgment sets that are deemed admissible as output judgment sets conditional on the given input profile. Formally, this can be modelled as the application of a ‘focusing function’, which maps the input profile to a set of admissible output judgment sets. The latter set should ideally have the property that any profile constructible from it leads to consistent majority judgments. In the second stage, each individual selects an output judgment set from the identified set of admissible ones. Formally, this can be modelled as the application of a ‘minimal judgment revision policy’, under which each individual chooses an output judgment set from the set of admissible ones that is as close as possible to his or her input judgment set, relative to some distance metric over judgment sets (on the notion of minimal revision, see Schulte 2005). This fixes the output profile and thereby completes the definition.

Informally, the first stage involves the identification of the opinions that can reasonably be held after deliberation given the opinions before deliberation; and the second stage involves a change of individual opinions such that each individual ends up holding one of the reasonable ones. Whether an individual’s post-deliberation opinion coincides with his or her pre-deliberation opinion depends on whether it was already among the ones identified as

reasonable in the first stage.

Crucially, two ingredients of this definition allow a number of different specifications: the group's focusing function in the first stage, and each individual's distance metric over judgment sets in the second. Thus the definition specifies an entire class of transformation functions, one for each possible specification of these two ingredients.

To illustrate how a constrained minimal revision function can bring about a 'meta-agreement', suppose again a group deliberates about how to rank three or more policy options in an order of social preference (example 3), as in the deliberative polls studied empirically. The following constrained minimal revision function generates single-peaked output profiles. For a given pre-deliberation profile, it is first determined which left-right ordering of the options renders a maximal number of individuals' ranking judgment sets single-peaked, as defined above.²⁴ Now a ranking judgment set is deemed admissible if and only if it is single-peaked relative to the identified left-right ordering. This specifies the group's focusing function and completes the first stage. Each individual then minimally revises his or her ranking judgment set so as to adopt one of the admissible ones; here an individual's distance metric could be the 'Hamming distance', whereby the distance between any two judgment sets is the number of propositions on the agenda on which these judgment sets disagree.²⁵ This determines the post-deliberation profile and completes the second stage. By construction, this transformation function guarantees a single-peaked output profile.

Further empirical research is needed to confirm that a suitable constrained minimal revision function can explain the precise patterns of deliberation-

²⁴In List, Luskin, Fishkin and McLean (2000/2006), such a left-right ordering is called a 'largest structuring dimension'.

²⁵The 'Hamming distance' has been employed in the literature on judgment aggregation by Pigozzi (2006).

induced opinion change observed in deliberative polls and other experiments. To the best of my knowledge, however, the present approach is the first attempt to model a deliberation-induced ‘meta-agreement’ formally.

From a normative perspective, further questions need to be asked on whether a suitable constrained minimal revision function captures the requirements of a good deliberative process as discussed in the literature on deliberative democracy. A constrained minimal revision function satisfies universal domain, rational co-domain, unanimity preservation and minimal relevance, while violating independence/systematicity, but does it also satisfy some other desiderata of good deliberation?

Obviously, it does not satisfy unanimity generation, except when the focusing function picks out only one admissible output judgment set for each input profile. What about other conditions? As already illustrated, a constrained minimal revision function may satisfy *cohesion generation*, where a profile is defined to be *cohesive* if it generates consistent majority judgments. Cohesion generation is a particularly appealing condition when the deliberative process is expected to lead to a democratic decision.

Another condition is *stability under repeated rounds of deliberation*, the requirement that the transformation function map any output profile (that is, any profile in the function’s range) to itself, or equivalently, that repeated applications of the function lead to the same output as a single application. Any transformation function satisfying both unanimity preservation and unanimity generation also satisfies this condition, because its first application leads to unanimity and subsequent applications preserve this unanimity. By contrast, if the members of a group sit around a circular table and each individual defers to his or her neighbour on the right in forming post-deliberation judgments, then the resulting transformation function – a special kind of opinion leader function – violates the stability condition, as each round of deliberation yields a further permutation of the given pro-

file across individuals. A constrained minimal revision function, however, is stable under repeated rounds of deliberation so long as the group’s focusing function and the individuals’ distance metrics are sufficiently well-behaved.²⁶

Whether stability under repeated rounds of deliberation is a plausible requirement depends on how the judgment transformation function is interpreted. If it is meant to capture opinion change in a single round of deliberation, then there is no reason to expect subsequent rounds of deliberation to leave opinions fixed. But if it is meant to capture a complete deliberative process up to the point of ‘reflective equilibrium’ (Rawls 1971), then the stability condition is plausible – arguably more so than consensus generation.

8 Conclusion

I have formulated the problem of judgment transformation and presented a baseline impossibility result. The framework and result allow us to determine which combinations of conditions on a deliberative process are simultaneously attainable and which are not. Among the five basic conditions introduced, the first four, I have argued, can be relaxed only in special cases. The most plausible candidate for relaxation – under both normative and positives interpretations of the model – is the fifth condition: independence/systematicity.

This observation suggests that realistic deliberative processes must exhibit a certain kind of holism: Individuals cannot form their post-deliberation judgments on each proposition based on pre-deliberation judgments on that proposition alone, but must take into account pre-deliberation judgments on other propositions too. This holistic property of group deliberation is anal-

²⁶The focusing function must have the property that it maps any output profile constructed from any set of judgment sets in its range to a new set of judgment sets that still contains all the judgment sets in the given output profile. The distance metric must have the standard property that the distance of any judgment set from itself is uniquely zero.

ogous to the holistic property of theory testing in science. According to the Duhem-Quine ‘holism’ thesis in the philosophy of science (Quine 1951), one cannot generally test a single proposition in isolation, but only in conjunction with other propositions in a larger web of beliefs.

Let me conclude with some remarks about how the present approach is related to game-theoretic approaches to studying deliberation. While game-theoretic approaches seek to come up with a behavioural theory of group deliberation, analysing individuals’ incentives in deliberation and predicting their behaviour on that basis, the goal of the present approach is to illuminate the logical space of possible functional relations between pre- and post-deliberation judgments. The present approach is thus more akin to Arrowian social choice theory than to game theory, and the relationship between the two approaches is similar to that between social choice theory and the theory of mechanism design. The former investigates possible functional relations between individual preferences and social decisions and the latter investigates the various mechanisms available (or unavailable) for implementing these functional relations under certain incentive constraints. Therefore I see the present approach as complementary to game-theoretic approaches, not as competitive.

In addition, there are important bridges between the two approaches. As already noted, the condition of independence can be motivated game-theoretically, namely as a necessary condition (under certain assumptions) for the incentive-compatibility of truthfulness in deliberation. Thus my conclusion that realistic deliberative processes are likely to violate independence reinforces a central question in game-theoretic work on deliberation (e.g., Austen-Smith and Banks 2006): How can we design deliberative processes that induce participants to reveal their judgments truthfully? Broadly, there are at least two approaches towards this question. One may either go along the mechanism-design route and ask what deliberative processes

ensure truthfulness even when individuals are strategically motivated. Or one may go along a psychological route and ask under what conditions individuals are truthful even in the presence of strategic incentives. Which of these routes – or perhaps which combination of them – is most promising remains an important question for future research.

9 References

Alchourrón, C. E., P. Gärdenfors and D. Makinson. 1985. “On the Logic of Theory Change: Partial Meet Contraction and Revision Functions.” *Journal of Symbolic Logic* 50 (2, June): 510-530.

Arrow, K. 1951/1963. *Social Choice and Individual Values*. New York: Wiley.

Austen-Smith, D., and T. J. Feddersen. 2006. “Deliberation, Preference Uncertainty, and Voting Rules.” *American Political Science Review* 100 (2, May): 209-217.

Black, D. 1948. “On the Rationale of Group Decision-making.” *Journal of Political Economy* 56 (1, February): 23-34.

Bohman, J., and W. Rehg, eds. 1997. *Deliberative Democracy: Essays on Reason and Politics*. Cambridge, MA: MIT Press.

Bovens, L., and W. Rabinowicz. 2006. “Democratic Answers to Complex Questions – An Epistemic Perspective.” *Synthese* 150 (1, May): 131-153.

Bradley, R. 2007. “Reaching a Consensus.” *LSE Choice Group Working Papers* 3 (3).

Calvert, R. 2006. “Deliberation as Coordination through Cheap Talk.” Working paper, Washington University, St. Louis.

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

Cohen, J. 1989. “Deliberation and Democratic Legitimacy.” In A. Ham-

lin and P. Pettit, eds. *The Good Polity: Normative Analysis of the State*. Oxford: Blackwell.

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

Dietrich, F. Forthcoming. "A generalised model of judgment aggregation." *Social Choice and Welfare*.

Dietrich, F., and C. List. 2006. "Judgment aggregation without full rationality." Working paper, London School of Economics.

Dietrich, F., and C. List. Forthcoming-a. "Arrow's theorem in judgment aggregation." *Social Choice and Welfare*.

Dietrich, F., and C. List. Forthcoming-b. "Judgment aggregation by quota rules." *Journal of Theoretical Politics*.

Dietrich, F., and C. List. Forthcoming-c. "Strategy-proof judgment aggregation." *Economics and Philosophy*.

Dokow, E., and R. Holzman. 2005. "Aggregation of binary evaluations." Working paper, Technion Israel Institute of Technology.

Dokow, E., and R. Holzman. 2006. "Aggregation of binary evaluations with abstentions." Working paper, Technion Israel Institute of Technology.

Dryzek, J. 1990. *Discursive Democracy: Politics, Policy and Political Science*. New York: Cambridge University Press.

Dryzek, J. 2000. *Deliberative Democracy and Beyond*. Oxford: Oxford University Press.

Dryzek, J., and C. List. 2002. "Social Choice Theory and Deliberative Democracy: A Reconciliation." *British Journal of Political Science* 33 (1, January): 1-28.

Elster, J. 1986. "The Market and the Forum." In Elster, J., and A. Hylland, eds. *Foundations of Social Choice Theory*. Cambridge: Cambridge University Press: 103-132.

Fishkin, J. S. 1991. *Democracy and Deliberation*. New Haven: Yale

University Press.

Gärdenfors, P. 2006. "A Representation Theorem for Voting with Logical Consequences." *Economics and Philosophy* 22 (2, July): 181-190.

Gauwin, O., S. Konieczny and P. Marquis. 2005. "Conciliation and Consensus in Iterated Belief Merging." In *Proceedings of the Eighth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty* (ECSQARU'05): 514-526.

Gutman, A., and D. Thompson. 1996. *Democracy and Disagreement*. Cambridge, MA: Harvard University Press.

Hafer, C., and D. Landa. Forthcoming. "Deliberation as Self-Discovery and Institutions for Political Speech." *Journal of Theoretical Politics*.

van Hees, M. Forthcoming. "The limits of epistemic democracy." *Social Choice and Welfare*.

Kasher, A., and A. Rubinstein. 1997. "On the Question 'Who is a J?': A Social Choice Approach." *Logique et Analyse* 40 (160): 385-395.

Knight, J., and J. Johnson. 1994. "Aggregation and Deliberation: On the Possibility of Democratic Legitimacy." *Political Theory* 22 (2, May): 277-296.

Konieczny, S. 2004. "Belief base merging as a game." *Journal of Applied Non-Classical Logics* 14 (3): 275-294.

Konieczny, S. and R. Pino Pérez. 2002. "Merging Information Under Constraints: A Logical Framework." *Journal of Logic and Computation* 12 (5, October): 773-808.

Kornhauser, L. A., and L. G. Sager. 1986. "Unpacking the Court." *Yale Law Journal* 96 (1, November): 82-117.

Landa, D., and A. Meirowitz. 2006. "Information, Game Theory, and Deliberative Democracy." Working paper, New York University.

List, C. 2002. "Two Concepts of Agreement." *The Good Society* 11 (1): 72-79.

List, C. 2003. “A Possibility Theorem on Aggregation over Multiple Interconnected Propositions.” *Mathematical Social Sciences* 45 (1, February): 1-13.

List, C. 2004. “A Model of Path Dependence in Decisions over Multiple Propositions.” *American Political Science Review* 98(3, August): 495-513.

List, C. 2006a. “The Discursive Dilemma and Public Reason.” *Ethics* 116 (2, January): 362-402.

List, C. 2006b. “When to defer to supermajority testimony – and when not.” Working paper, London School of Economics.

List, C., R. C. Luskin, J. Fishkin and I. McLean. 2006/2006. “Deliberation, Single-Peakedness, and the Possibility of Meaningful Democracy: Evidence from Deliberative Polls.” Working paper, London School of Economics.

List, C., and P. Pettit. 2002. “Aggregating Sets of Judgments: An Impossibility Result.” *Economics and Philosophy* 18 (1, April): 89-110.

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

Mongin, P. 2005. “Spurious Unanimity and the Pareto Principle.” *LSE Choice Group Working Papers* 1 (5).

Miller, D. 1992. “Deliberative Democracy and Social Choice.” *Political Studies* 40 (Special Issue): 54-67.

Nehring, K. 2003. “Arrow’s theorem as a corollary.” *Economics Letters* 80 (3, September): 379-382.

Nehring, K. 2005. “The Impossibility of a Paretian Rational.” Working paper, University of California at Davis.

Nehring, K., and C. Puppe. 2002. “Strategyproof Social Choice on Single-Peaked Domains: Possibility, Impossibility and the Space Between.” Working paper, University of California at Davis.

Nehring, K., and C. Puppe. 2005. “Consistent Judgement Aggregation:

A Characterization.” Working paper, University of Karlsruhe.

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

Pettit, P. 2001. “Deliberative Democracy and the Discursive Dilemma.” *Philosophical Issues* 11 (October): 268-299.

Pettit, P. 2006. “When to defer to majority testimony – and when not.” *Analysis* 66 (3, July): 179-187.

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

Quine, W. V. O. 1951. “Two Dogmas of Empiricism.” *Philosophical Review* 60 (1, January): 20-43.

Rawls, J. 1971. *A Theory of Justice*. Oxford: Oxford University Press.

Riker, W. H. 1986. *The Art of Political Manipulation*. New Haven: Yale University Press.

Rubinstein, A., and P. C. Fishburn. 1986. “Algebraic Aggregation Theory.” *Journal of Economic Theory* 38 (1, February): 63-77.

Schulte, O. 2005. “Minimal belief change, Pareto-optimality and logical consequence.” *Economic Theory* 19 (1, January): 105-144.

Sunstein, C. R. 1994. “Political Conflict and Legal Agreement.” Tanner Lectures on Human Values, Harvard University.

Sunstein, C. R. 2002. “The Law of Group Polarization.” *Journal of Political Philosophy* 10 (2, June): 175-195.

Wilson, R. 1975. “On the Theory of Aggregation.” *Journal of Economic Theory* 10 (1, February): 89-99.