LSE Research Online

Article (refereed)

A model of path-dependence in decisions over multiple propositions

Christian List

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (http://eprints.lse.ac.uk) of the LSE Research Online website.

You may cite this version as:
List, Christian (2004). A model of path-dependence in decisions over multiple propositions [online]. London: LSE Research Online.
Available at: http://eprints.lse.ac.uk/archive/00000666

This is an electronic version of an Article published in American political science review, 98 (3), pp. 495-513 © 2004 Cambridge University Press.
http://titles.cambridge.org/journals/journal_catalogue.asp?mnemonic=psr

http://eprints.lse.ac.uk
Contact LSE Research Online at: Library.Researchonline@lse.ac.uk

A Model of Path-Dependence in Decisions over Multiple Propositions

Christian List
Lecturer in Political Science
Department of Government, LSE

Revised in February 2004

Address for correspondence:
C. List
Department of Government
London School of Economics
Houghton Street
London WC2A 2AE, U.K.
c.list@lse.ac.uk

Acknowledgments: I wish to thank many friends and colleagues for comments, discussions and encouragement: Kanyarat Bhanthumnavin, Luc Bovens, Richard Bradley, Steven Brams, Geoffrey Brennan, Bruce Chapman, Franz Dietrich, John Dryzek, Christian Elsholtz, Wulf Gaertner, Natalie Gold, Robert Goodin, Thomas Kittsteiner, Margaret Meyer, David Miller, Klaus Nehring, Philip Pettit, Ashley Piggins, Clemens Puppe, Alec Stone Sweet, Kamakshya Trivedi, Bertil Tungodden. I am grateful to the anonymous reviewers and the editor of the APSR for very helpful comments and suggestions. My research was supported by the Social and Political Theory Program at the Australian National University, the London School of Economics, and Nuffield College, Oxford.

A Model of Path-Dependence in Decisions over Multiple Propositions

Abstract. I model sequential decisions over multiple interconnected propositions and investigate path-dependence in such decisions. The propositions and their interconnections are represented in propositional logic. A sequential decision process is path-dependent if its outcome depends on the order in which the propositions are considered. Assuming that earlier decisions constrain later ones, I prove three main results: First, certain rationality violations by the decision-making agent – individual or group – are necessary and sufficient for path-dependence. Second, under some conditions, path-dependence is unavoidable in decisions made by groups. Third, path-dependence makes decisions vulnerable to strategic agenda setting and strategic voting. I also discuss escape-routes from path-dependence. My results are relevant to discussions on collective consistency and reason-based decision-making, focusing not only on outcomes, but also on underlying reasons, beliefs and constraints.

A Model of Path-Dependence in Decisions over Multiple Propositions

Social choice theory has improved our understanding of many collective decision problems: elections, referenda, decisions in legislatures, committees, judiciaries. Classical models of social choice represent collective decision-making as the aggregation of individual preferences or votes into collective outcomes (Arrow 1953; Sen 1970). While illuminating many aspects of collective decision-making, these models focus mainly on the ranking of, or choice among, alternative *outcomes*: candidates, policies, actions. They do not capture the *reasons* given in support of choices, the *beliefs* choices are based on and the *constraints* choices impose on other choices. The importance of all these is emphasized, for example, in theories of democratic deliberation, which often criticize classical social-choice-theoretic models for their narrow focus on preference aggregation.[1]

In this paper, I develop a model which seeks to capture such reasons, beliefs and constraints. Combining social choice theory and propositional logic, my model addresses decisions on multiple propositions based on one or several individuals' views, beliefs or judgments on these propositions. The propositions are *interconnected*: The views on some propositions logically constrain the views on others. The decisions are *sequential*: Different propositions are considered not simultaneously, but one after another – as in many governments, political parties, legislatures, committees, judiciaries – and earlier decisions constrain later ones. I investigate whether such decisions are *path-dependent*, i.e. whether their outcome depends on the order in which propositions are considered.

For example, if a political body decides first to endorse a tax cut, it may reject subsequent proposals to increase expenditure, although it might have accepted these proposals if it had considered them before deciding the tax issue. In an individual decision, there might be two arguments where both begin with propositions the individual finds plausible, both use valid

inferences, and yet they support opposite conclusions. Political rhetoric or "heresthetics" – the art of political manipulation – might exploit this phenomenon (Riker 1986).

Interest in decisions over multiple interconnected propositions was sparked by the "doctrinal paradox" in law and economics (e.g. Kornhauser and Sager 1986; Chapman 1998). A three-member court has to decide three propositions. *P*: The defendant did some action X. *Q*: The defendant had a contractual obligation not to do action X. *R*: The defendant is liable. Legal doctrine requires that *R* (the conclusion) be accepted if and only if both *P* and *Q* (the premises) are accepted. The first judge accepts both *P* and *Q* and therefore accepts *R*. The second accepts *P* but not *Q* and therefore rejects *R*. The third accepts *Q* but not *P* and therefore also rejects *R*. This leads to the apparent paradox that a majority (of 2/3) accepts *P*, a majority (of 2/3) accepts *Q*, and yet a majority (of 2/3) rejects *R*, although *P* and *Q* imply *R* and all judges' views are individually consistent. There is more than one way of resolving the issue. The court might privilege the majority votes on the premises and find liability, overruling the majority vote on the conclusion. Or it might privilege the majority vote on the conclusion and find no liability, overruling the majority votes on (one or both of) the premises. The order in which votes are taken and privileged may matter.

The example highlights not only the final outcome (*R*), but also the supporting reasons (*P*, *Q*) and the constraint linking reasons and outcomes (*R* if and only if *P* and *Q*). In general, such reasons and constraints need not be exogenously fixed, say by legal doctrine; they may themselves be contested. A model of decisions over multiple propositions seeks to capture cases where decisions are made on reasons and constraints in addition to outcomes – more broadly, cases where a system of beliefs or judgments is built up, constrained by consistency and mutual support, a requirement sometimes called "integrity" (Dworkin 1986).

Are such cases politically relevant? Are there any political bodies that do *not* just make separate decisions – at best mutually detached, at worst mutually inconsistent – *but* that build

up a coherent system of beliefs or judgments? Some scholars think that, at most, judiciaries have this property. Others, particularly deliberative democrats, see the construction of such a coherent system as central to democratic politics (Pettit 1997, 2001a). Here the "doctrinal paradox" has been interpreted as a paradox of "collective coherence" or "discursive dilemma", illustrating a trade-off between majoritarian responsiveness and consistency in collective decisions (Brennan 2001; Pettit 2001a). Many social choice theorists have argued that a consistent collective set of views – more strongly, "collective will" – is democratically infeasible (e.g. Riker 1982). Others have defended its feasibility or argued that the issue is still open (e.g. Coleman and Ferejohn 1986). But the debate has been centered around results on preference aggregation, such as Arrow's theorem (1951) and its descendants. For a richer analysis of collective consistency, it is desirable to develop a model of how groups can form collective views on multiple propositions. I aim to contribute to that development.

The first model of decisions over multiple propositions inspired by the "doctrinal paradox" is presented in List and Pettit (2002) and has been generalized by Pauly and van Hees (2003).[2] Earlier precursors are Guilbaud (1966) and Murakami (1968).[3] But so far only *simultaneous* decisions on multiple interconnected propositions have been modeled.

My model represents *sequential* decisions – individual and collective – and allows the study of *path-dependence*. I show that the order in which propositions are considered may affect the decision outcome even when the initial views of the decision-maker(s) on each proposition are held fixed. In individual decisions, certain violations of perfect rationality are necessary and sufficient for path-dependence. A boundedly rational individual may be susceptible to path-dependence. In collective decisions, path-dependence may occur even when all individuals are perfectly rational. This follows from an impossibility result analogous to Arrow's theorem (1951). Path-dependence may be the price a group has to pay for achieving collective consistency.

I also investigate two types of strategic manipulation opened up by path-dependence. Unsurprisingly, path-dependence may enable an agenda-setter to manipulate the outcome by controlling the decision-path (Riker 1982). More surprisingly, the mere existence of an alternative decision-path that *would* change the outcome – even if that path is never adopted – may create incentives for strategic expression of untruthful views. This finding is related to the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975), albeit in a sequential context.[4] Finally, I explore some escape-routes from path-dependence.

Previous social-choice-theoretic work on path-dependence has addressed cases where a winning outcome or ordering is determined in sequential pairwise majority voting over multiple alternatives (e.g. Plott 1973), but not cases where a system of accepted propositions with logical interconnections is built up over time. In a judicial context, Kornhauser (1992) has observed that if judges' preferences are non-separable – their judgments on different cases are interdependent – then sequential decisions may be path-dependent. This finding motivates developing explicit models of such interconnections.[5] Page (1997; 2003) has modeled sequential decisions on multiple projects made by a single decision-maker. The projects create externalities for each other, and negative externalities may lead to path-dependence. Page's results also emphasize constraints between decisions, but these are due to externalities between projects, not logical interconnections between propositions.

My model is a first approach to representing how decisions on multiple interconnected propositions evolve over time. It makes many simplifications that call for generalization in future work. But, methodologically, I use propositional logic to model a class of decision problems not adequately captured by classical models. Substantively, I aim to identify how path-dependence arises in such decision problems and what its implications are.

TWO STYLIZED EXAMPLES

A Collective Decision Problem

The first example concerns a multi-member government making multiple decisions that constrain each other across time (Pettit 2001b). The propositions under consideration are:

P: A new (costly) education project shall be implemented.

Q: A new (costly) health care project shall be implemented.

R: A new (costly) defense project shall be implemented.

T: Taxes shall be increased.

The pursuit of all three projects leads to a budget deficit, given current tax revenues, while the pursuit of two or fewer projects is possible within the existing budget. I here assume that the government members unanimously accept a balanced budget constraint: *P*, *Q* and *R* can be accepted simultaneously *only if T* is also accepted, formally:

"If *P* and *Q* and *R*, then *T*."

My model also allows disagreement on such a constraint. Let the government have three members, with views as shown in Table 1. Each member's views are consistent.

[Table 1 about here]

The propositions are considered one by one. At each point, a majority vote is taken on the proposition under consideration. If the majority opinion is consistent with all prior decisions, the government decides the new proposition according to that opinion; if it is not, the government decides the new proposition based on what those prior decisions imply. So earlier decisions constrain later ones. Consider two alternative sequences in which the propositions might come up.

Case 1. The balanced budget constraint is considered in January and accepted unanimously. The tax increase proposal is considered in February and rejected unanimously. The education and health proposals are considered in March and April and are each accepted

by a 2/3 majority. When the defense proposal is considered in May, it is supported by a 2/3 majority, but the government is already committed to two costly projects, to not increasing taxes and to balancing the budget. To avoid losing credibility (and the next election), the government rejects the defense proposal, overruling its majority opinion on that proposal. Overall, the government accepts *P* and *Q* (and the budget constraint) and rejects *R* and *T*.

Case 2. Again, the balanced budget constraint is considered and accepted in January. But the education, health and defense proposals are considered in February, March and April and are each accepted by a 2/3 majority. As a tax increase has not been ruled out yet, accepting all three proposals is still consistent with the balanced budget constraint. But in May a dilemma arises. All government members oppose a tax increase, but the prior commitments to balancing the budget and to pursuing all three costly projects necessitate accepting a tax increase. Overall, the government accepts *P*, *Q*, *R* and *T* (and the budget constraint).

Cases 1 and 2 lead to mutually contradictory outcomes. The individual views are identical in both cases, as is the government's method of deciding each proposition by a majority vote, unless prior commitments intervene. The two cases differ only in the order in which the propositions are considered: The decision process is path-dependent. Could the government use a different decision rule to avoid such path-dependence?

An Individual Decision Problem

The second example concerns political argumentation. It is related to work by Aragones et al (2002) on rhetoric, suggesting that one can change an agent's beliefs without giving her new information, just by organizing existing information differently. While Aragones et al analyze the effects of rhetorical analogies, the present example suggests that similar effects can be achieved by adjusting the order in which propositions are presented to the agent(s) for consideration. As this phenomenon can be strategically exploited, it is relevant to political manipulation (Riker 1986; McLean 2002). Imagine an individual agent who has a disposition

to accept each of the following propositions: If she were to consider each proposition *in isolation*, then she would find the proposition sufficiently plausible to accept.

P: Young people are entitled to unrestricted freedom after school.

Q: Compulsory national service reduces crime among young people.

R: Compulsory national service is justifiable.

"If *P*, then not *R*": If young people are entitled to unrestricted freedom after school, then compulsory national service is not justifiable.

"If *Q*, then *R*": If compulsory national service reduces crime among young people, then compulsory national service is justifiable.

One may notice an inconsistency between these propositions, but I leave this point aside until later.[6] The agent considers the propositions one by one. At each point, she finds the proposition under consideration initially plausible, as assumed. So if the proposition is consistent with her previously accepted propositions, she accepts it. If it is not, she rejects it, to respect her prior decisions. As in the government example, the agent's earlier decisions constrain her later ones. Consider two alternative sequences in which the propositions might come up – or in which a Rikerian "heresthetician" might present them to the agent:

Case 1. At times 1 and 2, the agent considers whether young people are entitled to unrestricted freedom after school (*P*) and whether this implies that compulsory national service is unjustifiable ("If *P*, then not *R*"). She finds each proposition plausible and accepts it. So, at time 3, when she considers whether compulsory national service is justifiable (*R*), she rejects that proposition. At time 4, she considers and accepts that compulsory national service would reduce crime among young people (*Q*). At time 5, she considers whether such crime reduction would be sufficient to justify compulsory national service ("If *Q*, then *R*") and rejects the proposition, to respect her previous decisions. Overall, the agent accepts *P*, *Q* and "If *P*, then not *R*", and rejects *R* and "If *Q*, then *R*".

Case 2. At time 1, the agent considers and accepts that, if compulsory national service were to reduce crime among young people, this would make compulsory national service justifiable ("If *Q*, then *R*"). At time 2, she considers and accepts that it would reduce crime (*Q*). So, at time 3, when she considers whether compulsory national service is justifiable (*R*), she concludes that it is. At time 4, she considers the proposition that, if young people were entitled to unrestricted freedom after school, this would make compulsory national service unjustifiable ("If *P*, then not *R*"). She accepts this proposition, having not yet considered whether young people are entitled to such freedom. At time 5, she considers this issue (*P*) and rejects the proposition, following her prior decisions. So she concludes that compulsory national service takes priority over young people's freedom after school. Overall, the agent accepts *Q*, "If *Q*, then *R*", *R*, and "If *P*, then not *R*" and rejects *P*.

Cases 1 and 2 lead to mutually contradictory outcomes. Only a single agent is involved, who in both cases is initially disposed to accept each proposition. In each case, the agent uses only valid logical inferences. The two cases differ only in the order in which the propositions are considered: The decision process is path-dependent. What has happened?

THE MODEL

The Propositions

I use a simple language of propositional logic, which includes both *atomic* and *compound* propositions. Atomic propositions are the most basic ones; they are represented by the (finitely many) letters *P*, *Q*, *R*, *S*, … Compound propositions state interconnections between other propositions; they involve *logical connectives*, such as $\neg$ (not), $\wedge$ (and), $\vee$ (or), $\rightarrow$ (implies), $\leftrightarrow$ (if and only if).

Definition. The set of all propositions of the language, ***L***, is defined by three rules:

- Each of *P*, *Q*, *R*, *S*, … is a proposition.
- If ϕ and ψ are propositions, then so are $\neg\phi$, $(\phi \wedge \psi)$, $(\phi \vee \psi)$, $(\phi \rightarrow \psi)$, $(\phi \leftrightarrow \psi)$.

- There are no other propositions.

Examples of atomic and compound propositions. Examples of atomic propositions are the premises (P, Q) and the conclusion (R) in the court example, the policy and tax proposals (P, Q, R, T) in the government example, and the propositions about freedom (P), crime reduction (Q) and the justifiability of national service (R) in the argumentation example; examples of compound propositions are the legal rule "R if and only if P and Q", formally ($R \leftrightarrow (P \wedge Q)$); the balanced budget constraint "If P and Q and R, then T", formally, $((P \wedge Q \wedge R) \rightarrow T)$; and the propositions relating freedom and crime to the justifiability of national service, "If P, then not R" and "If Q then R", formally, $(P \rightarrow \neg R)$ and $(Q \rightarrow R)$.

Truth-values. An assignment of truth-values (true or false) to the propositions has the following properties. For any two propositions ϕ, ψ,

- $\neg\phi$ is true if and only if ϕ is false;
- $(\phi \wedge \psi)$ is true if and only if *both* ϕ and ψ are true;
- $(\phi \vee \psi)$ is true if and only if *at least one of* ϕ or ψ is true;
- $(\phi \rightarrow \psi)$ is true if and only if it is *not* the case that [ϕ is true *and* ψ is false];
- $(\phi \leftrightarrow \psi)$ is true if and only if ϕ and ψ are both true *or* both false.[7]

Properties of propositions. A proposition ϕ is a *tautology* if it is true for *every* assignment of truth-values; ϕ is a *contradiction* if it is true for *no* assignment of truth-values. A set of propositions Φ *logically entails* a proposition ϕ – denoted $\Phi \models \phi$ – if, for *every* assignment of truth-values, [if *all* propositions in Φ are true, then ϕ is also true]. Two propositions ϕ and ψ are *logically equivalent* if each entails the other. A set of propositions Φ is *logically consistent* if there exists an assignment of truth-values for which every proposition in Φ is true; Φ is *logically inconsistent* if there exists no such assignment.

The propositions under decision. The set of propositions on which decisions are to be made is a finite subset X of $\boldsymbol{L}$. The propositions in X must be neither tautologies nor contradictions (so they are "open"). I assume that X contains at least two distinct atomic propositions, P and Q, and the compound proposition $(P \wedge Q)$. The use of conjunction ($\wedge$) implies no loss of generality; using other logical connectives would yield similar results.[8] For every $\phi \in X$, $\neg\neg\phi$ is identified with ϕ. I further assume that X contains proposition-negation pairs: Whenever $\phi \in X$, then $\neg\phi \in X$. In the government example, X is the set containing the propositions P, Q, R, T, $((P \wedge Q \wedge R) \rightarrow T)$ and their negations.

Propositions in politics. Individual proposals offered for acceptance or rejection in political decisions, such as the policy and tax proposals above, can be represented as atomic propositions. There are also many instances of compound propositions in politics, although they are not usually so construed. I have interpreted a budget constraint rule as a compound proposition in a political decision-making context, and propositions on the implications of freedom and crime for the justifiability of national service as compound propositions in a political argumentation context. For further examples, think of some proposal, P, and some amendment clause, A. Then the proposal with that amendment can be represented as the compound proposition $(P \wedge A)$.[9] Or think of some political goal, G, and some action, A. Then the claim that A is a necessary means for achieving G can be represented as the compound proposition $(G \rightarrow A)$. Political argumentation frequently involves claims about how different political propositions, values, goals, or actions are interconnected. Such claims can be represented as compound propositions. Crucially, compound propositions need not be exogenously fixed; they can themselves be the subject of disagreement and decision.

The Propositional Attitudes of an Agent

I first assume that there is one agent. The agent can be an individual or a group of individuals acting collectively. I discuss the collective case explicitly below.

Propositional attitudes. I assume that, for each proposition ϕ, the agent has a *propositional attitude* towards ϕ, which can be interpreted in (at least) two ways. On one interpretation, it is the agent's *fully endorsed view* on ϕ: acceptance or non-acceptance of ϕ. On another, it is the agent's *initial disposition* on ϕ: the verdict (acceptance or non-acceptance) she *would* give on ϕ if she *were* to consider ϕ in isolation, not referring to other propositions. Which interpretation applies depends on whether the agent has already considered ϕ. If she has, she may have a fully endorsed view on ϕ. But even if not, she may still have an initial disposition on ϕ if she *were* to consider ϕ. That initial disposition might be a function of the *prima facie* plausibility of ϕ from the agent's perspective. The agent's propositional attitudes are represented by an acceptance/rejection function.

Definition. An *acceptance/rejection function* ("*AR-function*") is a function $\delta: X \rightarrow \{1, 0\}$.

For each proposition $\phi \in X$, $\delta(\phi)=1$ means that the agent accepts ϕ (or has an initial disposition to accept ϕ); $\delta(\phi)=0$ means that the agent does not accept ϕ (or has an initial disposition not to accept ϕ). Note that $\delta(\phi)=0$ does not by itself mean that the agent accepts the negation of ϕ (or has such an initial disposition). The agent accepts that negation (or has such an initial disposition) if and only if $\delta(\neg\phi)=1$. For brevity, I usually refer to $\delta(\phi)=1$ as "acceptance" of ϕ, and to $\delta(\phi)=0$ as "rejection" of ϕ, bearing in mind the more precise interpretation offered here.[10] As discussed below, an agent's propositional attitudes may or may not satisfy certain rationality conditions.[11] I also discuss various rationality violations.

The propositional attitudes of a group. If the agent is a group, ascribing propositional attitudes to that group does not presuppose any metaphysical assumptions about "group minds". As illustrated by the government example and discussed below, the group's propositional attitude on a proposition can be its majority verdict on that proposition, or the outcome of some other aggregation over the propositional attitudes of the group members.

A Sequential Decision Process

A decision-path. A decision-path is the order in which the propositions are considered in a sequential decision process. It can be interpreted in (at least) two ways. On one interpretation, it is the *temporal order* in which the propositions are considered: Earlier propositions in the path come up earlier *in time* than later ones. On another, it is the *order of importance or priority* assigned to the propositions: Earlier propositions in the path are more important or "weightier" than later ones.

Definition. A *decision-path* on X is a one-to-one function $\Omega : \{1, 2, ..., k\} \rightarrow X$, where k is the number of propositions in X.

Here $\Omega(1)$, $\Omega(2)$, ..., $\Omega(k)$ are the first, second, ..., kth propositions to be considered.[12]

The notion of a sequential decision-process. In defining a sequential decision process, it is useful to interpret the agent's propositional attitudes as initial dispositions: The agent enters the decision process without fully endorsed views on the propositions, but with initial dispositions. The agent considers the propositions in the order given by a decision-path Ω, and decides at each point whether or not to accept the proposition under consideration, say ϕ. The outcome of the decision process can then be interpreted as the *set of fully endorsed views* the agent has formed on the propositions *after* having considered them one by one along the path Ω. How does the agent decide whether or not to accept each proposition in the sequence? Suppose proposition ϕ is under consideration. There are two cases. *Either* the agent's initial disposition on ϕ is consistent with the propositions (if any) the agent has accepted earlier: If so, the agent accepts or rejects ϕ according to that initial disposition. *Or* there is a logical inconsistency between the initial disposition on ϕ and some previously accepted propositions: If so, the agent requires a method of resolving this inconsistency.

A conflict resolution rule. I call a method by which the agent can resolve an inconsistency between her disposition on a new proposition and previously accepted

propositions a *conflict resolution rule*. Under the *priority-to-the-past* rule, as in the two examples, the agent resolves the inconsistency by accepting the logical implications of previously accepted propositions and overruling the initial disposition on the new proposition. So earlier decisions constrain later ones, but not vice-versa. Precedent-based decision making is a form of priority-to-the-past decision making. For simplicity, I only consider the priority-to-the-past rule here. This is a restriction, as there are cases where other conflict resolution rules are more plausible, but the model can be generalized.[13]

When is the priority-to-the-past rule plausible? Let me suggest an illustrative list of cases (neither exhaustive, nor mutually exclusive) in which the priority-to-the-past rule is plausible. First, suppose the decision-path is the temporal order in which the propositions are considered. Then the priority-to-the-past rule is plausible in cases where decisions, once made, are to hard to overrule. This may be because:

- Earlier decisions may have created commitments that cannot (easily) be overruled, for several reasons: (i) Legal reasons: Earlier decisions may be legally binding; they may have involved passing certain legislation, changing constitutional rules, or they may serve as precedents. (ii) Social or cognitive reasons: Earlier decisions may be so entrenched in the system of beliefs or judgments of the agent(s) that they cannot be overruled without a substantial cost. (iii) Strategic reasons: If earlier decisions are publicly known, the agent may lose credibility by not respecting those decisions; this loss of credibility may be costly; the agent might lose the next election, for example.
- Earlier decisions may have led to actions that (i) are irreversible, or (ii) are perceived by the agent(s) to be irreversible, or (iii) are too costly to reverse in practice.

Second, suppose the decision-path is the order of importance or priority among the propositions. Then the priority-to-past rule is a natural way of making decisions based on

that order: Decisions on more important propositions constrain decisions on less important ones. My model is neutral between different interpretations of the priority-to-the-past rule.

Definition. A *priority-to-the-past decision process* is the following procedure. Consider the propositions along the decision-path Ω: $\phi_1 := \Omega(1)$ at time 1, $\phi_2 := \Omega(2)$ at time 2, …, $\phi_k := \Omega(k)$ at time k, where k is the number of propositions in X. For each $t = 1, 2, \ldots, k$, Φ_t is the set of all propositions accepted *up to* time t, where Φ_t is defined inductively as follows:

***t* = 0** (added): Φ_0 is the empty set.

***t* > 0**: Proposition ϕ_t is under consideration. There are two cases:

Case I: (i) Previously accepted propositions entail ϕ_t ($\Phi_{t-1} \models \phi_t$) or (ii) they entail $\neg\phi_t$ ($\Phi_{t-1} \models \neg\phi_t$). Then

$$\Phi_t := \begin{cases} \Phi_{t-1} \cup \{\phi_t\} & \text{if (i)} \\ \Phi_{t-1} \cup \{\neg\phi_t\} & \text{if (ii).}^{14} \end{cases}$$

Case II: ϕ_t and $\neg\phi_t$ are each consistent with previously accepted propositions (neither $\Phi_{t-1} \models \phi_t$ nor $\Phi_{t-1} \models \neg\phi_t$). Then

$$\Phi_t := \begin{cases} \Phi_{t-1} \cup \{\phi_t\} & \text{if } \delta(\phi_t) = 1 \\ \Phi_{t-1} & \text{if } \delta(\phi_t) = 0.^{15} \end{cases}$$

Define the *outcome set* of the decision process, where the agent's AR-function is δ and the decision-path is Ω, to be $M(\delta, \Omega) := \Phi_k$.

By definition, a proposition that is inconsistent with previously accepted propositions is never accepted in a priority-to-the-past decision process:

Proposition 1. For any δ and any Ω, $M(\delta, \Omega)$ is logically consistent.

Finally, I define a priority-to-the-past decision process to be decisive if, for every proposition in X, either the proposition or its negation is accepted in the process.

Definition. $M(\delta, \Omega)$ is *decisive* if, for every $\phi \in X$, either $\phi \in M(\delta, \Omega)$ or $\neg\phi \in M(\delta, \Omega)$.

An illustration from politics. In a discussion of credible commitment and path-dependence, Miller and Schofield (2003) give two empirical illustrations of how prior

commitments can constrain political decisions. First, they argue that in 1972, as a result of the Republican Party's socially conservative and anti-civil-rights decisions in the 1960s (including Goldwater's presidential candidacy in 1964), Nixon had little choice but to continue his party's ideological course and to pursue a "Southern strategy" that included a socially conservative law-and-order position. They suggest that Nixon would have been so constrained even if, hypothetically, he had considered an alternative "Northern strategy" that included a pro-civil-rights position – a position that Republicans had been inclined towards up to the early 1960s, to appeal to black voters. Second, they argue that in the mid-1990s, "Clinton was under a great deal of pressure from liberals in his party to restore the Democratic Party to the economic liberalism of the New Deal." (p. 258) But to do so, Clinton would have had to appeal to Southern and other social conservatives, which "was no longer credible for the party that had supported civil rights and affirmative action for 30 years" (p. 258). So Clinton had little choice but to pursue a more moderate economic policy. Miller and Schofield argue that "candidates are constrained by recent historical events that introduce an asymmetry in [their] calculations" (p. 258), where those events include their parties' previous decisions and commitments. The analysis suggests that Nixon and Clinton had strong strategic reasons – in the sense defined above – for respecting their parties' prior commitments. But note that the presence of constraints on decisions due to prior commitments is not the same as path-dependence and leads to path-dependence only under certain conditions, as shown below.

The Notion of Path-Dependence

A priority-to-the-past decision process is invariant under changes of the decision-path if its outcome is the same for all decision-paths. It is weakly path-dependent if there exist (at least) two decision-paths with *different* outcomes. It is strongly path-dependent if there exist (at least) two decision-paths with *contradictory* outcomes. Let δ be given.

Definition. $M(\delta, \Omega)$ is *invariant under changes of the decision-path* if, for any two decision-paths Ω_1 and Ω_2, $M(\delta, \Omega_1) = M(\delta, \Omega_2)$.

Definition. $M(\delta, \Omega)$ is *weakly path-dependent* if, for some proposition $\phi \in X$, there exist two decision-paths Ω_1 and Ω_2 such that $\phi \in M(\delta, \Omega_1)$ and $\phi \notin M(\delta, \Omega_2)$.

If ϕ is such a proposition, I say that $M(\delta, \Omega)$ is weakly path-dependent *with respect to* ϕ.

Definition. $M(\delta, \Omega)$ is *strongly path-dependent* if, for some proposition $\phi \in X$, there exist two decision-paths Ω_1 and Ω_2 such that $\phi \in M(\delta, \Omega_1)$ and $\neg\phi \in M(\delta, \Omega_2)$.

If ϕ is such a proposition, I say that $M(\delta, \Omega)$ is strongly path-dependent *with respect to* ϕ. Strong path-dependence implies weak path-dependence. The converse holds only in special conditions:

Proposition 2. Suppose that $M(\delta, \Omega)$ is decisive for all decision-paths Ω. Then $M(\delta, \Omega)$ is strongly path-dependent if and only if it is weakly path-dependent.[16]

Weak and strong path-dependence may differ when $M(\delta, \Omega)$ is indecisive for some decision-path Ω. Below I state necessary and sufficient conditions for both kinds of path-dependence, but I focus mainly on strong path-dependence.

RATIONALITY VIOLATIONS AND PATH-DEPENDENCE

Rationality Conditions on Propositional Attitudes

I introduce four rationality conditions which an agent's propositional attitudes may or may not satisfy. Completeness requires that the agent should accept *at least* one member of each proposition-negation pair. Weak consistency requires that the agent should accept *at most* one member of each proposition-negation pair. Strong consistency requires that the propositions accepted by the agent can be simultaneously true. Deductive closure requires that the agent should accept all implications of other propositions she accepts. These

concepts provide rationality criteria over and above the standard criterion of transitivity of preferences. Let δ be the agent's AR-function:

Definition. δ is *complete* if, for any $\phi \in X$, $\delta(\phi)=1$ or $\delta(\neg\phi)=1$.

Definition. δ is *weakly consistent* if there exists no $\phi \in X$ such that $\delta(\phi)=1$ and $\delta(\neg\phi)=1$.

Definition. δ is *strongly consistent* if the set $\{\phi \in X : \delta(\phi)=1\}$ is logically consistent.[17]

For any subset $\Phi \subseteq X$, I write $\delta(\Phi) = 1$ to mean $[\delta(\phi)=1$ for all $\phi \in \Phi]$.

Definition. δ is *deductively closed* if the following holds: For any logically consistent $\Phi \subseteq X$ and any $\phi \in X$, if $\delta(\Phi) = 1$ and $\Phi \models \phi$, then $\delta(\phi)=1$.

The conditions are interrelated. Strong consistency implies weak consistency, but not vice-versa. To illustrate, if the agent accepts P, $(P \rightarrow Q)$ and $\neg Q$, then weak, but not strong, consistency is satisfied: The agent accepts no proposition and its negation simultaneously, but P, $(P \rightarrow Q)$ and $\neg Q$ cannot be simultaneously true. Weak consistency and deductive closure jointly imply strong consistency.[18] In the Appendix, I prove the following result:

Lemma 1. δ violates strong consistency if and only if there exist two logically consistent subsets $\Psi_1, \Psi_2 \subseteq X$ and a proposition $\phi \in X$ such that $[\delta(\Psi_1) = 1$ and $\Psi_1 \models \phi]$ and $[\delta(\Psi_2) = 1$ and $\Psi_2 \models \neg\phi]$.

So a necessary and sufficient condition for a violation of strong consistency is the existence of two logically consistent sets of propositions Ψ_1 and Ψ_2 and a proposition ϕ such that the agent accepts all the propositions in Ψ_1 and all in Ψ_2, but Ψ_1 logically entails ϕ, while Ψ_2 logically entails $\neg\phi$. If this condition is met, I say that δ violates strong consistency *with respect to* ϕ. I say that δ violates weak consistency *with respect to* ϕ if $\delta(\phi)=1$ and $\delta(\neg\phi)=1$.[19] I say that δ violates deductive closure *with respect to* ϕ if there exists a logically consistent set of propositions $\Phi(\subseteq X)$ such that the agent accepts all the propositions in Φ, Φ logically entails ϕ, but the agent does not accept ϕ, i.e. $\delta(\Phi) = 1, \Phi \models \phi$ and $\delta(\phi)=0$.

Necessary and Sufficient Conditions for Path-Dependence

Let me revisit the two examples. Consider the argumentation example. Are the agent's propositional attitudes consistent? They are weakly, but not strongly, consistent: The agent is disposed to accept Q and $(Q \rightarrow R)$, which entail R, and to accept P and $(P \rightarrow \neg R)$, which entail $\neg R$. They also violate deductive closure, as she is disposed to accept P and $(P \rightarrow \neg R)$, but not $\neg R$. Consider the government example. Each government member's views are complete, weakly and strongly consistent, and deductively closed. But what about the propositional attitudes of the government acting collectively, as determined by majority voting on each proposition? There are majorities for P, Q, R and $((P \wedge Q \wedge R) \rightarrow T)$, which entail T, but there is no majority for T. So the majority views violate deductive closure with respect to T. They also violate strong consistency, as there is a majority (unanimity) for $\neg T$. This suggests that violations of the rationality conditions, particularly of strong consistency, might lead to path-dependence. The following results confirm this suggestion. Consider any agent – individual or group – making decisions on multiple propositions, where δ is the agent's AR-function.

Theorem 1. For any $\phi \in X$, $M(\delta, \Omega)$ is strongly path-dependent with respect to ϕ if and only if δ violates strong consistency with respect to ϕ.

Proof: The proofs of all theorems are in the Appendix. ■

By Theorem 1, if (and only if) the agent's propositional attitudes violate strong consistency with respect to some proposition ϕ, there exist (at least) two decision-paths such that under one path ϕ is accepted, while under the other $\neg\phi$ is accepted. An agent's propositional attitudes violate strong consistency *either* when they violate weak consistency *or* when they are weakly consistent but an inconsistency is "hidden" by a violation of deductive closure:

Proposition 3. Suppose δ is complete and weakly consistent. For any $\phi \in X$, δ violates strong consistency with respect to ϕ if and only if δ is not deductively closed with respect to one of ϕ or $\neg\phi$.

The conjunction of Theorem 1 and Proposition 3 yields a necessary and sufficient condition for strong path-dependence in the case where the agent's propositional attitudes are complete and weakly consistent. In this case, there exist (at least) two decision-paths with mutually inconsistent outcomes on some proposition ϕ if and only if the agent's propositional attitudes violate deductive closure with respect to ϕ:

Theorem 2. Suppose δ is complete and weakly consistent. For any $\phi \in X$, $M(\delta, \Omega)$ is strongly path-dependent with respect to ϕ if and only if δ is not deductively closed with respect to one of ϕ or $\neg\phi$.

While strong path-dependence is ruled out when the agent's propositional attitudes are strongly consistent (by Theorem 1), weak path-dependence is still possible:

Theorem 3. Suppose δ is strongly consistent. For any $\phi \in X$, $M(\delta, \Omega)$ is weakly path-dependent with respect to ϕ if and only if δ is not deductively closed with respect to ϕ.

Theorem 3 states that, if the agent's propositional attitudes are strongly consistent, then there exist (at least) two decision-paths with different outcomes on some proposition ϕ if and only if these propositional attitudes violate deductive closure with respect to ϕ.[20]

If the agent is an individual, the conditions for avoiding path-dependence are clear. If this individual's propositional attitudes are weakly consistent and deductively closed (hence strongly consistent), she is immune to path-dependence in a priority-to-the-past decision process. By contrast, if her propositional attitudes violate strong consistency or deductive closure, she may be manipulable by a "heresthetician" presenting the propositions to her in some strategic order. While it is unsurprising that *irrationality* can make individuals manipulable, my results imply that even a *boundedly rational* individual may be manipulable. As shown, a violation of deductive closure is sufficient for path-dependence, even when weak consistency is satisfied.[21] Deductive closure violations are empirically plausible: While individuals might be sufficiently rational not to accept a proposition and its

negation simultaneously, they may lack the computational resources needed to derive all logical implications of the beliefs they hold. My results therefore suggest that bounded rationality might be of relevance to certain phenomena of political manipulation. This suggestion raises interesting empirical questions on whether individuals satisfy or violate the identified rationality conditions and on whether there are cases of political manipulation targeted specifically at the fact that agents are boundedly, but not perfectly, rational.

If the agent is a group, the conditions for avoiding path-dependence are less clear. It would be desirable to find a method of aggregating the group members' views into a collective AR-function that satisfies completeness, weak consistency and deductive closure (and thus strong consistency). By Theorems 2 and 3, this would rule out path-dependence. The multi-member government example has illustrated that majority voting does not generally have these properties. I next show that this is not an accidental property of majority voting: Under certain conditions, no aggregation function with the desired properties exists.

PATH-DEPENDENCE AT THE COLLECTIVE LEVEL: A GENERAL RESULT

The Setting

I now assume that the agent is a *group* of individuals, represented by the set $N = \{1, 2, \ldots, n\}$ ($n \geq 2$). The views of each individual, $i \in N$, on the propositions in X are represented by an *AR-function* $\delta_i : X \rightarrow \{1, 0\}$, defined as before. I here interpret the individuals' propositional attitudes as fully endorsed views, and assume, as a best-case scenario, that each individual's AR-function δ_i is complete, weakly consistent and deductively closed (and thus strongly consistent).

Definition. A *profile of individual AR-functions* ("*profile*") is a vector of AR-functions across the individuals, $<\delta_i> = <\delta_1, \delta_2, \ldots, \delta_n>$.

Different individuals may hold different views not only on atomic propositions, but also on compound ones.[22] This captures the fact that in politics individuals may disagree not only

on specific policies, actions, or reasons, but also on relations between reasons and actions, means-ends relations, constraints between different policies, budget constraints, and so on.

Aggregation Functions

To determine the group's propositional attitude to a proposition ϕ – its *initial disposition* on ϕ – a method of aggregating the n individual views on ϕ into a single collective propositional attitude on ϕ is needed. For each ϕ, the vector of 0s and 1s, $<\delta_1(\phi), \delta_2(\phi), \ldots, \delta_n(\phi)>$, must be aggregated into a single value of either 0 (non-acceptance of ϕ) or 1 (acceptance of ϕ).

Definition. An *aggregation function* is a function δ: $\{0,1\}^n \rightarrow \{0,1\}$.

Examples of aggregation functions are majority voting, as used above, and a dictatorship of one individual, analogous to the definition in Arrow's preference-based model (1951):

Definition. *Majority voting* is the aggregation function δ: $\{0,1\}^n \rightarrow \{0,1\}$ defined as follows. For any $<d_1, d_2, ..., d_n> \in \{0,1\}^n$,

$$\delta(d_1, d_2, ..., d_n) = \begin{cases} 1 \text{ if } \sum_{i \in N} d_i > {}^n/_2 \\ 0 \text{ otherwise.} \end{cases}$$

Definition. A *dictatorship of individual* $i \in N$ is the aggregation function δ: $\{0,1\}^n \rightarrow \{0,1\}$ defined as follows. For any $(d_1, d_2, ..., d_n) \in \{0,1\}^n$, $\delta(d_1, d_2, ..., d_n) = d_i$.

Majority voting has the attractive property of giving all individuals equal weight in determining the group's propositional attitude on any proposition; it satisfies *anonymity*. A dictatorship violates not only anonymity, but also the weaker condition of *non-dictatorship*:

Definition. An aggregation function δ is *anonymous* if, for any $<d_1, d_2, ..., d_n> \in \{0,1\}^n$ and any permutation σ: $N \rightarrow N$, $\delta(d_1, d_2, ..., d_n) = \delta(d_{\sigma(1)}, d_{\sigma(2)}, ..., d_{\sigma(n)})$; δ is *non-dictatorial* if it is not a dictatorship of any $i \in N$.

How can an aggregation function generate a single collective AR-function based on a profile of individual AR-functions? For each profile $<\delta_i>$, an aggregation function δ induces a collective AR-function $\delta_{<\delta_i>}$: $X \rightarrow \{0,1\}$, where,

for each $\phi \in X$, $\delta_{<\delta_i>}(\phi) := \delta(\delta_1(\phi), \delta_2(\phi), \ldots, \delta_n(\phi))$.

The collective AR-function $\delta_{<\delta_i>}$ represents the group's propositional attitudes, as determined by applying the aggregation function δ to the profile $<\delta_i>$.[23]

A restriction is built into my definition of an aggregation function: The aggregation over the *n* individual views is performed on a proposition-by-proposition basis. This mirrors the restriction introduced by independence of irrelevant alternatives in Arrow's classical model. Under Arrow's independence condition, preferences are aggregated based on pairwise comparisons between alternatives, i.e. based on (what might be seen as) pairwise ranking propositions. Whether such a restriction is plausible depends on the context. In many sequential decision problems, collective decisions are taken on a proposition-by-proposition basis, by voting separately on each proposition as it arises. But, as discussed below, a possible generalization of my model is to relax propositionwise aggregation.

Impossibility Results on Determining Collective Propositional Attitudes

I noted that majority voting may sometimes fail to induce a complete, weakly consistent and deductively closed collective AR-function. Are there other aggregation functions which always generate collective AR-functions with these properties? I now show that, under some minimal conditions, the answer is negative.

Definition. The *universal domain*, **U**, is the set of all logically possible profiles of individual AR-functions satisfying completeness, weak consistency and deductive closure.

Proposition 4. (Corollary of List and Pettit 2002) There exists no aggregation function δ (satisfying anonymity) which induces, for every $<\delta_i> \in \mathbf{U}$, a complete, weakly consistent and deductively closed collective AR-function $\delta_{<\delta_i>}$.

Pauly and van Hees (2003) have recently shown that, if we drop anonymity, the unique aggregation function satisfying the remaining conditions is a dictatorship of one individual:

Proposition 5. (Corollary of Pauly and van Hees 2003) An aggregation function δ induces, for every $<\delta_i>\in \mathbf{U}$, a complete, weakly consistent and deductively closed collective AR-function $\delta_{<\delta_i>}$ if and only if δ is a dictatorship of some individual $i\in N$.

These results are analogous to Arrow's theorem on preference aggregation. They are not corollaries of Arrow's theorem, as there is no straightforward mapping from views on multiple interconnected propositions into preferences (see also List and Pettit 2003).[24]

Impossibility Results on Path-(In)Dependence

Consider the implications of Propositions 4 and 5. *Any* anonymous or just non-dictatorial (propositionwise) aggregation function generating complete and weakly consistent collective AR-functions *necessarily* produces deductive closure violations for *some* profiles in the universal domain. By Theorem 2, for these profiles, there exist different decision-paths with mutually inconsistent outcomes. This mechanism underlies the following result.

Theorem 4. There exists no aggregation function δ (satisfying anonymity or non-dictatorship) such that, for every $<\delta_i>\in \mathbf{U}$, $M(\delta_{<\delta_i>}, \Omega)$ is decisive and invariant under changes of the decision-path.

Moreover, the unique aggregation function that guarantees decisiveness and path-*in*dependence for all profiles in the universal domain is a dictatorship of one individual:

Theorem 5. Let δ be an aggregation function. For every $<\delta_i>\in \mathbf{U}$, $M(\delta_{<\delta_i>}, \Omega)$ is decisive and invariant under changes of the decision-path if and only if δ is a dictatorship of some individual $i\in N$.

In summary, suppose a priority-to-the-past decision process satisfies two conditions. Universal domain: Any logically possible profile of (rational) individual views on the propositions is acceptable as input. Decisiveness: The decision process produces a determinate decision on every proposition. Then the unique (propositionwise) aggregation

function for which the decision process is generally path-*in*dependent is a dictatorship of one individual. Add a further condition. Anonymity: All individuals have equal weight in the decision on each proposition. Or, even more weakly, non-dictatorship: There does not exist a dictatorial individual. Then there exists no (propositionwise) aggregation function for which the decision process is generally path-*in*dependent. This does *not* mean that the decision process is path-dependent for *every* profile. But no anonymous or non-dictatorial aggregation function *guarantees* path-*in*dependence. Under *any* such aggregation function, there exist *some* profiles for which the decision process is (strongly) path-dependent.[25]

Does this imply that, at a collective level, path-dependence is *in principle* unavoidable? And, if so, what is the cost of this? I first address the second question, and show that path-dependence opens up two types of strategic manipulation. I then turn to the first question, and explore some escape-routes from path-dependence.

THE COST OF PATH-DEPENDENCE: STRATEGIC MANIPULABILITY

Manipulation by Agenda Setting

Whenever the agent's propositional attitudes violate strong consistency with respect to ϕ, an agenda-setter – whoever chooses the decision-path – may have power to determine whether the outcome of the decision process will be ϕ or $\neg\phi$. By Theorem 1, given a violation of strong consistency with respect to ϕ, there exists one decision-path leading to ϕ and another leading to $\neg\phi$. Given sufficient information and computational power, an agenda-setter can determine the decision-path required to bring about the preferred outcome. In the multi-member government example, an agenda-setter who cares strongly about the defense proposal (proposition *R*) might advocate the decision-path of case 2, which results in the acceptance of *R*. An agenda-setter who opposes the defense proposal might advocate the decision-path of case 1, which results in the rejection of *R*. This parallels the problem of

agenda-dependence in preference-based models, albeit within the new domain of decisions on multiple interconnected propositions (e.g. Plott and Levine 1978; Riker 1982, ch. 7).

Manipulation by Expression of Untruthful Views

Suppose a decision process is (strongly) path-dependent, but a particular decision-path has been fixed. If some individual (or group) cares strongly about certain *later* propositions in the decision-path, they might strategically express untruthful views on certain *earlier* ones, as decisions on the later propositions will be affected by decisions on those earlier ones. In the government example, suppose individual 3 cares most about the defense proposal (proposition R), and is willing to sacrifice her support for the health proposal (proposition Q), to get her way on the defense issue. Suppose the decision-path is as in case 1. In April, when Q is considered, individual 3 might untruthfully vote against Q, bringing about a majority *rejection* of Q. Then, in May, when R is considered, there would no longer be a conflict between prior commitments and the majority verdict on R; the government would be able to follow that majority verdict and accept R. Without individual 3's strategic intervention, the outcome set of the decision process would have been $\{((P \wedge Q \wedge R) \rightarrow T), P, Q, \neg R, \neg T\}$. With that intervention, the outcome set is $\{((P \wedge Q \wedge R) \rightarrow T), P, \neg Q, R, \neg T\}$, an outcome individual 3 prefers: Individual 3 has an incentive to express an untruthful view on Q.[26]

I now define strategic incentives formally. Let $\boldsymbol{M}$ be the set of all possible outcome sets $M(\delta, \Omega)$ ($\subseteq X$) of a priority-to-the-past decision process. I assume that each individual has certain preferences over these possible outcome sets. Individual i's most preferred outcome set – for brevity, denoted δ_i like the individual's AR-function – is the one which includes precisely those propositions she individually accepts, $\{\phi \in X : \delta_i(\phi) = 1\}$. The closer an outcome set is to this most preferred one – in a sense defined below – the more she prefers that outcome set. The individual's preferences over the outcome sets are represented by an

ordering R_{i,δ_i} on $\boldsymbol{M}$ (reflexive, transitive and connected). The subscripts i and δ_i indicate that the ordering depends not only on individual i, but also on her AR-function δ_i. For any two outcome sets Φ_1, Φ_2, $\Phi_1 R_{i,\delta_i} \Phi_2$ means that individual i weakly prefers Φ_1 to Φ_2. I write $\Phi_1 P_{i,\delta_i} \Phi_2$ to mean that individual i strictly prefers Φ_1 to Φ_2, formally [$\Phi_1 R_{i,\delta_i} \Phi_2$ and not $\Phi_2 R_{i,\delta_i} \Phi_1$]. I now formalize the assumption that the closer an outcome set is to individual i's most preferred one, the more she will prefer that outcome set.

Definition. An outcome set Φ_1 is *at least as close to* δ_i as an outcome set Φ_2 if, for *every* $\phi \in X$, $\delta_i(\phi)=\Delta_2(\phi)$ implies $\delta_i(\phi)=\Delta_1(\phi)$, where, for j=1, 2, $\Delta_j(\phi)=1$ if $\phi \in \Phi_j$ and $\Delta_j(\phi)=0$ if $\phi \notin \Phi_j$.

Informally, one outcome set is at least as close to an individual's most preferred outcome set as another if there is no proposition which the individual accepts and which is included in the second outcome set but not in the first, and no proposition which the individual rejects and which is excluded from the second outcome set but included in the first.

Assumption 1. If an outcome set Φ_1 is at least as close to δ_i as another outcome set Φ_2, then individual i weakly prefers Φ_1 to Φ_2, i.e. $\Phi_1 R_{i,\delta_i} \Phi_2$.

As the "at least as close" relation is only a partial, not generally complete, ordering over the possible outcome sets, this assumption is typically consistent with more than one preference ordering; so it is not maximally restrictive. My definition of strategic incentives is a translation of the classical definition into the framework of decisions on multiple propositions (compare Gibbard 1973; Satterthwaite 1975).

Definition. In the decision process $M(\delta, \Omega)$ defined on the domain $\mathbf{D}$, an individual $i \in N$ has an *incentive to express an untruthful AR-function* at the profile $<\delta_i> \in \mathbf{D}$ if there exists an AR-function δ^*_i ($\neq \delta_i$) (where $<\delta_1,\ldots,\delta^*_i, \ldots,\delta_n> \in \mathbf{D}$) such that $\Phi^* P_{i,\delta_i} \Phi$, where $\Phi^* = M(\delta_{<\delta_1, \ldots, \delta^*_i, \ldots, \delta_n>}, \Omega)$ and $\Phi = M(\delta_{<\delta_i>}, \Omega)$.[27]

Informally, individual *i* has an incentive to express an untruthful AR-function at the profile $<\delta_i>$ if three conditions hold: (i) If individual *i* expresses her truthful AR-function δ_i (holding the other individuals' AR-functions fixed), the decision process leads to the outcome set Φ. (ii) If individual *i* expresses the "untruthful" AR-function δ^*_i (holding the other individuals' AR-functions fixed), the decision process leads to the outcome set Φ^*. (iii) Individual *i* strictly prefers Φ^* to Φ.

The government example shows that path-dependent decision processes may give individuals incentives to express untruthful AR-functions. Whether or not an individual has such an incentive in a given case depends on several factors: the individual's preference ordering, the decision-path, whether the individual's views are pivotal for the collective outcome on some propositions. Below I prove that (weak) path-dependence is a *necessary* condition for the existence of individuals with incentives to express untruthful views. It is not a *sufficient* condition. Even in cases of (strong) path-dependence there may not exist a single individual who is pivotal for the outcome on a relevant proposition; so there may not exist an individual who can single-handedly manipulate the outcome. A more technical analysis might be used to show that, under some conditions, (strong) path-dependence implies that there exists a *coalition of individuals* with an incentive to express untruthful views.

Avoiding Strategic Manipulation

Neither type of strategic manipulation is possible when a decision process is invariant under changes of the decision-path. In the case of manipulation by agenda setting this is obvious. Agenda setting – determining the order in which the propositions are considered – has no effect when the decision process is path-*in*dependent. In the case of manipulation by expression of untruthful views, the following result holds.

Definition. An aggregation function δ is *weakly monotonic* if, for any $(d_1, d_2, ..., d_n)$, $(e_1, e_2, ..., e_n) \in \{0,1\}^n$, [for every $i \in N$, $d_i \geq e_i$] implies $\delta(d_1, d_2, ..., d_n) \geq \delta(e_1, e_2, ..., e_n)$.

Propositionwise majority voting is a weakly monotonic aggregation function.

Theorem 6. Suppose δ is a weakly monotonic aggregation function, and $M(\delta, \Omega)$ is invariant under changes of the decision-path for every profile $<\delta_i>$ in some domain **D**. For every individual $i \in N$ and every profile $<\delta_i> \in$ **D**, the following holds. For every AR-function δ^*_i ($\neq \delta_i$) (with $<\delta_1,\ldots,\delta^*_i, \ldots,\delta_n> \in$ **D**), $\Phi = M(\delta_{<\delta_i>}, \Omega)$ is at least as close to δ_i as $\Phi^* = M(\delta_{<\delta_1,\ldots,\delta^*_i, \ldots,\delta_n>}, \Omega)$.

Informally, suppose the priority-to-the-past decision process is path-*in*dependent for every profile in the domain **D**, and the aggregation function used satisfies the weak monotonicity condition just introduced. Then, for every individual *i*, the following holds. Regardless of the AR-functions expressed by the other individuals (where all resulting profiles are in **D**), if individual *i* expresses her truthful AR-function δ_i the outcome set produced by the decision process is at least as close to her most preferred outcome set as the one produced if she expresses any alternative, "untruthful" AR-function δ^*_i.

So, if all individuals' preferences satisfy Assumption 1, no individual has an incentive to express an untruthful AR-function at any profile in **D**. More strongly, in the domain **D**, expression of truthful AR-functions is a (weakly) dominant strategy for every individual. Thus, in **D**, expression of truthful AR-functions by all individuals is a (weakly-)dominant-strategy Nash equilibrium. In short, if a priority-to-the-past decision process $M(\delta, \Omega)$ is path-*in*dependent for every profile in the domain **D**, then it is strategy-proof in **D**.

Strategy-Proofness in D. For every individual $i \in N$ and every profile $<\delta_i> \in$ **D**, the following holds. For every AR-function δ^*_i ($\neq \delta_i$) (with $<\delta_1,\ldots,\delta^*_i, \ldots,\delta_n> \in$ **D**), $\Phi R_{i,\delta_i} \Phi^*$, where $\Phi = M(\delta_{<\delta_i>}, \Omega)$ and $\Phi^* = M(\delta_{<\delta_1, \ldots, \delta^*_i, \ldots, \delta_n>}, \Omega)$.

Corollary of Theorem 6. Suppose δ is a weakly monotonic aggregation function, and **D** is a domain of profiles. If $M(\delta, \Omega)$ is invariant under changes of the decision-path for every $<\delta_i> \in \mathbf{D}$, then $M(\delta, \Omega)$ is strategy-proof in **D**.

Although related to the Gibbard-Satterthwaite theorem, Theorem 6 and its Corollary make a subtly different point. The Gibbard-Satterthwaite theorem shows that certain minimal conditions on preference aggregation imply a *violation* of strategy-proofness.[28] My result shows that path-*in*dependence of a sequential decision process *implies* strategy-proofness. A question for future research is whether closer analogues of the Gibbard-Satterthwaite theorem can be established in the context of decisions on multiple propositions.

Despite the earlier impossibility results (Theorems 4 and 5), Theorem 6 and its Corollary are *not* vacuous. As discussed below, if the aggregation function is the unanimity rule or a dictatorship of one individual, the resulting priority-to-the-past decision process is always path-*in*dependent – and thus strategy-proof. If the aggregation function is majority voting, but the domain of admissible profiles is suitably restricted, then the resulting decision-process is also path-*in*dependent, and thus strategy-proof, in this restricted domain.

ESCAPE-ROUTES FROM PATH-DEPENDENCE AT A COLLECTIVE LEVEL

Suppose again that all individuals' views are complete, weakly consistent and deductively closed. By Theorems 4 and 5, there exists no anonymous or even just non-dictatorial propositionwise aggregation function such that a priority-to-the-past decision process is decisive and path-*in*dependent for every profile in the universal domain. To avoid path-dependence it is necessary to relax at least one of the conditions underlying this impossibility result: universal domain, non-dictatorship, decisiveness, or propositionwise aggregation. By Theorems 1 to 3, these conditions must be relaxed so as to allow the existence of an aggregation function that always generates a weakly consistent and deductively closed (hence strongly consistent) collective AR-function.

Relaxing Decisiveness: The Special Support Approach

This escape-route relaxes the requirement that the decision process should produce a determinate verdict on every proposition, and allows that, for some propositions, neither the proposition nor its negation are accepted. The group's propositional attitude on each proposition can then be defined by the unanimity rule or by a supermajority rule.

Definition. The *unanimity rule* is the aggregation function δ: $\{0,1\}^n \rightarrow \{0,1\}$ defined as follows. For any $(d_1, d_2, ..., d_n) \in \{0,1\}^n$,

$$\delta(d_1, d_2, ..., d_n) = \begin{cases} 1 \text{ if } d_i = 1 \text{ for every } i \in N \\ 0 \text{ otherwise.} \end{cases}$$

On this rule, for each proposition ϕ, the group has a disposition to accept ϕ if and only if *every* individual accepts ϕ. The group's propositional attitudes are then weakly consistent and deductively closed (hence strongly consistent), but not necessarily complete.[29] They are incomplete whenever there is a lack of unanimity on some propositions *and* their negations. The resulting priority-to-the-past decision process is invariant under changes of the decision-path, but not generally decisive: It fails to produce a verdict on those propositions which are neither unanimously accepted nor unanimously rejected by the individuals. The approach gives veto power to every individual, and is thus prone to "stalemate". Can a less demanding supermajority requirement be used instead?

Definition. *Supermajority voting* with parameter q is the aggregation function δ: $\{0,1\}^n \rightarrow \{0,1\}$ defined as follows. For any $(d_1, d_2, ..., d_n) \in \{0,1\}^n$,

$$\delta(d_1, d_2, ..., d_n) = \begin{cases} 1 \text{ if } \sum_{i \in N} d_i > qn \\ 0 \text{ otherwise.} \end{cases}$$

If $q = ½$, supermajority voting reduces to simple majority voting.

For a result on supermajority voting, note that the set of propositions X can be partitioned into $2m$ equivalence classes of logically equivalent propositions, where $m>1$. The number of equivalence classes is even because, as assumed, X always contains proposition-negation pairs. It is greater than 1 because, as assumed, X contains more than one such pair.

Theorem 7. Let δ be supermajority voting with parameter (m-1)/m (where there are 2m equivalence classes of logically equivalent propositions in X). Then δ satisfies anonymity and induces, for every $<\delta_i>\in \mathbf{U}$, a strongly consistent (but not generally complete and deductively closed) collective AR-function $\delta_{<\delta_i>}$.

Under supermajority voting with parameter (m-1)/m, the group has a disposition to accept each proposition ϕ if and only if a proportion of *more than* (m-1)/m of the individuals accept ϕ. The group's propositional attitudes are then strongly consistent, but not generally complete and deductively closed. They are incomplete whenever some propositions *and* their negations lack the required supermajority. They violate deductive closure whenever there is no supermajority on some *implications* of propositions that each get a supermajority. The resulting priority-to-the-past decision process is not strongly path-dependent (as the group's propositional attitudes are strongly consistent). It may be weakly path-dependent (as the propositional attitudes may violate deductive closure), and it may be indecisive (as they may violate completeness). The approach does not give veto power to every individual,[30] but it gives veto power to every *group* of a proportion of 1/m or more of the individuals. It may thus still be prone to "stalemate", albeit less so than the unanimity approach.

Theorem 7 mirrors results by Craven (1971) and Ferejohn and Grether (1974) in the context of preferences showing that suitable supermajority rules can generate acyclical collective preferences; it is not a direct corollary of these results, because of the differences between views on multiple interconnected propositions on the one hand and preferences on the other, and between the associated consistency conditions.[31]

Relaxing Non-Dictatorship: The Dictatorship Approach

This escape-route not only relaxes the requirement that all individuals should have equal weight in determining the group's propositional attitudes, but also allows the existence of a

dictator. By Theorem 5, a dictatorship of one individual is the unique aggregation function that guarantees decisiveness and path-*in*dependence for all profiles in the universal domain (where the dictator's views are complete, weakly consistent and deductively closed). Under a dictatorship, the outcome set of a priority-to-the-past decision process is the set of propositions accepted by the dictator and will typically fail to reflect the views of any other individuals. Here path-*in*dependence is achieved at the cost of democratic responsiveness.

Relaxing Universal Domain: The Domain Restriction Approach

This escape-route relaxes the requirement that the decision process should accept as admissible input any logically possible profile of individual AR-functions. Suppose that not all such profiles will occur in collective decisions, for one of the following reasons: (i) Coercive reasons: Certain profiles are explicitly ruled out by restrictions on the views individuals are permitted to express. (ii) Empirical reasons: Certain profiles do not occur in practice, as a matter of contingent fact. (iii) Deliberative reasons: While all logically possible profiles can in principle occur, the individuals engage in group deliberation prior to a collective decision, where such group deliberation transforms the individual views and thereby restricts the domain of post-deliberation profiles.

Parallel mechanisms of domain restriction have been discussed in relation to classical models of social choice (Miller 1992; Knight and Johnson 1994; Dryzek and List 2003). I here remain neutral on how feasible such mechanisms are. I identify a structure condition with the property that, if the domain of admissible profiles includes only ones satisfying that condition, then a priority-to-the-past decision process is path-*in*dependent. The condition – called *unidimensional alignment* – is similar in spirit, but not in form, to Black's (1948) condition of single-peakedness in the context of preferences. A profile satisfies *unidimensional alignment* if there exists a single ordering of all individuals from left to right – a *structuring ordering* – such that, for *every* proposition in X, the individuals accepting the

proposition are either all to the left, or all to the right, of those rejecting it. The profile shown in Table 2 satisfies unidimensional alignment. The corresponding structuring ordering is 4, 1, 5, 2, 3. Individual 5 is the *median individual*, as defined below.

[Table 2 about here]

For a formal definition, fix a profile $<\delta_i>$. For each $\phi \in X$, let $N_{accept\text{-}\phi}$ be the set of individuals accepting ϕ, and $N_{reject\text{-}\phi}$ the set of individuals not accepting ϕ, i.e. $N_{accept\text{-}\phi} := \{i \in N : \delta_i(\phi)=1\}$ and $N_{reject\text{-}\phi} := \{i \in N : \delta_i(\phi)=0\}$. A *structuring ordering* is a linear ordering ω on the set of individuals N, where $i\omega j$ means that individual i is to the left of individual j. For any two subsets $N_1, N_2 \subseteq N$, I write $N_1 \omega N_2$ as an abbreviation for [for all $i \in N_1$ and all $j \in N_2$, $i\omega j$], i.e. the individuals in N_1 are all to the left of those in N_2.

Definition. (List 2003) A profile $<\delta_i>$ satisfies *unidimensional alignment* if there exists a structuring ordering ω such that, for every $\phi \in X$, either $N_{accept\text{-}\phi} \omega N_{reject\text{-}\phi}$ or $N_{reject\text{-}\phi} \omega N_{accept\text{-}\phi}$.[32]

Definition. The *unidimensional alignment domain*, **R**, is the set of all logically possible profiles of individual AR-functions satisfying completeness, weak consistency, deductive closure and unidimensional alignment.

Suppose a profile $<\delta_i>$ satisfies unidimensional alignment. Order the n individuals along a corresponding structuring ordering, say ω. There are two cases:

***n* is odd.** Individual m is the *median individual* with respect to ω if there are as many individuals to m's left as to m's right, formally $|\{i \in N : i\omega m\}| = |\{i \in N : m\omega i\}|$. By unidimensional alignment, the median individual shares the majority view on every proposition. Hence the collective AR-function induced by majority voting is the AR-function of the median individual, δ_m. Provided the median individual's AR-function is complete, weakly consistent and deductively closed, so is the collective AR-function.

***n* is even.** Here there exists no single median *individual*, but a median *pair*. Individuals m_1 and m_2 are the *median pair* with respect to ω if they are adjacent (with m_1 left of m_2) and

there are as many individuals to m_1's left as to m_2's right, formally (i) $m_1 \omega m_2$, (ii) there is no $i \in N$ such that $m_1 \omega i$ and $i \omega m_2$, and (iii) $|\{i \in N: i \omega m_1\}| = |\{i \in N: m_2 \omega i\}|$. By unidimensional alignment, if (and only if) the median *pair* of individuals agree on a proposition, then their view is also the majority view on that proposition. So the collective AR-function induced by majority voting is the product (the "intersection") of the AR-functions of the median pair of individuals, $\delta_{m_1} \delta_{m_2}$.[33] Provided the AR-functions of the median pair are each complete, weakly consistent and deductively closed, the collective AR-function is also weakly consistent and deductively closed. It may violate completeness, namely when the median pair of individuals disagree on a proposition ϕ.[34] But this happens *only if* there is a majority tie, i.e. only if $|N_{accept\text{-}\phi}| = |N_{reject\text{-}\phi}|$. Thus the collective AR-function is "almost" complete:

Definition. Given a profile $<\delta_i>$, $\delta_{<\delta_i>}$ is *almost complete* if, for any $\phi \in X$, $|N_{accept\text{-}\phi}| \neq |N_{reject\text{-}\phi}|$ implies that $\delta_{<\delta_i>}(\phi)=1$ or $\delta_{<\delta_i>}(\neg\phi)=1$.

Proposition 6. (List 2003) Let δ be propositionwise majority voting. Then δ satisfies anonymity and induces, for every $<\delta_i> \in \mathbf{R}$ (with structuring ordering ω), an almost complete, weakly consistent and deductively closed collective AR-function $\delta_{<\delta_i>}$, where:

(i) If n is odd, $\delta_{<\delta_i>} = \delta_m$, where m is the median individual with respect to ω.

(ii) If n is even, $\delta_{<\delta_i>} = \delta_{m_1} \delta_{m_2}$, where m_1 and m_2 are the median pair with respect to ω.

This result mirrors Black's median voter theorem (1948), albeit in a new context; it is not a corollary of Black's result, as unidimensional alignment and single-peakedness are formally different (for a more general domain restriction result in a preference context, see Sen 1966). By Proposition 6, for any profile satisfying unidimensional alignment, the group's propositional attitudes, determined by majority voting, are almost complete, weakly consistent and deductively closed. I next show that the resulting priority-to-the-past decision process is path-*in*dependent. It is not generally decisive if n is even, but "almost" decisive:

Definition. Given a profile $<\delta_i>$, $M(\delta_{<\delta_i>}, \Omega)$ is *almost decisive* if, for every $\phi \in X$, $|N_{accept\text{-}\phi}| \neq |N_{reject\text{-}\phi}|$ implies that $\phi \in M(\delta_{<\delta_i>}, \Omega)$ or $\neg\phi \in M(\delta_{<\delta_i>}, \Omega)$.

Theorem 8. Let δ be propositionwise majority voting. Then δ satisfies anonymity and, for any $<\delta_i> \in \mathbf{R}$ (with structuring ordering ω), $M(\delta_{<\delta_i>}, \Omega)$ is almost decisive and invariant under changes of the decision-path, where:

(i) If n is odd, $M(\delta_{<\delta_i>}, \Omega) = \{\phi \in X : \delta_m(\phi)=1\}$, where m is the median individual with respect to ω.

(ii) If n is even, $M(\delta_{<\delta_i>}, \Omega) = \{\phi \in X : \delta_{m_1}(\phi)\delta_{m_2}(\phi)=1\}$, where m_1 and m_2 are the median pair with respect to ω.

So if the group uses majority voting as its aggregation function, then the resulting priority-to-the-past decision process is almost decisive and path-*in*dependent for every profile in the restricted domain **R**. The outcome set of the process is the set of propositions accepted by the median individual (if n is odd) or the intersection of the sets of propositions accepted by the median *pair* of individuals (if n is even).

Is unidimensional alignment just an artificial condition, or can it be met in plausible situations? Suppose, first, different individuals disagree on the propositions, but they reach some "meta"-agreement on what their disagreement is about: They agree on a single dimension (e.g. from "most liberal" to "most conservative") on which their positions can all be placed; each individual takes a position on that dimension. I call it a left-right dimension, but different interpretations are possible. Suppose further that, for each proposition, the extreme positions on the left-right dimension correspond to either clear acceptance or clear rejection of the proposition, and there exists an "acceptance threshold" on the dimension (possibly different for different propositions) such that all individuals to the threshold's left accept the proposition and all individuals to its right reject it (or vice-versa). These

conditions entail unidimensional alignment; i.e. "meta"-agreement of the form described can induce unidimensional alignment. This suggests that, to the extent that such "meta"-agreement is feasible, the escape-route from path-dependence opened up by unidimensional alignment may have some promise. But, for my purposes here, the specific interpretation of unidimensional alignment is less relevant than the general insight that path-dependence can be avoided if there is a sufficient level of structure among the individuals' views.

Relaxing Propositionwise Aggregation: The Generalized Aggregation Approach

In my model the group's initial dispositions are defined by *propositionwise* aggregation, in analogy to Arrow's independence requirement. The aggregation function maps each vector of 0s and 1s across individuals to a corresponding single collective disposition of either 0 or 1. Although plausible in sequential decisions where separate votes are taken on each proposition, this is a restrictive definition of aggregation.[35] A *generalized aggregation function*, F, maps each profile $<\delta_i>$ to a collective AR-function $\delta: X \rightarrow \{0,1\}$. Such a generalized aggregation function F provides an escape-route from path-dependence if and only if it generates a weakly consistent and deductively closed AR-function δ.

A function F with this property can be defined by "feeding" each profile $<\delta_i>$ into an appropriate sequential decision process. For each profile $<\delta_i>$, define $F(<\delta_i>) := \delta_F$ by:

$$\text{for each } \phi \in X,\ \delta_F(\phi) = \begin{cases} 1 \text{ if } \phi \in M(\delta_{<\delta_i>}, \Omega_F) \\ 0 \text{ otherwise,} \end{cases}$$

where $\delta_{<\delta_i>}$ is the result of applying propositionwise majority voting to $<\delta_i>$ and Ω_F is some fixed decision-path. By definition, δ_F is (almost) complete, weakly consistent and deductively closed for any profile $<\delta_i> \in \mathbf{U}$. Suppose now that δ_F is itself used as an AR-function in a sequential decision process. The weak consistency and deductive closure of δ_F implies that $M(\delta_F, \Omega)$ is path-*in*dependent. In this *technical* sense, the generalized aggregation function F provides an escape-route from path-dependence.

Some aggregation functions that have been proposed for solving the "doctrinal paradox" are of this form. The "premise-based procedure", for example, applies majority voting to the premises and decides other propositions based on what the votes on the premises imply (Pettit 2001a; Chapman 2002). This can be represented as a generalized aggregation function F, defined in terms of a decision-path Ω_F where the premises occur before the conclusion.

Now the decision process $M(\delta_F, \Omega)$ is path-*in*dependent *after* δ_F has been determined in the manner outlined. But the path-dependence problem has not been resolved; it has been shifted one level up. The path-*in*dependence of $M(\delta_F, \Omega)$ has been achieved at the cost of making the aggregation function F used for defining δ_F itself "internally" path-dependent. As F is defined by a sequential decision process – i.e. $\delta_F(\phi) = 1$ if and only if $\phi \in M(\delta_{<\delta_i>}, \Omega_F)$ – the collective AR-function δ_F is not generally invariant *under changes of the path* Ω_F *used for defining* F. All the problems I have discussed – particularly strategic manipulability – will therefore recur at the level of the aggregation function F.

So, if the aim is not just to shift path-dependence one level up, the escape-route via relaxing propositionwise aggregation seems not very promising. An open challenge is to find a generalized aggregation function that *both* relaxes the assumption of propositionwise aggregation *and* is not itself "internally" path-dependent in any problematic way.

CONCLUDING REMARKS

I have modeled sequential decisions over multiple interconnected propositions and illustrated my model by investigating path-dependence in such decisions. The model is intended to complement, not replace, existing social-choice-theoretic models. It seeks to capture a class of decisions that are not straightforwardly captured by classical models: reason-based decisions, where the focus is not only on outcomes, but also on underlying reasons, beliefs and constraints. Let me briefly summarize my main substantive results.

I have shown that certain violations of perfect rationality by the relevant agent are necessary and sufficient for path-dependence. I have discussed the implications of this result for both individuals and groups. While a perfectly rational individual can avoid path-dependence, a boundedly rational individual may be susceptible to it. An example is an agent who never accepts a proposition and its negation simultaneously, but who fails to foresee all the logical implications of the propositions she accepts. As path-dependence makes the agent manipulable by the presentation of the propositions in some strategic order, this finding suggests a possible link between bounded rationality and Rikerian "heresthetics", the art of political manipulation. Exploring this link may be an interesting avenue for future work.

Path-dependence is particularly serious at the collective level. While individuals might try to avoid path-dependence through a self-imposed "discipline" of rationality, no such option is easily available to groups. Under certain conditions, any group that determines its propositional attitudes by aggregation over its members' views will necessarily run the risk of those rationality violations that lead to path-dependence. Path-dependence makes collective decisions vulnerable to manipulation both by agenda setting and by expression of untruthful views. I have identified some formal escape-routes from path-dependence at the collective level, and future research might ask how substantively promising these routes are.

My results are also relevant to the debates on collective consistency referred to above. As noted, social choice theorists have traditionally been skeptical towards the feasibility of a consistent collective set of views. Related to the classical results on collective inconsistencies – usually defined as voting cycles – there are also several results suggesting that collective outcomes are inherently unstable. This, in turn, can be strategically exploited. Schwartz's universal instability theorem (1981), for instance, shows that, under mild assumptions, any collective outcome involving certain generalized exchange among agents is unstable: Some group of agents has the preference and power to overturn that outcome in favor of some other

outcome. Further, some results show that, even when equilibria exist in politics, these may not be unique, which raises the problem of equilibrium selection. The Muller-Satterthwaite theorem (Muller and Satterthwaite 1977; Myerson 1996), for instance, suggests that, if we interpret a political system as a voting game and there are three or more possible outcomes, then, absent a dictator, the game does not generally have a unique equilibrium (if any).

Are my results yet another addition to the large set of social-choice-theoretic inconsistency and instability results? Not quite. Given a decision-path, the outcomes of a priority-to-the-past decision process in my model are neither inconsistent nor unstable. They are *consistent* as the priority-to-the-past rule does not permit the acceptance of propositions that conflict with propositions accepted earlier. They are *stable* as propositions, once accepted, are not overruled by the priority-to-the-past rule; moreover, each separate decision along the path is itself binary. The constraints implemented by the priority-to-the-past rule may thus reduce the continual drift of collective outcomes that classical models predict in unconstrained settings. All this comes at a price – path-dependence: A group can generate consistent collective outcomes in a priority-to-the-past decision process; it can do so in a relatively democratic manner, by using an aggregation function such as majority voting; *but* the decision outcomes may be path-dependent and thus affected by the problems I have identified, such as strategic manipulation.

One might speculate whether my results suggest another possible answer to Tullock's question "Why so much stability?" (1981). Perhaps political outcomes are stabilized not only by constraints created by institutional arrangements – as in a structure-induced equilibrium (Shepsle 1979) – but also by constraints created by prior decisions or commitments – as in a priority-to-the-past decision process. Based on my research here, however, it would be premature to propose a new equilibrium concept, but there are interesting avenues for future research. For example, see Page's work on path-dependent equilibria (2003).

How threatening is path-dependence? Some factors might make it less of a threat. A decision's subject-matter might itself single out the appropriate decision-path. Some propositions might be unambiguously "weightier than", or "prior to", others, and the order in which the propositions are to be considered might be uncontroversial. Alternatively, history or empirical contingencies might determine the decision-path, with little opportunities for intervention by agents with a strategic interest. But even then it is worth asking whether or not the given path makes a difference to the outcome. If there is no path-dependence, the legitimacy of an outcome will be under no threat, whether or not the decision-path is disputed; that path is simply irrelevant. But, if there is path-dependence, a justification of the chosen path is crucial. Further, in such cases, even if agents agree on a decision-path, this will not rule out manipulation by the expression of untruthful views. As shown, the mere existence of an alternative decision-path that would change the outcome – although that path is not adopted – may create incentives for strategic expression of untruthful views. So, curiously, path-dependence may matter even when the decision-path is fixed.

I have already mentioned several questions for further research. Others include how far the present results can be generalized to conflict resolution rules other than the priority-to-the-past rule; under what restrictions on the set of admissible, or naturally occurring, decision-paths there is no path-dependence; and whether any domain restrictions such as unidimensional alignment are empirically realistic.

Improving our understanding of path-dependence is an important challenge in the theory of democracy. Many democratic decision processes are sequential, and hence it is important to learn whether, and how, the decision-path matters, and what the implications of path-dependence are.

APPENDIX

Proof of Lemma 1. Let (i) and (ii) denote the left and right sides of the biconditional.

(i) implies (ii): Suppose (i) holds. Then $\Phi:=\{\phi\in X : \delta(\phi)=1\}$ is logically inconsistent. Let Ψ_2 be a maximal consistent subset of Φ. Then $\Psi_2 \neq \varnothing$, as X contains no contradictions; $\Psi_2 \neq \Phi$, as Φ is not consistent. Choose any $\phi\in \Phi\setminus\Psi_2$. As Ψ_2 is a *maximal* consistent subset of Φ, $\Psi_2\cup\{\phi\}$ is not logically consistent; so $\Psi_2 \models \neg\phi$. Let $\Psi_1 := \{\phi\}$; Ψ_1 is logically consistent, as ϕ is not a contradiction. Then Ψ_1 and Ψ_2 have the properties required by (ii).

(ii) implies (i): Suppose (ii) holds. As $\Psi_1 \models \phi$ and $\Psi_2 \models \neg\phi$, the set $\Psi_1\cup\Psi_2$ is logically inconsistent. But $\Phi\supseteq\Psi_1\cup\Psi_2$. So Φ is also logically inconsistent, and (i) holds. ■

Lemma 2. For any $\phi\in X$, (i) there exists a decision-path Ω such that $\phi\in M(\delta, \Omega)$ if and only if (ii) there exists a logically consistent subset $\Psi\subseteq X$ such that $\delta(\Psi)=1$ and $\Psi \models \phi$.

Proof of Lemma 2. (i) implies (ii): Suppose (i) holds. Let Ω be a decision-path such that $\phi\in M(\delta, \Omega)$. Choose t such that ϕ is accepted at time t in $M(\delta, \Omega)$. Let $\Psi = \{\psi\in X : \delta(\psi)=1$ *and* ψ *is accepted at some time* $s<t$ *in* $M(\delta, \Omega)\}$. As ϕ is accepted at time t, *either* $\delta(\phi)=1$ *or* $\Psi\models\phi$. If $\delta(\phi)=1$, then $\{\phi\}$ has the properties required by (ii): $\{\phi\}$ is logically consistent, as ϕ is not a contradiction. If $\Psi\models\phi$, then Ψ has the properties required by (ii): Ψ is logically consistent, as $\Psi\subseteq M(\delta, \Omega)$, which is logically consistent.

(ii) implies (i): Suppose (ii) holds. Define Ω as follows. Let $t = |\Psi\cup\{\phi\}|$. On $\{1, 2, \ldots, t\}$, let Ω be any one-to-one mapping from $\{1, 2, \ldots, t\}$ onto $\Psi\cup\{\phi\}$ such that $\Omega(t) = \phi$. To continue the path Ω on $\{t+1, ..., k\}$, let Ω be any one-to-one mapping from $\{t+1, ..., k\}$ onto $X\setminus(\Psi\cup\{\phi\})$, where $k = |X|$. Then Ω has the properties required by (i). ■

Proof of Theorem 1. Theorem 1 follows immediately from Lemmas 1 and 2. ■

Lemma 3. Suppose $M(\delta, \Omega)$ is invariant under changes of Ω. Then, for every $\phi\in X$ (and any decision-path Ω), $\phi\in M(\delta, \Omega)$ if and only if $\delta(\phi)=1$.

Proof of Lemma 3. Let $\phi \in X$. Suppose $M(\delta, \Omega)$ is path-*in*dependent. We can use any decision-path Ω. For each $\phi \in X$, define Ω_ϕ such that $\Omega_\phi(1):=\phi$, and $\Omega_\phi(2):=\neg\phi$; on $\{3, \ldots, k\}$, let Ω_ϕ be any one-to-one mapping from $\{3, \ldots, k\}$ onto $X\backslash\{\phi, \neg\phi\}$, where $k = |X|$. Suppose $\delta(\phi)=1$. Then, by the definition of $M(\delta, \Omega_\phi)$, $\phi \in M(\delta, \Omega_\phi)$. Suppose, conversely, $\phi \in M(\delta, \Omega_\phi)$, but $\delta(\phi)=0$. As Φ_0 does not entail ϕ, $\phi \notin \Phi_1$ and $\phi \notin \Phi_2$. But ϕ or $\neg\phi$ do not occur elsewhere in Ω_ϕ, so $\phi \notin M(\delta, \Omega_\phi)$, a contradiction. Hence $\delta(\phi)=1$, as required. ■

Proof of Theorem 3. Let (i) and (ii) denote the left and right sides of the biconditional. Suppose $\delta: X \rightarrow \{0,1\}$ is strongly consistent.

(i) implies (ii): Suppose there exist Ω_1 and Ω_2 such that $\phi \in M(\delta, \Omega_1)$ and $\phi \notin M(\delta, \Omega_2)$, but δ is deductively closed with respect to ϕ. As δ is strongly consistent, $M(\delta, \Omega)$ is weakly but not strongly path-dependent. As ϕ and Ω_1 satisfy (i) in Lemma 2, there exists a logically consistent $\Phi \subseteq X$ such that $\delta(\Phi)=1$ and $\Phi \models \phi$. Then $\delta(\phi)=1$, by deductive closure of δ. Choose t such that $\Omega_2(t) = \phi$. Under path Ω_2, at time t, ϕ is *not* accepted, as $\phi \notin M(\delta, \Omega_2)$. As $\delta(\phi)=1$, this requires $\Phi_{t-1} \models \neg\phi$. By the definition of $M(\delta, \Omega_2)$, $\neg\phi \in M(\delta, \Omega_2)$; so $M(\delta, \Omega)$ is strongly path-dependent, a contradiction. Hence δ is not deductively closed with respect to ϕ.

(ii) implies (i): Suppose δ violates deductive closure with respect to $\phi \in X$, but $M(\delta, \Omega)$ is path-*in*dependent with respect to ϕ. There exist a logically consistent $\Phi \subseteq X$ and $\phi \in X$ such that $\delta(\Phi)=1$, $\Phi \models \phi$ but $\delta(\phi)=0$. By Lemma 2, there exists Ω such that $\phi \in M(\delta, \Omega)$. As $M(\delta, \Omega)$ is path-*in*dependent, $\phi \in M(\delta, \Omega)$ for *every* Ω. Consider Ω_ϕ as in the proof of Lemma 3. As $\delta(\phi)=0$ and not $\Phi_0 \models \phi$, $\phi \notin \Phi_1$ and $\phi \notin \Phi_2$ for Ω_ϕ. But ϕ or $\neg\phi$ do not occur elsewhere in Ω_ϕ, so $\phi \notin M(\delta, \Omega_\phi)$, contradicting $\phi \in M(\delta, \Omega)$ for *every* Ω. Hence $M(\delta, \Omega)$ is weakly path-dependent with respect to ϕ. ■

Proof of Theorems 4 and 5. Let (i) and (ii) denote the left and right sides of the biconditional in Theorem 5. Theorem 4 follows from the proof that (i) implies (ii).

(i) implies (ii): Let δ be an aggregation function such that, for every $<\delta_i>\in \mathbf{U}$, $M(\delta_{<\delta_i>}, \Omega)$ is decisive and path-*in*dependent. By Lemma 3, for every $\phi\in X$, $\phi\in M(\delta_{<\delta_i>}, \Omega)$ if and only if $\delta_{<\delta_i>}(\phi)=1$. As $M(\delta_{<\delta_i>}, \Omega)$ is decisive and logically consistent, $\delta_{<\delta_i>}$ is complete and strongly consistent, hence weakly consistent and deductively closed. So δ induces, for every $<\delta_i>\in \mathbf{U}$, a complete, weakly consistent and deductively closed $\delta_{<\delta_i>}$. By Proposition 4, δ violates anonymity. More strongly, by Proposition 5, δ is a dictatorship of some $i\in N$.

(ii) implies (i): Let δ be a dictatorship of some $i\in N$. Take any $<\delta_i>\in \mathbf{U}$. Then $\delta_{<\delta_i>} = \delta_i$. Since δ_i is complete, weakly consistent and deductively closed (hence strongly consistent), so is $\delta_{<\delta_i>}$. By Theorems 2 and 3, $M(\delta_{<\delta_i>}, \Omega)$ is invariant under changes of the decision-path. The decisiveness of $M(\delta_{<\delta_i>}, \Omega)$ follows from the completeness of $\delta_{<\delta_i>}$. ■

Proof of Theorem 6. Suppose δ is weakly monotonic, and $M(\delta, \Omega)$ is path-*in*dependent for every $<\delta_i>\in \mathbf{D}$. Take any $i\in N$, any $<\delta_i>\in \mathbf{D}$, and any $\delta^*_i \neq \delta_i$ (with $<\delta_1,\ldots,\delta^*_i, \ldots,\delta_n>\in \mathbf{D}$). I prove that $\Phi = M(\delta_{<\delta_i>}, \Omega)$ is at least as close to δ_i as $\Phi^* = M(\delta_{<\delta_1,\ldots,\delta^*_i,\ldots,\delta_n>}, \Omega)$. By path-*in*dependence of $M(\delta, \Omega)$ in $\mathbf{D}$, Lemma 3 implies that, for every $\phi\in X$, [$\phi\in M(\delta, \Omega)$ if and only if $\delta(\phi)=1$] and [$\phi\in M(\delta^*, \Omega)$ if and only if $\delta^*(\phi)=1$], where $\delta := \delta_{<\delta_i>}$ and $\delta^* := \delta_{<\delta_1,\ldots,\delta^*_i,\ldots,\delta_n>}$. I must show that, for every $\phi\in X$, $\delta_i(\phi)=\delta^*(\phi)$ implies $\delta_i(\phi)=\delta(\phi)$, i.e. $|\delta(\phi)-\delta_i(\phi)| \leq |\delta^*(\phi)-\delta_i(\phi)|$. Take any $\phi\in X$. Let $d_1:=\delta_1(\phi)$, $d_2:=\delta_2(\phi)$, …, $d_n:=\delta_n(\phi)$, and $d^*_i:=\delta^*_i(\phi)$. Then $\delta(\phi) = \delta(\delta_1(\phi),\ldots,\delta_i(\phi),\ldots,\delta_n(\phi)) = \delta(d_1,\ldots,d_i,\ldots,d_n)$ and $\delta^*(\phi) = \delta(\delta_1(\phi),\ldots,\delta^*_i(\phi),\ldots,\delta_n(\phi)) = \delta(d_1, \ldots,d^*_i,\ldots, d_n)$. The weak monotonicity of δ implies that $|\delta(d_1, \ldots, d_i, \ldots, d_n) - d_i| \leq |\delta(d_1, \ldots, d^*_i, \ldots, d_n) - d_i|$, and thus $|\delta(\phi) - \delta_i(\phi)| \leq |\delta^*(\phi) - \delta_i(\phi)|$, as required. ■

Proof of Corollary of Theorem 6. Suppose δ is weakly monotonic, and $M(\delta, \Omega)$ is path-*in*dependent for all $<\delta_i>\in \mathbf{D}$. Take any $i\in N$, any $<\delta_i>\in \mathbf{D}$, and any δ^*_i ($\neq\delta_i$) (with $<\delta_1,\ldots,\delta^*_i,\ldots,\delta_n>\in \mathbf{D}$). I show $\Phi R_{i,\delta_i}\Phi^*$, where $\Phi=M(\delta_{<\delta_i>},\Omega)$ and $\Phi^*=M(\delta_{<\delta_1, \ldots,\delta^*_i, \ldots, \delta_n>}, \Omega)$. By Theorem 6, Φ is at least as close to δ_i as Φ^*. By Assumption 1, $\Phi R_{i,\delta_i}\Phi^*$. ■

Proof of Theorem 7. Assume Theorem 7's conditions. Define $N_{accept\text{-}\phi}$ and $N_{reject\text{-}\phi}$ as in the section on domain restriction. Let δ be supermajority voting with parameter $(m\text{-}1)/m$; δ satisfies anonymity. Take any $<\delta_i>\in\mathbf{U}$. Let $\Phi:=\{\phi\in X : \delta_{<\delta_i>}(\phi)=1\}$ be the set of propositions accepted under the aggregation function δ for $<\delta_i>$. I show that, for some $i\in N$, $\Phi\subseteq\{\phi\in X : \delta_i(\phi)=1\}$. As each δ_i is weakly consistent, there is no $\phi\in X$ such that ϕ and $\neg\phi$ both obtain $(m\text{-}1)/m$ supermajorities; so $\delta_{<\delta_i>}$ is weakly consistent. Partition Φ into m^* equivalence classes of logically equivalent propositions. As $\delta_{<\delta_i>}$ is weakly consistent, Φ contains at most one member of each proposition-negation pair, so $m^*\leq m\leq 2m$. Let $\phi_1, \phi_2, ..., \phi_{m^*}\in\Phi$ be representatives for these equivalence classes. Now $|N_{accept\text{-}\phi_1}|, |N_{accept\text{-}\phi_2}|, ..., |N_{accept\text{-}\phi_{m^*}}| > n(m\text{-}1)/m$. Then $|N_{accept\text{-}\phi_1}\cap N_{accept\text{-}\phi_2}| > n(m\text{-}1)/m\text{-}n/m = n(m\text{-}2)/m$. Also, $|N_{accept\text{-}\phi_1}\cap N_{accept\text{-}\phi_2}\cap N_{accept\text{-}\phi_3}| > n(m\text{-}2)/m - n/m = n(m\text{-}3)/m$. Continuing, $|N_{accept\text{-}\phi_1}\cap N_{accept\text{-}\phi_2}\cap...\cap N_{accept\text{-}\phi_{m^*}}| > n(m\text{-}m^*)/m$. Since $m^*\leq m$, $n(m\text{-}m^*)/m\geq 0$, hence $|N_{accept\text{-}\phi_1}\cap N_{accept\text{-}\phi_2}\cap...\cap N_{accept\text{-}\phi_{m^*}}| > 0$. So for some $i\in N$, $i\in N_{accept\text{-}\phi_1}\cap N_{accept\text{-}\phi_2}\cap...\cap N_{accept\text{-}\phi_{m^*}}$. But every $\phi\in\Phi$ is equivalent to one of ϕ_1, ϕ_2, ..., ϕ_{m^*}. As δ_i is complete, weakly consistent and deductively closed, it follows that, for every $\phi\in\Phi$, $i\in N_{accept\text{-}\phi}$; hence $\Phi\subseteq\{\phi\in X:\delta_i(\phi)=1\}$. As δ_i is strongly consistent, so is $\delta_{<\delta_i>}$. ■

Proof of Theorem 8. Let δ be majority voting. By Proposition 6, δ satisfies anonymity and, for any $<\delta_i>\in\mathbf{R}$, $\delta_{<\delta_i>}$ is almost complete, weakly consistent and deductively closed, thus strongly consistent. Take any $<\delta_i>\in\mathbf{R}$. By Theorem 3, $M(\delta_{<\delta_i>}, \Omega)$ is path-*in*dependent. By Lemma 3, $M(\delta_{<\delta_i>}, \Omega)=\{\phi\in X : \delta_{<\delta_i>}(\phi) = 1\}$. The almost-completeness of $\delta_{<\delta_i>}$ implies the almost-decisiveness of $M(\delta_{<\delta_i>}, \Omega)$; (i) and (ii) follow from (i) and (ii) in Proposition 6. ■

REFERENCES

Aragones E., I. Gilboa, A. Postlewaite, and D. Schmeidler. 2002. "Rhetoric and Analogies." Working paper, Department of Economics, Tel-Aviv University.

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

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

Brennan, G. 2001. "Collective Coherence?" *International Review of Law and Economics* 21 (2, June): 197-211.

Chapman, B. 1998. "More Easily Done than Said: Rules, Reason and Rational Social Choice." *Oxford Journal of Legal Studies* 18 (2, Summer): 293-330.

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

Coleman, J., and J. Ferejohn. 1986. "Democracy and Social Choice." *Ethics* 97 (1,Oct.): 6-25.

Craven, J. 1971. "Majority Voting and Social Choice." *Review of Economic Studies* 38 (2, April): 265-267.

Dietrich, F. 2003. "Judgment aggregation: (im)possibility theorems." Working paper, Group on Philosophy, Probability and Modeling, University of Konstanz.

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

Dworkin, R. M. 1986. *Law's empire*. London: Fontana.

Ferejohn, J. and D. Grether. 1974. "On a Class of Rational Social Decision Procedures." *Journal of Economic Theory* 8 (4, August): 471-482.

Gibbard, A. 1973. "Manipulation of Voting Schemes: A General Result." *Econometrica* 41 (4, July): 587-601.

Guilbaud, G. Th. 1966. “Theories of the General Interest, and the Logical Problem of Aggregation.” In Lazarsfeld, P. F., and N. W. Henry, eds. *Readings in Mathematical Social Science*. Cambridge/MA: MIT Press: 262-307.

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

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

Kornhauser, L.A. 1992. “Modeling Collegial Courts I: Path-Dependence.” *International Review of Law and Economics* 12 (2, June): 169-185.

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

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. 2003. “Aggregating Sets of Judgments: Two Impossibility Results Compared.” *Synthese*, forthcoming.

McLean, I. 2002. “William H. Riker and the Invention of Heresthetic(s).” *British Journal of Political Science* 32 (3, July): 535-558.

Miller, G., and N. Schofield. 2003. “Activists and Partisan Realignment in the United States.” *American Political Science Review* 97 (2, April): 245-260.

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

Muller, E., and M. Satterthwaite. 1977. “The Equivalence of Strong Positive Association and Strategy-Proofness.” *Journal of Economic Theory* 14 (April): 412-418.

Myerson, R. B. 1996. “Fundamentals of Social Choice Theory.” Working paper, CMS-EMS, Northwestern University.

Murakami, Y. 1968. *Logic and Social Choice*. London: Routledge.

Nehring, K., and C. Puppe. 2002. "Strategy-Proof Social Choice on Single-Peaked Domains." Working paper, Department of Economics, University of California at Davis.

Page, S. E. 1997. "An appending efficient algorithm for allocating public projects with positive complementarities.", *Journal of Public Economics* 64 (3, June): 291-444.

Page, S. E. 2003. "A Path Dependence Primer." Working paper, Department of Political Science, University of Michigan, Ann Arbor.

Pauly, M., and M. van Hees. 2003. "Some General Results on the Aggregation of Individual Judgments." Working paper, Department of Computer Science, University of Liverpool.

Pettit, P. 1997. *Republicanism: A Theory of Freedom and Government*. Oxford: O.U.P.

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

Pettit, P. 2001b. *A Theory of Freedom*. New York: O.U.P.

Plott, C. 1973. "Path Independence, Rationality, and Social Choice." *Econometrica* 41 (6, November): 1075-1091.

Plott, C., and M. Levine. 1978. "A Model of Agenda Influence on Committee Decisions." *American Economics Review* 68 (1, March): 146-160.

Riker, W. H. 1982. *Liberalism Against Populism*. San Francisco: W. H. Freeman.

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

Rott, H. 2001. *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford: O.U.P.

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 (April): 187-217.

Schwartz, T. 1981. "The universal-instability theorem." *Public Choice* 37 (3): 487-501.

Sen, A. K. 1966. “A Possibility Theorem on Majority Decisions.” *Econometrica* 34 (2, April): 491-499.

Sen, A. K. 1970. *Collective choice and social welfare*. Amsterdam: North-Holland.

Shepsle, K. A. 1979. “Institutional Arrangements and Equilibrium in Multidimensional Voting Models.” *American Journal of Political Science* 23 (1, February): 27-59.

Stone Sweet, A. 2002. “Path Dependence, Precedent, and Judicial Power.” In Shapiro, M., and A. Stone Sweet. *On Law, Politics, and Judicialization*. Oxford: O.U.P.: 112-135.

Taylor, A. D. 2002. “The Manipulability of Voting Systems.“ *American Mathematical Monthly* 109 (April): 321-337.

Taylor, A. D. 2004. *Social Choice and the Mathematics of Manipulation*. Cambridge: C.U.P., forthcoming.

Tullock, G. 1981. “Why so much stability.” *Public Choice* 37 (2): 189-205.

NOTES

[1] Beliefs enter such models either not at all, or only as probabilities individuals assign to different states of the world. Regarding constraints between choices, only minimal rationality conditions are imposed on individual and collective preferences. On deliberation and aggregation, see Miller (1992), Knight and Johnson (1994) and Dryzek and List (2003).

[2] Nehring and Puppe (2002), discussed below, have also modeled decisions on multiple propositions, but with less general interconnections. A recent contribution is Dietrich (2003).

[3] Both reformulated preference aggregation in terms of the aggregation of *ranking propositions*, where different ranking propositions (e.g. "A>B", "B>C", "A>C") are logically interconnected. Both observed that majority voting, or some other propositionwise aggregation, over such ranking propositions does not generally preserve consistent interconnections between these propositions. Guilbaud abstracted this effect (which he called the "Condorcet effect") from ranking propositions to general binary propositions; so he arguably anticipated the more recent results on the "doctrinal paradox". Neither Guilbaud nor Murakami developed a general model of decisions on multiple interconnected propositions based on propositional logic (e.g. allowing decisions on atomic *and* compound propositions).

[4] Taylor (2002, 2004) discusses recent developments on strategic voting.

[5] Stone Sweet (2002) discusses path-dependence in precedent-based legal reasoning.

[6] The example is not affected by assuming that the agent has no initial view at all on R (instead of assuming that she is disposed to accept R in isolation). The agent's initial views on all other propositions – that these propositions are plausible in isolation – are essential.

[7] Formally, such an assignment is represented by a *truth-function* $v : \boldsymbol{L} \rightarrow \{1, 0\}$, where, for any proposition ϕ, $v(\phi)=1$ means that ϕ is true, and $v(\phi)=0$ means that ϕ is false.

[8] The set of connectives $\{\neg, \wedge\}$ is *expressively adequate*: For any proposition from $\boldsymbol{L}$, there exists a logically equivalent proposition whose only connectives are $\neg$ and $\wedge$.

[9] The famous killer amendment problem can be re-stated as a path-dependence problem in a decision on multiple propositions. Suppose there are three individuals: The first supports the proposal P and the amendment A, the second P but not A, the third A but not P. If all individuals reveal their views truthfully, there is a 2/3 majority for P, a 2/3 majority for A, but a 2/3 majority against ($P \wedge A$), the amended proposal. If the group votes on P first and on A next, then both are accepted; by implication ($P \wedge A$) is also accepted. If the group votes on A first and on ($P \wedge A$) next (having accepted A), then ($P \wedge A$), and by implication P, is rejected.

[10] The agent's propositional attitudes are modelled as acceptance or rejection and do not allow of *degrees* of belief. This is realistic in a political context where propositions are ultimately accepted or rejected, particularly by voting. My model might be generalized by using a logical system with more than two truth-values or by using real-valued credence functions of the form $\delta: X \rightarrow [0, 1]$. Pauly and van Hees (2003) have extended the analysis of judgment aggregation to many-valued logics, allowing different (discrete) degrees of belief, and shown that results similar to Propositions 4 and 5 below continue to hold. This suggests that results similar to Theorems 4 and 5 may also continue to hold.

[11] If the agent is perfectly rational, her AR-function can be extended to a truth-function.

[12] In the examples, I assumed that each proposition-negation pair is considered at the same time in the decision-path, and that the individual(s) reject(s) a proposition if and only if they accept its negation. My formal model is more general: It allows ϕ and $\neg\phi$ to be considered at different (possibly even non-adjacent) times in the decision-path.

[13] Other rules include the *priority-to-the-present* rule, by which the inconsistency is resolved by accepting the initial disposition on the new proposition and suitably revising previously accepted propositions. The agent might use this rule in cases where, given new evidence, she would rather reject some previously held beliefs than deny the new evidence. There are interesting research avenues here; see the literature on belief revision, e.g. Rott (2001). My

model also allows generalizations that preserve the priority-to-the-past rule. For example, suppose certain propositions are *non-negotiable*: The agent would never overrule them even given a logical conflict. Examples might be strongly believed factual premises or propositions considered *a priori* truths. The constraint that the priority-to-the-past rule should never overrule such propositions can be formalized in my model by restricting the set of admissible decision-paths to those where non-negotiable propositions occur first. If the set of non-negotiable propositions is logically consistent, under the modified model the agent would never overrule a non-negotiable proposition in a priority-to-the-past decision process.

[14] If we already have $\phi_t \in \Phi_{t-1}$ or $\neg\phi_t \in \Phi_{t-1}$, then $\Phi_t = \Phi_{t-1}$ under this definition.

[15] The reason for *not* defining $\Phi_t := \Phi_{t-1} \cup \{\neg\phi_t\}$ if $\delta(\phi)=0$ is to allow consideration of ϕ_t and $\neg\phi_t$ at separate steps in the decision-path. If δ is incomplete (defined below) this allows indecisive outcome sets ($\phi \notin M(\delta, \Omega)$ *and* $\neg\phi \notin M(\delta, \Omega)$) when $\delta(\phi)=0$ *and* $\delta(\neg\phi)=0$ *and* neither ϕ nor $\neg\phi$ is entailed by propositions accepted before ϕ or $\neg\phi$ are considered.

[16] If $M(\delta, \Omega_1)$ and $M(\delta, \Omega_2)$ are both decisive, then the existence of $\phi \in X$ such that $\phi \in M(\delta, \Omega_1)$ and $\phi \notin M(\delta, \Omega_2)$ implies that $\phi \in M(\delta, \Omega_1)$ and $\neg\phi \in M(\delta, \Omega_2)$.

[17] δ is strongly consistent if and only if δ can be extended to a truth-function v on $\boldsymbol{L}$ such that, for all $\phi \in X$ ($\subseteq \boldsymbol{L}$), $\delta(\phi) = v(\phi)$.

[18] All these properties can easily be proved in the propositional calculus.

[19] All violations of weak consistency are violations of strong consistency, but not all violations of strong consistency are violations of weak consistency.

[20] Under Theorem 2's conditions – the agent's propositional attitudes are complete and weakly consistent – violations of deductive closure are necessary and sufficient for *strong* path-dependence; here violations of deductive closure imply violations of strong consistency. By contrast, under Theorem 3's conditions – the agent's propositional attitudes are strongly

consistent – violations of deductive closure are necessary and sufficient for *weak* path-dependence; here violations of deductive closure imply violations of completeness.

[21] If δ is strongly consistent, a deductive closure violation implies weak path-dependence, by Theorem 3. If δ is not strongly consistent, this implies strong path-dependence, by Theorem 1.

[22] Nehring and Puppe's model (2002) (N&P) can also represent decisions on multiple propositions, but with less general interconnections. In N&P, agents do not have AR-functions over propositions from propositional logic, but preferences over vectors of *properties*, $<a_1, a_2, \ldots, a_m>$ ($\in \{0,1\}^m$). If all individuals accept the *same* compound propositions, the logical structure in my model can be represented in N&P's property structure, by identifying each property in N&P with an atomic proposition, and representing the unanimously accepted compound propositions by restricting the set of alternatives, i.e. the set of admissible $<a_1, a_2, \ldots, a_m>$ vectors. For example, if all individuals accept $(R \leftrightarrow (P \wedge Q))$, then each vector $<a_1, a_2, a_3>$ corresponds to an assignment of truth-values to P, Q and R; the rule $(R \leftrightarrow (P \wedge Q))$ can be captured by restricting the set of alternatives to $\{<1,1,1>, <1,0,0>, <0,1,0>, <0,0,0>\}$. But if individuals disagree about compound propositions, then they also disagree about the set of admissible alternatives. If individual 1 accepts $(R \leftrightarrow (P \wedge Q))$ while individual 2 accepts $(\neg R \leftrightarrow (P \wedge Q))$, then for individual 1 the set of admissible alternatives is $\{<1,1,1>, <1,0,0>, <0,1,0>, <0,0,0>\}$, while for individual 2 it is $\{<1,1,0>, <1,0,1>, <0,1,1>, <0,0,1>\}$. N&P's model requires a single set of alternatives. Therefore my model can represent more general logical connections than N&P's model.

[23] When a profile $<\delta_i>$ has been fixed and there is no risk of ambiguity, I sometimes identify the aggregation function δ with the collective AR-function $\delta_{<\delta_i>}$ induced by δ for the given profile $<\delta_i>$. On a precise notation: (i) $M(\delta_{<\delta_i>}, \Omega)$ is the outcome set of a decision process for the collective AR-function $\delta_{<\delta_i>}$ (induced by the aggregation function δ for the profile

$<\delta_i>$) and the decision-path Ω. (ii) $M(\delta, \Omega)$ is a function mapping each profile $<\delta_i>$ to a corresponding outcome set $M(\delta_{<\delta_i>}, \Omega)$ as defined in (i). Again, I sometimes simplify the notation by dropping the subscript $<\delta_i>$ and writing $M(\delta, \Omega)$ for $M(\delta_{<\delta_i>}, \Omega)$.

[24] The exact relation between sets of views on multiple interconnected propositions and preferences is non-trivial. On one interpretation, preferences are a special case of sets of views on multiple propositions (from *predicate* logic). An agent's preference ordering is here identified with a set of binary ranking propositions the agent accepts; e.g. the preference ordering $xPyPz$ is identified with the set of propositions $\{xPy, yPz, xPz\}$. Propositions like $\forall x \forall y \forall z((xPy \wedge yPz) \rightarrow xPz)$ can represent conditions like transitivity. Preference aggregation is then a special case of judgment aggregation on multiple propositions. It is arguably harder to interpret judgment aggregation as a special case of preference aggregation. Identifying each proposition with a single alternative in a preference context is insufficient, as judgment aggregation requires determining a collective set of judgments on these propositions, not a preference ordering over them. Identifying a *set* of propositions with a single alternative might be another route, but it might fail to make explicit the logical structure of such a set of propositions. A third route might be to identify each proposition with a *set* of alternatives, but this leads to the problem that classical models typically take as the unit of choice (or ordering) not such *sets* but *single* alternatives. While this might be solved via aggregation functions that produce *winning sets* rather than *winning alternatives*, one would still require a corresponding way of representing logical interconnections. It is also unclear what *sets of propositions* would be on that route (e.g. *sets of sets* of alternatives). The most promising route might be Nehring and Puppe's, discussed above, but even there the representable logical interconnections are restricted. It is too early to draw a conclusion. Future work might identify more fundamental parallels between the different frameworks.

[25] Since I consider a *decisive* $M(\delta, \Omega)$ here, strong and weak path-dependence coincide.

[26] To formalize this example using the definitions below, one needs to assign to individual 3 a preference ordering with respect to which individual 3 strictly prefers the outcome set $\{((P\wedge Q\wedge R)\rightarrow T), P, \neg Q, R, \neg T\}$ to the outcome set $\{((P\wedge Q\wedge R)\rightarrow T), P, Q, \neg R, \neg T\}$.

[27] Here Φ is the outcome set when *every* individual $i\in N$ expresses the "truthful" AR-function δ_i; Φ^* is the outcome set when individual i expresses the "untruthful" AR-function δ^*_i and every other individual j expresses the "truthful" AR-function δ_j.

[28] The social choice function satisfies universal domain, non-dictatorship, has at least three alternatives in its range, and always produces a determinate winner.

[29] Given that individual views are complete, weakly consistent and deductively closed.

[30] At least if $n > m$ (i.e. if $n/m > 1$).

[31] E.g. strong consistency on the one hand, and acyclicity on the other. The relations between consistency conditions in the two frameworks are discussed in List and Pettit (2003).

[32] Note that this definition permits $N_{accept\text{-}\phi} = \emptyset$ or $N_{reject\text{-}\phi} = \emptyset$.

[33] Here $\delta_{m_1}\delta_{m_2}$ is defined as follows: for any $\phi\in X$, $\delta_{m_1}\delta_{m_2}(\phi) = \delta_{m_1}(\phi)\delta_{m_2}(\phi)$.

[34] In that case, $\delta_{m_1}(\phi)\neq\delta_{m_2}(\phi)$; so $\delta_{m_1}(\phi)\delta_{m_2}(\phi) = 0 = \delta_{m_1}(\neg\phi)\delta_{m_2}(\neg\phi)$.

[35] My assumption of propositionwise aggregation is demanding, requiring that (i) the group's propositional attitude on each proposition be determined by aggregation over the individual views on that proposition (independent of other propositions), and (ii) the same aggregation function be used for all propositions. Pauly and van Hees (2003) and Dietrich (2003) have recently investigated weaker assumptions of propositionwise aggregation – e.g. keeping (i) but dropping (ii) – and obtained results broadly similar to Propositions 4 and 5, suggesting that results similar to Theorems 4 and 5 continue to hold. So even a weaker assumption of propositionwise aggregation may not provide an escape-route from path-dependence.

TABLES

Table 1. A Multi-member Government Example

	P "implement education project"	*Q* "implement health care project"	*R* "implement defense project"	*T* "increase taxes"	If *P* and *Q* and *R*, then *T*. "keep the budget balanced"
Individual 1	Yes	Yes	No	No	Yes
Individual 2	Yes	No	Yes	No	Yes
Individual 3	No	Yes	Yes	No	Yes

Table 2. Unidimensional Alignment

	Individual 4	Individual 1	Individual 5	Individual 2	Individual 3
P	1	1	1	1	0
Q	1	1	1	0	0
R	1	1	0	0	0
T	1	1	0	0	0
$((P \wedge Q \wedge R) \rightarrow T))$	1	1	1	1	1