

Rule by Multiple Majorities: A New Theory of
Popular Control

Sean Ingham¹

September 6, 2018

¹Department of Political Science, University of California San Diego.

Contents

1	Introduction	3
2	First attempts	31
2.1	Rule by “the” majority	33
2.2	Rule by multiple majorities: a provisional statement	40
3	An objection from social choice theory	58
3.1	Condorcet’s paradox	59
3.2	Two kinds of “majority cycle”	64
3.3	The spatial model of politics	73
3.4	Does politics have just one dimension?	85
3.5	Deliberation and single-peaked preferences	91
3.6	Conclusion	94
4	The concept of control	96
4.1	A preliminary statement	98
4.2	A more formal statement	110
4.3	Electoral control and the role of intentions	132
4.4	Whether control must be robust	140

4.5	Indifferent and divided citizens	145
4.6	Conclusion	149
5	Rule by multiple majorities	150
5.1	The multiple majorities view of popular control	151
5.2	Why identify popular control with majorities' control?	159
5.3	Popular control and voting rules	170
5.4	Popular control and democracy	175
5.5	Why the impossibility objection fails	182
5.6	A model illustrating rule by multiple majorities	186
5.7	Objections	191
5.8	Conclusion	201
6	Retrospective voting and control	202
6.1	Evidence of retrospective voting	207
6.2	Popular control with myopic and partisan voters	215
6.3	Popular control with heterogeneous voters	228
6.4	Conclusion	236
6.5	Appendix	237
7	On a normative theory of democracy	245
7.1	Preferences and their intensity	247
7.2	Equal control	252
7.3	The status of minorities	270
7.4	Conclusion	279
8	Conclusion	281

CONTENTS

iii

Bibliography

293

List of Figures

3.1	Single-peaked preferences	75
3.2	Multidimensional spatial model	79
3.3	Radial symmetry	80
3.4	Estimated ideal points of the 89th U.S. House of Representatives	81
4.1	Newspaper owner and editor	111
4.2	Newspaper owner and editor, with varying preferences	118
4.3	Newspaper owner and editor, with job protections	126

Acknowledgements

This book grew out of the interests in democratic theory and formal political theory that I developed in graduate school. I was lucky to have had teachers and mentors who nurtured those interests and did not push me down a more conventional path—Eric Beerbohm, John Patty, Maggie Penn, Dennis Thompson, Richard Tuck, and Ken Shepsle.

The core idea of the book evolved from ideas in a paper presented to the Political Science Department at Washington University in St. Louis in 2013. I owe a special thanks to John and Maggie, in particular, for the invitation and the encouragement to develop those ideas further.

Since leaving graduate school, I have been fortunate to find two political science departments that welcomed my unusual mix of interests, first at the University of Georgia and now at the University of California San Diego. I am grateful to the people at each institution who have supported me and to the many students and colleagues who have taken the time to read my work and talk with me about it.

I would like to thank Lucas Stanczyk for organizing a roundtable on the manuscript, and thank him and the other participants—Brookes Brown, Michael Kates, Frank Lovett, and John McCormick—for their thoughtful

comments and encouragement, both at the roundtable and since.

Many people have taken the time to comment on earlier drafts or sit and talk about the book with me. Thanks to Jesse Beder, Renee Bowen, Steven Brams, Lawrence Broz, Kanchan Chandra, Scott Desposato, Keith Dougherty, Tiberiu Dragu, Patrick Egan, James Fowler, Bob Grafstein, Cathy Hafer, Seth Hill, Alex Kaufman, Thad Kousser, Matt Landauer, Gerry Mackie, Bernard Manin, Jane Mansbridge, Ryan Pevnick, Molly Roberts, Emma Saunders-Hastings, Melissa Schwartzberg, Christina Schneider, Itai Sher, Branislav Slantchev, David Stasavage, Don Tontiplaphol, and Tom Wong.

I owe a special thank you to Dimitri Landa and David Wiens, who each invested more time to understand the book and help me improve it than I could have reasonably asked.

I am thankful to three anonymous readers for Cambridge University Press for providing invaluable feedback; to Robert Dreesen for finding such good reviewers, and taking an interest in the project in the first place; to John Patty, Maggie Penn, and Jeff Frieden for proposing the book's inclusion in the PEID series; and to Shaheer Anwarali, Jackie Grant, Josh Penney, and Jared Wright for all their help during the book's copyediting and production.

Finally, I would like to thank my parents, John and Mary, and my dad in particular for what he has taught me about scholarship; Eleana, for her unflagging patience, love, and moral support, and for the countless sacrifices she has made for me; and Patrick, for making me laugh more during the last two years than the previous two decades. The book is dedicated to them.

Chapter 1

Introduction

In a democracy citizens ought to have some control over how they are governed. If they do not participate directly in making policy, they ought to maintain control over the public officials who design policy on their behalf. This book develops a novel theory of popular control: an account of what it is, why democracy's promise of popular control is compatible with what we know about actual democracies, and why it matters.

That democracy involves some kind of popular control is widely accepted. The assumption shows up in descriptions of representative democracy as a relationship between a "principal" and its "agents." The people of a representative democracy are described as though they were an individual, a "principal," who selects agents to act on its behalf. Just as the owner of a properly organized business retains some control over the manager whom he hires to run his business, the people of a properly organized democracy are supposed to retain some control over the legislators and public officials whom they elect to govern on their behalf.

This metaphor is compatible with different understandings of representation, the proper scope and degree of control that the people should retain

over their representatives, and the mechanisms that do or should undergird their control. Marx invoked this metaphor when describing the Paris Commune, where “universal suffrage was to serve the people, constituted in Communes, as individual suffrage serves every other employer in the search for the workmen and managers in his business.” The Parisians were to have the right of revoking elected representatives at any time, just as “companies, like individuals, in matters of real business generally know how to put the right man in the right place, and, if they for once make a mistake, to redress it promptly.”¹

The idea that elected representatives are agents of the people and subject to some degree of popular control does not commit one to the Commu-nards’ view of representation. One may accept that public officials should be leashed to the sentiments and opinions of the people but favor a long and slack leash. James Madison famously praised representation for its tendency “to refine and enlarge the public views, by passing them through the medium of a chosen body of citizens, whose wisdom may best discern the true interest of their county.” But even though he did not wish to see elected officials slavishly cater to their constituents’ given preferences, he still described the federal and state governments as “but different agents and trustees of the people.” His beliefs about representation did not keep him from criticizing the “adversaries of the Constitution [who] seem to have lost sight of the people altogether in their reasonings on this subject” and mistakenly viewed the federal and state governments as though they were

¹Karl Marx, *The Civil War in France*, in Marx (1977).

“uncontrolled by any common superior.”² They were, he thought, to be controlled by the people.

Accepting that the people should have some kind of control over their elected representatives leaves open questions about the proper means and scope of popular control. Thinkers as different as Madison and Marx can accept the need for popular control but disagree over whether it is desirable for the people to be able to recall their representatives at any moment or issue binding instructions. They can disagree over whether the people should have control over particular details of policy, or simply over the broad contours of policy, or the general objectives and values that inform policy-making.

The claim that democracy involves some form of popular control is thus widely accepted, but it is vulnerable to an objection. Meeting the objection is the primary aim of the book. The people are not a unified person, as the principal-agent metaphor pretends, but rather a multitude of separate individuals with different opinions and preferences. A principal has desires and goals that determine how it selects and disciplines its agents, but different citizens want different things from their elected representatives. How could elected representatives be viewed as agents acting on behalf of a collection of individuals who have conflicting preferences and interests? It is natural to think that in order to view the government as the people’s agent and under their control, there must be some sense in which the people have a common *will*, which they can get their agent to carry out. If it is illegitimate to attribute a common will to the people, then the identification of democracy

²*The Federalist* 10, 46.

with popular control may seem indefensible.

What prompts the objection is not the principal-agent metaphor in particular, but rather an intuition about the concept of control, which the principal-agent metaphor registers. Ordinarily, if a person has control over a variable, one expects the variable to conform to her preferences. Consider someone who has control over the temperature in her house, thanks to a thermostat. If she wanted it to be warmer or colder, then it would be. Analogously, if the people have control over the government, one would expect the government to conform to their preferences. But citizens normally have conflicting preferences over what the government should do. How can we understand the idea of popular control so that it does not require attributing a unified will to a diverse collection of separate individuals?

A standard answer is that popular control—or “popular rule,” to use a different idiom—consists in the rule of the majority. A government cannot be compelled to do what everyone wants, because different people want different things, but the majority can compel the government to do what the majority wants. The government can be made to respect “the will of the majority,” to use the familiar phrase.

The puzzle about popular control is not resolved so easily. Popular control cannot be identified with the control of a *single* majority, imposing its unified will on the government and political process, because typically there is no such majority. Instead, there are multiple, overlapping majorities—multiple subsets of the citizen body each of which constitutes a numerical majority of the whole. Each majority will share preferences over some, but not all of the political alternatives under consideration. No majority

is likely to share a “will” in any straightforward sense of the term. In a majoritarian democracy that distributes political power equally, one would expect the outcome of the political process to reflect each majority’s shared preferences. Recognizing this simple fact is the first step in developing a workable account of popular control.

Taking this step leads us to the right questions—is it possible to empower every majority at once? what would it mean for multiple majorities to have control over policy?—but a formidable obstacle still stands in the way. The shared preferences of different majorities may be impossible to satisfy simultaneously, because citizens’ preferences may be such that every feasible policy frustrates the shared preferences of some majority. The simplest example involves a choice between three policies, where one majority prefers the first to the second policy, another majority prefers the second to the third, and yet another majority prefers the third to the first. In these situations it would appear to be impossible for majorities to have control over policy, if control implies that majorities can make policy conform to their shared preferences.

This simple example and the insight it generates are at the foundation of social choice theory, a body of formal theorems that concern the aggregation of individual votes and preferences into collective decisions (or “social choices”). Some thinkers, most famously William Riker (1982), have drawn strong negative conclusions from social choice theory, arguing that we must abandon traditional claims about democracy as a form of popular rule. Political theorists’ initial reception of social choice theory was largely dismissive, even hostile, at least in part because Riker was its chief ambassador. Chris-

tiano (1993) argued that the “impossibility results of social choice theory have much less significance for normative democratic theory than is usually claimed.” According to Pildes and Anderson (1990), “social choice theory only muddies efforts to think clearly about democracy,” originating “in profound misconceptions about the structure of public values, the nature of democratic politics, and the concept of rationality itself” (p. 2213). Since the publication of Gerry Mackie’s (2003) influential attack on the empirical basis for Riker’s argument, most democratic theorists have proceeded on the assumption that social choice theory is of no more than marginal importance to democratic theory, judging by the amount of attention it receives in the literature.

The perspective presented here differs from the two typical responses. Social choice theory is indisputably relevant for how we understand ideas like popular control or popular rule. But the lesson to draw from it is not that these ideas are meaningless or incoherent or indefensible. Social choice theory shows us that some traditional interpretations of these ideas—indeed, any that invoke “the will of the majority”—rest on shaky ground, but it is also possible to put notions of democratic rule and popular control on a rigorous foundation, as I will try to show. The conception of popular control developed here is fully consistent with insights from social choice theory. In the remainder of this introduction, I give a short overview of the puzzle from social choice theory and the account of popular control that I will develop in response.

Social choice theory developed into a systematic, scholarly discipline with the publication of Arrow (1951)’s seminal theorem on the possibility of ag-

gregating individuals' preferences into a "collective" or "social" preference ordering. The theorem admits different interpretations, but under one familiar interpretation it speaks directly to questions about how to define "the popular will."³ Suppose that the will of the people, if such a thing exists, takes the form of an ordering of the possible objects of collective choice, from best to worst, and this ordering depends solely on how the individual citizens order the same objects. Suppose, further, that how the popular will ranks any two alternatives depends solely on citizens' preferences over *that* pair of alternatives—and not, for example, on their preferences over any *other* pair of alternatives. Finally, suppose that if everyone prefers one alternative to another, then it is also ranked as the better of the two according to the popular will. From these suppositions we may conclude, by Arrow's theorem, that there is a single citizen whose preferences unilaterally determine the content of the popular will, in the manner of a dictator.⁴ Thus, if the suppositions are part of any reasonable definition of the "pop-

³The two most common interpretations of the theorem, both present in Arrow's original monograph, treat it as either a statement about voting rules or a statement about judgments of social welfare; on the differences between these two interpretations, see Sen (1977*b*) and Pattanaik (2014). These two interpretations do not exhaust the possibilities, however; for an explanation of the theorem's broader reach, see Patty and Penn (2014) and Ingham (2018).

⁴The first supposition incorporates the assumption that the social ordering is transitive; the second and third assert "independence of irrelevant alternatives" and the Pareto condition, respectively.

ular will,” but no reasonable definition would identify it with the will of a single “dictator,” then the implication is that there is no satisfactory way to define “the popular will.” If we accept this conclusion, then references to what the people want are often vacuous, as meaningless as references to the largest natural number. If “the will of the people” often fails to refer to anything, then it makes no sense to claim that democracy should allow the people to rule by electing representatives who carry out their will.⁵

Riker (1982) is the most emphatic advocate for viewing social choice theory and its significance along these lines. According to Riker, Arrow’s theorem and other results from social choice theory show that even if citizens have rational, determinate preferences, their preferences do not typically combine to make up a coherent “popular will” to which elected representatives might respond—and in those cases where they do produce a popular will, there is no reason to expect voting or other democratic mechanisms to be capable of revealing it or producing decisions in conformity with it. “Social choice theory forces us to recognize that the people cannot rule as a corporate body in the way that populists suppose. Instead, officials rule, and they do not represent some indefinable popular will” (p. 244). Like Schumpeter (1942) before him, Riker concludes that democracy should be identified merely with the form of government in which rulers are selected through elections and the people have periodic opportunities to remove them

⁵For more on the implications of Arrow’s theorem for the concept of the popular will and the view that legitimacy requires respecting the popular will, see Ingham (2018).

from power. Since elections cannot be said to empower the people to sanction public officials for failing to do as the people want—“as the people want” is an empty phrase—this form of democracy, Riker tells us, is not “popular rule, but rather an intermittent, sometimes random, even perverse, popular veto” (Riker, 1982, pp. 238, 239, 244).⁶

In one form or another, Riker’s view about the implications of social choice theory is broadly shared among political scientists.⁷ In their recent book, for example, Christopher Achen and Larry Bartels concur with Riker, concluding from their review of social choice theory that the “populist ideal . . . suffers from grave logical and practical problems” (Achen and Bartels, 2016, p. 30).

This skepticism about popular rule rests on an implicit conceptual claim. The assumption is that popular rule implies, by definition of the concept, the existence of a popular will. At its core, Riker’s critique is not that popular rule is dangerous or ineffective, but rather that the demand for popular rule makes no sense because it implies an identifiable popular will, but citizens’ preferences do not reliably produce in the aggregate anything resembling a popular will, and even if they did, there would be no reliable way to identify it. One of my primary contributions in this book is to show that the conceptual claim is one we should reject. Popular control—and popular rule, to the extent that it just amounts to meaningful popular control over

⁶For fuller expositions and forceful critiques of Riker’s argument, see Mackie (2003) and Dowding (2006). I discuss one element of Mackie’s critique in chapter 3.

⁷See, for example, the “hall of quotations” in Mackie (2003).

important decisions—does not imply the existence of a popular will. Before giving a preview of this contribution, let me fill out the sketch of social choice theory and the reasons for Riker’s skepticism.

The insight that motivates Arrow’s impossibility theorem can be illustrated with a simple example, which also illustrates why Riker’s objection applies even, indeed especially, to the view that references to the will of the people can be understood as references to the will of the majority. Suppose a country must choose between escalating a war, maintaining the current level of engagement, or withdrawing all forces and ending the war, and public opinion is evenly split into three different points of view. Members of the antiwar coalition want to end the war and believe escalation would be the worst option. Hawkish liberals and conservatives who are optimistic about the current course believe it represents the best option and ending the war would be the worst option. And a more pessimistic contingent of hawks believe that the country should either end the war or, better yet, ramp up its effort, but the current policy must be abandoned. To summarize, the public’s preferences over the three options are:

Pessimistic hawks: escalation \succ withdrawal \succ status quo

Antiwar group: withdrawal \succ status quo \succ escalation

Optimistic hawks: status quo \succ escalation \succ withdrawal

Here and throughout the book, the symbol \succ is used to describe a preference ordering: pessimistic hawks prefer escalation to withdrawal and prefer each of these to the status quo, and so on. (And, to avoid misunderstanding,

let me add that the binary relation *prefers* ___ *to* ___ should be construed broadly, as roughly synonymous with the binary relation *would choose* ___ *over* ___, *given the choice*. It does not refer to a person's mere "tastes" or "inclinations." I say more on the terminology in chapter 2.)

Observe that no matter which course of action the government takes, its choice will be opposed by a large majority who believe that one of the other two options would have been better. A majority are opposed to escalating the war: the optimistic hawks and antiwar contingent both prefer the status quo to escalation. Another majority—the antiwar contingent and the pessimistic hawks—would rather withdraw all forces than stay the course. But another majority—the two hawkish contingents—would rather escalate the war than withdraw all forces. In this situation, it is impossible to satisfy all majorities.

How is it possible for the government's decision to be subject to popular control if its decision will necessarily run counter to the preferences of a large majority? Spain threw its support behind the 2003 U.S. invasion of Iraq despite massive opposition from its citizens. According to public opinion surveys conducted in the spring of 2003, around 90 percent of Spaniards opposed the war and around 80 percent opposed Prime Minister Aznar's position.⁸ Hundreds of thousands marched in protest. It would have been natural to view this event as evidence that the Spanish people lacked control over their government's foreign policy. But if we accept that inference—from the observation that the government has acted contrary to the wishes

⁸Goodman (2003).

of an overwhelming majority, to the conclusion that the people do not have control—then we will have to conclude that in our hypothetical example popular control is impossible. For in the example, the government *inevitably* acts contrary to the wishes of an overwhelming majority of its citizens, no matter how it acts.

This train of thought illustrates the pattern of argument leading to skepticism about popular control, popular rule, and like concepts. There is a conceptual claim, often left implicit: “if the people have control over what the government does, then the government will not act contrary to the wishes of an overwhelming majority.” Then a claim is made about citizens’ preferences: “citizens’ preferences over the government’s options are such that it will inevitably act contrary to the wishes of an overwhelming majority.” Together they produce the conclusion that popular control—or whatever the democratic ideal targeted for refutation may be—is impossible. To resist the skeptical conclusion, one must either contest the conceptual claim about the logical implications of popular control, popular rule, or whatever the target of the skeptical attack is,⁹ or, conceding the conceptual claim, maintain its

⁹ Riker described his target as the idea that democracy establishes *popular rule*, not the idea that it establishes *popular control*. In fact, he even suggested at points that in his favored conception of democracy (“liberalism”), voters would enjoy a “negative kind of control” (p. 243), and that their “random, even perverse popular veto” would “sometimes” make it possible to “restrain official tyranny” (p. 244). I explain in chapter 5 why Riker’s description of this veto as instituting even a limited, negative form of control is unjustified.

compatibility with what we know about citizens' preferences.

For the most part, Riker's critics have followed the second strategy. The most influential response to Riker is to question whether "majority cycles"—like the configuration of preferences in the hypothetical example—occur in practice.¹⁰ Another common response is to claim that deliberation can modify individuals' preferences in such a way that cycles will not arise or persist (Miller, 1992; Knight and Johnson, 1994; Cohen, 1997; Dryzek, 2000; List, 2002; Dryzek and List, 2003; Farrar et al., 2010; Knight and Johnson, 2011). Some of these writers concede Riker's claim that if majority cycles *did* occur, then it would be bad news for democracy (Farrar et al., 2010, p. 337).

I will argue that the central assertions in these responses to Riker—that majority cycles are rare, that deliberation has a tendency to prevent or eliminate them—are harder to defend than commonly realized. We should therefore ask whether another response is possible, one that contests the conceptual claim that popular control implies a popular will, while conceding—for the sake of argument, and perhaps truth also—the claim about majority preference cycles.

Whether one can defend the idea of popular control against the skeptics' accusations of logical incoherence should interest a broad audience. Political theorists regularly identify democratic popular control—over leaders and public officials, or over the general direction and contours of policy, or over consequential social outcomes downstream of policy, if not over partic-

¹⁰Mackie (2003). See also Regenwetter et al. (2006).

ular policy decisions themselves—as an important principle of democratic legitimacy or as a goal of institutional design (McCormick, 2011; Pettit, 2012; Shapiro, 2016; Landemore, 2017; Saunders-Hastings, 2018). Theorists of democratic socialism, and workplace democracy in particular, often invoke the concept of workers’ democratic control (Marx, 1977; Hsieh, 2005; Gourevitch, 2015, 2016), a concept exposed to the same skeptical attacks as popular control over the state. Where political theorists do not refer explicitly to popular control *per se*, they often use other terminology—popular rule (Christiano, 1996), popular sovereignty (Tuck, 2016), collective self-rule (Abizadeh, 2012)—to express some variation on the thought that citizens ought to have, in some sense, an effective “say” over what happens in a democracy. Sometimes, political theorists use the concept of popular control to explicate the content of related ideas like popular sovereignty (Lane, 2016). Yet, with few exceptions—Pettit (2012), for example—no one has attempted a careful, systematic account of the concept of popular, democratic control, one which explains why it is compatible with the absence of anything resembling a collective will and, in particular, compatible with cyclical majority preferences.

In their rejection of the idea that elections allow democratic citizens to exert some control over their leaders, skeptics like Riker and Achen and Bartels are challenging not just normative democratic theory but much applied research in political science, too. A large literature uses principal-agent models to explore the conditions under which elections enable voters to control public officials and hold them accountable (Barro, 1973; Ferejohn, 1986; Fearon, 1999; Persson and Tabellini, 2000; Maskin and Tirole, 2004;

Ashworth and Bueno de Mesquita, 2006; Besley, 2006).¹¹ Complementing this work is an empirical literature on retrospective voting and electoral accountability (Key, 1966; Fiorina, 1981; Healy and Malhotra, 2010; Alt, Bueno de Mesquita and Rose, 2011; Huber, Hill and Lenz, 2012; Achen and Bartels, 2016).¹² Research on popular referendums and initiatives asks whether they give voters more control over policy and whether the effects of increased control are desirable (Lupia and Matsusaka, 2004; Matsusaka, 2010; Le Bihan, 2016).¹³ If the skeptics are right, the fundamental question that many scholars contributing to this literature take themselves to be addressing—do elections and other democratic institutions allow citizens to control their leaders and policy?—is conceptually muddled.¹⁴

¹¹For a review, see Ashworth (2012).

¹²For a review, see Healy and Malhotra (2013).

¹³Terms like *accountability* or *responsiveness* sometimes substitute for *control*. Despite their different connotations, authors in this literature use them to express roughly the same thought that I am describing as a thought about popular control.

¹⁴Achen and Bartels themselves go on to present empirical evidence against the hypothesis that elections establish popular control, even though the implication of their arguments about social choice theory in chapter 2 appears to be that such hypotheses are conceptually confused. Charitably interpreted, their argument might be understood as: (1) social choice theory describes conditions under which questions about popular control are incoherent, (2) but even if one restricts attention to cases where those conditions do not hold and such questions are at least intelligible, the empirical evidence

Scholars working on these topics justifiably rely on simplifying assumptions and bracket the hard conceptual questions about what popular control means when citizens' conflicting preferences do not aggregate nicely into anything resembling a collective will. Principal-agent models examine the relationship between a single representative voter, such as a median voter, and elected officials. Empirical studies restrict attention to issues where one may plausibly assume "all citizens have common interests, such as a prosperous and peaceful society" (Healy and Malhotra, 2013, p. 287) or to binary policy decisions, where references to the will of the majority are basically unproblematic (Lax and Phillips, 2009; Matsusaka, 2010). What these assumptions simplify is the heterogeneity of democratic citizens and the potential complexity of collective decisions in the real world. With the heterogeneity and complexity safely assumed away, the puzzles about preference aggregation do not arise. Simplifying assumptions are justified for social scientific inquiry, but one would like to know whether the idea of popular control can be defended against skeptical attacks without these crutches.

This book offers a general account of popular control that does not presuppose that citizens' preferences aggregate into a popular will, or even that one can speak meaningfully of the will of the majority. I argue we should identify popular control with the control of all majorities, i.e., with the simultaneous control of every subset of the citizen population that makes up a majority of the whole. A majority subset has some control over a variable—such as the behavior of an elected official or a policy outcome—

weighs strongly against the hypothesis of popular control.

if the majority's shared preferences would constrain the variable provided that its members cared enough about it relative to other things they care about. This condition for control can hold for every majority subset of the population, at the same time, even if majorities' preferences exhibit cycles and citizens' preferences fail to aggregate into anything resembling a popular will.

The core of this account consists of two claims about control: first, a collection of individuals, such as a majority subset of the citizen population, can be said to have control even if the collection does not remotely resemble a group agent; and, second, whether such a collection has control over something depends on the truth of counterfactuals—what would happen if its individual members were to have certain preferences and attach sufficient importance to their satisfaction—and not on whether their actual preferences are satisfied. The first claim is what leads to the identification of democratic, popular control with the control of every majority subset of the citizen population—even majority subsets whose members do not conceive of themselves as a collective political force. The second claim is what allows one to attribute control to a majority even when—as must inevitably happen in certain cases of cyclical majority preferences—its shared preferences are frustrated. On this account, popular control does not imply that the people have a collective will, nor that any majority has a collective will, nor that majorities have transitive preferences. It is compatible with all of the discoveries from social choice theory that the skeptics have taken to discredit the idea of popular rule.¹⁵

¹⁵The view defended in the book is similar in spirit to the view set out in

As these remarks suggest, this account of popular control does not require us to imagine “the people” as a singular noun, as a kind of group agent with its own beliefs and desires. The book does not defend a theory of group agency.¹⁶ It is rather an account of how citizens, pursuing disparate and even conflicting goals, their preferences aligning over only a limited range of options, and bearing little more resemblance to a singular group agent than a crowd of pedestrians at an intersection, can nonetheless have power and control over public officials, policy, and other variables. The people, in this account, are emphatically a ‘they,’ not an ‘it.’

Partly for this reason, the book is not a defense of *populism* as many writers understand it. For Riker, populism was the doctrine that what the people want ought to become law. I agree with Riker that it is rarely meaningful to speak about ‘what the people want.’ Jan-Werner Müller argues we should understand populism as “a way of perceiving the political world that sets a morally pure and fully unified...people against elites who are deemed corrupt or in some other way morally inferior.” The problem with populism is that, as Müller argues, “the idea of a single, homogeneous, authentic people is a fantasy,” and a dangerous one, at that (Müller, 2016). I agree completely. In fact, I go further. We should reject not only the populist’s conceit of a fully unified people, but also the idea of a single, homogeneous *majority*. They are each fictional characters, useful only for ideological ends. Any majority subset of the population is a heterogenous

Ingham (2016*b*). But there are major differences, which I explain in chapter 5, footnote 3.

¹⁶For a theory of group agency, see List and Pettit (2011).

collection of citizens who may agree on some points but not on all, and its membership overlaps with other majority subsets, of which the same is also true. It is therefore problematic, in most cases, to speak of “the will of the majority” and to identify popular control or popular rule with the rule of the majority.

But after we see the populist’s fantasy for what it is and disabuse ourselves of the fiction of a monolithic majority, must we also abandon the idea of popular control as well? We need not. There is a perfectly intelligible sense in which all majority subsets can have simultaneous control over something, despite the heterogeneity of citizens’ preferences. We can identify popular control with an arrangement in which every majority has control. When majorities of ordinary citizens all have control, the distribution of control is ‘popular’ in character. The adjective plays the same semantic role as it does in the term *popular culture*; rather than indicating a mythical group agent, the People, to whom control belongs, it indicates the egalitarian and democratic flavor of the control exercised simultaneously by different majorities. Accordingly, I use the terms *popular control* and *democratic control* interchangeably.

I refer to this view as the multiple majorities conception of popular control, and refer to the arrangement it describes as the rule of multiple of majorities.¹⁷ Apart from freeing the idea of popular control from unrealis-

¹⁷Almost no political theorists put any importance on the fact that the population contains multiple majorities each of which must presumably be empowered in a democracy that gives the same political rights and resources to everyone. One exception is McGann (2006). On the difference between

tic assumptions about citizens' preferences, it gives us a novel perspective on various debates in normative democratic theory and political science, from long-standing questions about democracy and majority tyranny to recent debates about "myopic" retrospective voting and its implications for democratic accountability.

I should emphasize that this account of popular control is not a defense of majority voting. Other political theorists have considered the implications of social choice theory for the justification of collective decision-making procedures like majority rule (Risse, 2004; McGann, 2006; Knight and Johnson, 2011; Elster, 2013). The question I take up is not about decision procedures per se. A decision procedure, like majority rule, is a possible *instrument* for bringing a decision under popular control; it is not the same thing as popular control. To judge whether a particular voting rule is a good instrument, we first need an account of what it means for a decision to be under popular control. Such an account is the primary contribution of the book. Although I discuss, at length, its application to models of elections and, more briefly, its implications for voting rules, my primary concern is the conceptual puzzle about popular control that I have described: what could it even mean to say that this or that voting rule brings decisions under popular control, if we know—irrespective of the voting rule in place—a majority will prefer a different decision to whatever the voting rule produces.

To provide some motivation for the multiple majorities concept of popular control, I begin in chapter 2 by explaining the problems with standard his project and mine, see footnote 18 as well as the discussion in Chapter 7.

views that identify it with “the rule of the majority,” as if there were only one majority subset of the citizen population responsible for the decisions taken in a democracy. Once we recognize the problems with the idea of a singular majority, we are led to ask whether and in what sense it is possible for every majority subset of the population to have control, at the same time. In this chapter, I give a provisional statement of what it means for every majority to have control—provisional because it falls to an objection from social choice theory.

Chapter 3 elaborates this objection. When citizens have preferences like those in our hypothetical example involving the war, where majorities’ preferences form “cycles,” it is impossible for the object of popular control to conform to the shared preferences of every majority. I explain in this chapter why—contrary to what some political theorists have argued or assumed—there is little reason to believe that such configurations of preferences are rare in practice. If we wish to defend the ideal of popular control from the charge that majority preference cycles render it incoherent, then we should re-examine the underlying assumptions about the meaning of popular control, on which this charge implicitly rests. Rather than try to refute the skeptic’s factual claim that majority preference cycles occur, we should contest the skeptic’s conceptual claims about their implications for the ideal of popular control.¹⁸

¹⁸In this respect, my strategy for defending the ideal of popular control resembles Anthony McGann’s (2006) strategy for defending democracy against Riker’s (1982) skeptical arguments. But while each strategy concedes the same premise—that majority preference cycles occur—they do so in order to

In trying to meet this objection, one should be wary of *ad hoc* and question-begging solutions. The skeptic's argument fundamentally rests on the conceptual claim that if a decision is under popular control, then it will respect the preferences of majorities. We must *explain* why this claim is mistaken. If we just *assert* that it is false, or start our argument from premises that trivially imply as much, then we are just begging the question in dispute. To avoid the charge of being ad hoc, our explanation should rest on a general and *independently* plausible analysis of control, one that has something to be said for it besides the contribution it makes to meeting the challenge from social choice theory, and which explains why we should reject the key conceptual claim about popular control. We should start with an explanation of what it means *in general* for a group to have control over something. This explanation will not refer to anything narrowly political, much less democratic, like elections or a specific voting rule. After all, nobody would think that voting rules or other formal procedures should enter into an explanation of what it means for, say, two parents to have control over their child. To arrive at a satisfactory, general analysis of control, we should reflect on non-political cases where we confidently judge that a person or group has control and then find a definition that fits with these judgments as far as possible. The definition must agree with these judgments to some defend different conclusions. None of the conclusions about democracy that McGann defends is a conclusion about *popular control* per se; McGann's goal is to justify democratic procedures like majority rule, not to defend the idea of popular control.

extent, because otherwise our conception of popular control will not be recognizable as a conception *of* popular control as opposed to something else, and any conclusions we draw will not be recognizable as contributions to the debates of interest to us.¹⁹

I undertake this task in chapter 4. Abstracting from the special case of popular control in democracy, I consider what it means in general for an individual or group to have control over something. A group—such as a majority of citizens—can be said to have control over a variable, I claim, if their shared preferences constrain the variable, provided the group’s members attach enough importance to the variable relative to other variables they care about. The smaller the minimum amount of importance they must attach to the variable in order for their shared preferences to constrain it, the more control they have. This analysis provides the conceptual foundation for a revised majoritarian account of popular control. The analysis is justified by reflection on ordinary, non-political examples of control, so the resulting conception of popular control, built on top of it, is not an *ad hoc* or question-begging solution to the problem posed by social choice theory. The analysis in this chapter also yields other conclusions with important implications for a theory of democratic control, such as the thesis that an agent or group of agents can exercise power and control without intending to do so, and the claim that whether an agent or group of agents has control usually depends on the dispositions of other agents.

¹⁹My understanding of conceptual analysis has been influenced by the discussion in Lovett (2010).

With this conceptual foundation in place, I formulate in chapter 5 the multiple majorities account of popular control. I explain how every majority can have control over a variable even if the variable is bound to take on a value that frustrates the preferences of some majority—and thus even if majorities’ preferences are cyclical. What is necessary is that for each majority, its members’ shared preferences would constrain the variable if its members cared enough about it relative to other things they care about. After explaining why this conception of popular control is not vulnerable to the objection from social choice theory, I explain why it is plausibly taken as a description of one of the essential features of a democratic regime, and consider a possible mechanism by which democratic elections could, in theory, establish popular control over policy, irrespective of whether majorities’ preferences are cyclical.

In the remaining chapters I consider the implications of accepting the concept of popular control for the explanatory and evaluative aims of political science and normative democratic theory. In chapter 6, I turn to the theory of retrospective voting—the idea that voters choose whether to support an incumbent officeholder’s bid for re-election in part on the basis of retrospective assessments of performance—and ask whether it is a plausible mechanism by which majorities could acquire control over elected officials. The chapter addresses several sources of skepticism about the idea, using the novel account of popular control defended in earlier chapters. Achen and Bartels (2016) and other writers have argued that retrospective voting is “blind,” a reflection of non-rational, knee-jerk psychological mechanisms, such as shark attacks or bad weather causing voters to feel antipathy towards

incumbent officeholders. Others have wondered whether retrospective voting could produce meaningful electoral control if voters use heterogeneous performance standards—for example, each voter using his own personal welfare as his standard (Hibbs, 2006). And there is an even more fundamental, conceptual question—whatever it is that retrospective voting produces, even in the best of cases, does it count as genuine *popular control* or popular rule? One democratic theorist scoffs at the notion that “accountability through retrospective voting amounts to genuine rule by the people,” wondering “today, at the beginning of the twenty-first century, who can still make this [argument] with a straight face?” (Landemore, 2017, p. 147). Another criticizes the idea that retrospective voting could produce anything like popular rule because the outcome of retrospective voting may not “approximate the popular will,” as “there may not be an underlying popular will for many issues in the first place” (Green, 2010, p. 106).

The purpose of chapter 6 is to explain why retrospective voting could be a mechanism by which majorities acquire control over elected officials *even if* one grants the premises of the arguments skeptics make against this view. By the arguments of chapter 5, there need not be any popular will in order for retrospective voting to produce popular control. And if popular control over the agents acting on citizens’ behalf suffices for “genuine rule by the people,” then one can say, with a straight face, that retrospective voting could be a mechanism of popular rule. I use simple formal models to explain why majorities could have control over elected officials, in virtue of the electoral incentives their retrospective voting behavior creates, even if voters’ decisions are largely explained by non-instrumental partisan attach-

ments, and not rational assessments of what their elected representatives have done for them; their retrospective assessments are influenced by arbitrary, irrelevant events beyond the incumbent's control; and voters judge the incumbent by heterogeneous performance standards. I do not claim that the formal models describe actual democracies in all relevant respects; they are meant rather as thought experiments, which show that the reasons skeptics have given for doubting that actual retrospective voting produces popular control are not good reasons.

In chapter 7, I consider the concept of rule by multiple majorities from the perspective of normative democratic theory. The extent to which a democratic regime gives majorities control, either over policy or over the public officials who choose policy, is plausibly one of the reasons to prefer it to alternative regimes. The goal of this chapter is not to give a full defense of that claim, but instead to reply to several objections that the claim is likely to meet with. One worry is that popular control on my account is too easily achieved, so I explain in this chapter how the concept of popular control can be made suitably demanding, by articulating the sense in which a system of popular control ought to be egalitarian. A second objection is that the concept of popular control inappropriately treats facts about preference intensity as relevant to democratic decisions when what matters, in a democracy, is “how many people want something rather than how much they want it” (Shapiro, 2016, p. 48). The main worry is that counting how much people want something will end up empowering intense minorities. In reply, I explain why popular control on my account does not require that the minorities with sufficiently intense preferences be allowed

to prevail over majorities. Finally, I give examples of how the concept of rule by multiple majorities generates novel insights into a familiar reservation with majoritarian democracy, namely that such regimes carry the risk of “majority tyranny.” Every minority belongs to various majorities, and giving all of those majorities control over political decisions therefore has important implications for how we understand the place of minorities in a majoritarian democracy.

The book leaves important questions about popular control, the conditions under which it emerges, and the reasons to value it unanswered, but it provides the conceptual tools for investigating them. Articulating new and useful concepts for the study of politics—concepts that help us recognize possible explanations for political phenomena and possible reasons for accepting or criticizing political institutions and practices—is one service, perhaps *the* service, that political theory can and should perform. That view of political theory’s aims is not unusual, but in its use of formal models, the book illustrates an approach to that project that remains unorthodox. Or rather, it is unorthodox within “political theory” understood as a field of inquiry distinct from “formal political theory,” a demarcation that is useful for an ethnography of political science as a discipline, but less useful for an understanding of the intellectual contributions that the different theoretical parts of the discipline make. One hope for the book is to blur that demarcation in the reader’s mind by showing how formal models are useful for political theory. Models are valuable in political theory, not as sources of empirically testable predictions, but as “tools for conceptual exploration” and aids in the communication of abstract ideas (Johnson, 2014). As John-

son argues, models play this role throughout political theory, even if the use of formal, game-theoretic models remains rare. I comment further on this approach to political theory in the book's concluding chapter.

Chapter 2

First attempts

If a captain has control over a ship, then one would expect the ship to go where the captain wants. Control is a power to make a feature of the world conform to one's will. This, anyway, is an intuition that guides judgments about popular control and popular rule. The intuition leads to the view that if the people have control over some feature of the world, then it ought to behave as they want. If the people have control over the state, or over public officeholders, then these agents must respect the will of the people, or so this intuition suggests.

How could one legitimately attribute a single will to a heterogeneous body of citizens whose members have varied and conflicting political values and preferences? Different groups wish to use the powers of the state for conflicting purposes; different citizens want their elected representatives to pursue incompatible goals. Claiming that the actions of elected representatives or policy outcomes or other features of the world respect "the will of the people" appears suspect, unless this phrase is acknowledged as mere rhetorical decoration for something more prosaic, like the claim that *most* citizens get what they want.

This more prosaic equation, between the will of the people and the will of the majority, is the usual solution to puzzles about popular control, popular rule, and related concepts. “In all countries where the people reign, the majority rules in the name of the people,” as Tocqueville says.¹ Different citizens want different things from their government, but popular control requires only that the majority can get the government to do what the majority wants.

Yet the puzzle about popular control is not resolved so easily. In any group, there are multiple subsets that make up majorities of the whole. In a large body of citizens, none of these majorities is likely to share a “will” in any natural sense of the term. The members of a large subset of the population – majority or otherwise – will share preferences over some, but not all pairs of alternatives. In light of this fact, we should abandon the traditional view of majority rule and popular control, on display in Tocqueville’s quote, which identifies it with the rule of a single majority.

In its place, I will defend a majoritarian conception of popular control according to which something is under popular control if every majority subset of the citizen population has simultaneous control over it. In this chapter, I start with a provisional statement of the view: a majority has control if it can impose its shared preferences on the outcome. First, I explain why we should reject the identification of popular control, or popular rule, with “the rule of the majority.”

¹de Tocqueville (2000, p. 173).

§2.1 Rule by “the” majority

Every group of three or more people contains multiple majorities – multiple subsets of the group whose members make up a majority of the whole. For example, if there were just three citizens, *A*, *B*, and *C*, then four distinct groups would make up majorities of the whole: *A* and *B*, *B* and *C*, *A* and *C*, and the group consisting of everyone. If the rules for reaching collective decisions empower any one of these majorities to make a decision, then they presumably empower all of them.

Yet almost all political theorists who write about democracy and the place of majority rule in it refer to a singular majority. They write of the rule of *the* majority, as if there were only one that mattered. Democracy, Hobbes says, rests on an agreement by each citizen “to submit his will to the will of the majority” (Hobbes, 1998, p. 95). Locke claims that “the majority have a right to act and conclude the rest” when men unite into a political body (Locke, 1988, §95). Tocqueville holds that “the essence of democratic government” is the “sovereignty of the will of the majority” (de Tocqueville, 2000, p. 247). More recent political theory is no different. John Rawls argues that “there is nothing to the view . . . that what the majority wills is right” (Rawls, 1971, §54). Ronald Dworkin characterizes the majoritarian conception of democracy by its insistence on “a match between political decision and the will of the majority” (Dworkin, 2000, 357). Brian Barry, articulating a view about voters’ power over democratic governments, suggests that “where ‘the people’ are divided, . . . the majority should be taken as a proxy for ‘the will of the people’” (Barry, 2002, p. 158).

What fixes the meaning of “the majority”?² Taken out of context, the statement “Democracy is the rule of the majority” is no more informative than “Monarchy is the rule of the man or woman.” Which one? Obviously, references to “the majority” in typical descriptions of majority rule must be shorthand. Monarchy is rule by the individual *who inherits the throne*, or who uniquely satisfies some other identifying condition that no other individual satisfies. Likewise, references to *the majority* must be elliptical references to the majority that uniquely satisfies some identifying condition,

²Some writers acknowledge the ambiguity of “the majority,” but they usually have a different source of ambiguity in mind. For example, Thompson (2002) refers to the “problem of many majorities” while examining a majoritarian principle of popular sovereignty:

But which majority? The identity of the majority is problematic in two dimensions. First, spatially: there is reasonable disagreement about which unit – local, state, national – should be taken as the proper locus of the majority. . . . Second, temporally: there is reasonable disagreement about whether earlier or later majorities should have precedence.

But the assumption here is still that, holding fixed the spatial and temporal coordinates, there is a unique majority whose power is in question: “Because the meaning of the majority will is problematic for these reasons, the principle of popular sovereignty does not once and for all grant *the majority* at any particular level or any particular time final authority. . .” (Thompson, 2002, p. 125), emphasis added.

where the context is supposed to make the identifying condition clear.

One interpretation of the ellipsis is that “the majority” refers to the (largest) majority that shares a common will, that is, the (largest) majority whose members all want the same thing. (We must say *the largest*, because any majority whose members all want the same thing and which is not a minimal majority will contain within itself another majority whose members all want the same thing.) When an author claims that popular control or democracy requires decisions on policy or legislation to conform to the will of the majority, the thought is perhaps that these decisions must conform to the shared will of the largest majority whose members share a common will. Any population contains multiple “latent” majorities, but the specific majority that matters, and to which authors are implicitly referring, is the “actualized” majority that coalesces around some political program or party, united by a shared will to see the program carried out or the party brought to power. For example, Shapiro (2016) credits “critics of majority rule since Kenneth Arrow” with pointing out that “there are many potential majorities – and hence many potential minorities – in any even modestly pluralistic population.” The conclusion he draws from the observation is that “just who is a majority and who a minority is an artifact of the issues that are pitted against one another at any given time” (Shapiro, 2016, p. 50). First this latent majority is actualized, then the other, as the issues on the political agenda change.

Yet there is no reason to expect a large group, including a majority, to possess a shared will when there are at least three alternatives to choose from – at least three feasible policies, at least three relevant ways in which

the government or a public official can act, etc. If a latent majority is “actualized” when its members all want the same policy to be implemented, or all want the government to take the same course of action, then almost no latent majorities will ever become actual. Thus, a variable that is constrained to respect the will of the majority in this sense – one that must conform to what the majority wants *if* there is a majority whose members all want the same outcome – will hardly be constrained at all. A government that must respect the will of the majority, when there is a majority whose members have a common will, can normally do as it likes, because there will normally be no one course of action which every member of a majority wants it to undertake.

To illustrate the point, imagine that an elected legislator is considering whether to draft a bill concerning the role of the federal government in education policy. Specifically, he is considering whether to draft a bill that would expand its role or instead one that would contract it; he could also choose to do nothing. His constituents divide into three groups and have the following preferences over the three alternatives:

Expansionists : expansion \succ status quo \succ contraction

Status quo critics : contraction \succ expansion \succ status quo

Moderates : status quo \succ expansion \succ contraction

Recall that the symbol \succ refers to a preference ordering, e.g., “expansionists” prefer expanding the government’s role to the status quo, and prefer

the status quo to contracting its role.³

Assume that any two of the groups make up a majority. If “the will of the majority” refers to the alternative that a majority all want, then none of the majorities possesses a will. One majority – expansionists and moderates – prefers expanding the government’s role to contracting it; a different majority – expansionists and status quo critics – prefers expanding its role to the status quo. But there is no majority whose members all want the role of the federal government to be expanded: status quo critics would prefer for its role to be contracted, and moderates would prefer for no change at all to be made. Of no majority could one say that expanding the government’s role is its “will.” Only a minority of the population want to see this happen.

Adding more realism to the example would only further underscore the point. There are more than three possibilities when considering the role of the federal government in education, as there are for any complex policy issue. And there are bound to be more than just three types of preferences represented within real-world populations. It is exceedingly unlikely, in most real-world contexts, that a majority of a large population will all want the very same policy, among the virtually infinite possibilities between which elected legislators or government officials can choose. If the “will of the majority” refers to the alternative that each member of a majority wants, then it typically refers to nothing at all.

One context in which a majority *is* likely to share a common will is

³The example is inspired by the survey results reported in Egan (2014).

when a group must choose between just two options, such as whether to approve or reject proposed legislation or which of two candidates for office to elect. When writers describe democracy as “rule by the majority,” they presumably have such a context in mind. They may be thinking of decisions on whether to pass proposed legislation, for example. In democracies, such decisions ought to reflect “the will of the majority”; they ought to reflect whichever preference a majority shares, either for passing or rejecting the proposed legislation. The thought might be that the most consequential decisions are binary decisions with respect to which a majority can be expected to have a common will. This is arguably the most charitable explanation for why so many writers persist in identifying majoritarian democracy with the rule of *the* majority.

Yet this thought – that the most consequential decisions are binary – is wrong. Choices between two mutually exclusive and exhaustive options do not fall out of the sky. They are almost never forced as a matter of logic or necessity. Binary choices are political artifacts. The only thing that compels a choice between just two options on a decision-making body is a person or group with power over the agenda. Legislatures and electorates face binary choices only after some political process has winnowed the space of possibilities down to two. At that point, much of the important action has often already taken place. Whether the final binary decision is taken by majority vote tells us, on its own, almost nothing about the distribution of power. In particular, it does not give enough information to conclude whether there is any plausible sense in which the regime establishes “popular rule” or popular control over collective decisions. Imagine a regime in which

the people's elected representatives vote on proposed legislation according to simple majority rule, but a single unelected and unaccountable executive has total control over which proposals are put to a vote. Or imagine a one-party state in which two-candidate elections are held and decided by majority rule, but the party controls which two of its members are allowed to stand for office. No one could say that the majority of the people "rule," even if the decisions on proposed legislation, or the choices of candidate from those eligible to stand for office, respect the will of the majority.⁴

The conclusion to draw from these observations is that a satisfactory account of popular control and majoritarian democracy will not identify it with the "rule of the majority." If we consider citizens' attitudes towards

⁴Landemore defends a focus on binary decisions by appeal to the actual use of binary agendas in practice and the possibility of normatively acceptable agenda-setting mechanisms, such as one "ensure[ing] the political accountability of the agenda setter to the larger community" Landemore (2012, pp. 149, 150). The problem is that if we try to explain what it would mean for an agenda setter to be accountable to citizens, we will find ourselves referring to the idea that the agenda setter is compelled to take into account citizens' preferences over the entire set of feasible alternatives – just imagine an agenda setter that is free to exclude from the agenda options that an overwhelming majority of citizens strongly prefer to the options he places on the agenda. At bottom, we care whether and how, and under what conditions, citizens' preferences over the entire set of feasible alternatives affect the outcome of the comprehensive process of decision-making.

the entire range of feasible policies, then no majority is likely to share a common will. If popular control over the government and elected officials meant only that they must respect the will of the majority, *when* a majority had a common will, then governments and elected officials subject to popular control would normally be free to do whatever they liked. If instead we consider only binary decisions, such as whether to pass or reject a proposed law, a majority will likely have a common will, but whether the decision respects their common will tells us almost nothing about the democratic character of the larger decision-making process. It will be implausible to identify popular control merely with a constraint on how such decisions are taken, as if the ideal of popular control imposed no constraints on the larger process by which the agenda is formed.

§2.2 Rule by multiple majorities: a provisional statement

What are the alternatives to identifying popular control with the rule of “the majority”? According to the conception of popular control I defend in chapter 5, a decision, elected leader, or any other variable is under popular control if every majority subset of the citizen population has control over it. The critical question is what it means for a group, such as a majority, to have control over something. In this chapter I consider a provisional answer to this question, and, for the sake of concreteness, stick to the case of popular control over a policy decision (as opposed to just any generic variable).

Candidate definition of ‘control.’ A decision is under a group’s control if the decision is constrained to respect the

group's shared *preferences*: whenever each member of the group prefers an option x to an alternative option y , option y is not chosen.

By this definition, a majority – meaning a group whose members make up a numerical majority of the larger reference group – has control over a decision if their shared preferences constrain the decision taken, in the specific sense that an option y is not taken whenever each member of the majority prefers an alternative option x to y . Provisionally stated, the multiple majorities conception of popular control identifies it with every majority's having control in this sense.

A brief remark on terminology may be helpful before proceeding. When I say that an agent prefers one option to another, all I mean is that, given a choice between the two, the agent would choose the first over the second.⁵ When I say that the members of a group (such as a majority) have a shared preference over a pair of alternatives, I mean that each member of the group would be disposed to choose the same option given a choice between the two. So the provisional account of control is that a majority has control over a decision if, whenever its members would all be disposed to choose one option x over another y , given a hypothetical choice between x and y , then *that* fact about their dispositions constrains the decision actually taken, in the sense of preventing y from being chosen. Saying that an agent, or each member

⁵For discussion of how the term *preference* should be understood in economics and related disciplines, see Broome (1999, p. 4) or Pettit (1991, p. 153).

of a group of agents, has certain preferences is a way of describing these “choice dispositions,” but it implies nothing about the reasons that would inform the hypothetical choices. I emphasize this definitional point because the term *preferences* is sometimes used to refer to unprincipled motivations, mere “tastes” or “inclinations.” But here – and in social choice theory and game theory generally – the term should be treated as a technical term of art without any such connotations. An agent’s preferences – her dispositions to make certain choices – might reflect her respect for moral principles, altruistic concern for others, narrow self-interest, or (most plausibly) some combination of reasons of different kinds.

Given this definition of preferences, it is clear why the concept would appear in a definition of control. If an agent has perfect control over a decision, then either the agent must have an ability to make the decision directly or it should be *as if* the agent had an ability to make it directly. In the case of perfect control, that the agent would, hypothetically, choose one option over another – equivalently, that the agent prefers one option to another – should suffice to prevent the latter option from being chosen, just as if the agent were actually presented with the choice.

This account of popular control avoids the shortcomings of the traditional view described above. The members of a large group are unlikely to all be disposed to choose the same option, given a choice from the entire set of feasible alternatives. In that sense, they are unlikely all to want the same option – or to share a common “will” regarding the choice from this set. But there will likely be some pairs of options from which they are all disposed to choose the same thing – in other words, they are likely to share

preferences over some pairs of options.

To illustrate, take again the example involving education policy from the previous section. No majority wants to adopt the same education policy, but a majority – expansionists and moderates – prefers expanding the federal government’s role to contracting it. Another majority – expansionists and status quo critics – prefers expansion to the status quo. No majority prefers an alternative to the policy of expansion. In the jargon of social choice theory, expansion is a *Condorcet winner*, an alternative such that no majority prefers a different alternative to it.⁶

Now suppose that the majority comprising expansionists and moderates can compel the government to respect their shared preference for the policy of expansion over contraction, and the government does not choose the latter; while the majority comprising expansionists and status quo critics can compel the government to respect *their* shared preference for expansion

⁶This is one of two possible definitions of the term, which are equivalent if citizens have complete preferences but can come apart otherwise. Let $x \succsim y$ mean that a majority considers the alternative x at least as good as the alternative y , and let $x \succ y$ mean that a majority strictly prefers x to y , i.e., considers x strictly better than y . On our definition, an alternative w is a Condorcet winner if and only if there is no alternative z such that $z \succ w$. On the other definition, w is a Condorcet winner if and only if, for every alternative z , $w \succsim z$. The definitions are equivalent if for any pair of alternatives x and y and any citizen, she either considers x at least as good as y or considers y at least as good as x .

over the status quo, so the government is unwilling to abide the status quo. It seems plausible to say of either majority that it has some control over the government's choice of policy in such a case. But expanding the government's role in education does not reflect "the will of the majority"; only a minority want to see this policy chosen. Instead, it respects every majority's pairwise preferences: no majority prefers an alternative policy to it. And if that outcome results not from dumb luck, but because the government is constrained to respect majorities' shared preferences, then one might reasonably describe each majority as having some control over the government policy, just as the candidate definition invites us to do.

Not only is this description intuitively plausible, but it yields a more demanding account of popular control than a formulation in terms of the will of the majority. In this hypothetical example and in the real world more generally, popular control would be a trivial constraint on choices from menus containing three or more options if all it required was respect for the "will of the majority" when a majority all want the same option. But if popular control over a decision requires instead that every majority's shared preferences constrain the decision, then it will represent a meaningful requirement, just as it does here, where it forces the selection of the unique Condorcet winner. Indeed, the problem with our candidate definition of majority control, as I explain in the next chapter, is that it is *too* demanding, because Condorcet winners may fail to exist.

I have so far been concerned to explain the focus on shared *preferences* rather than a shared *will*. But why identify popular, democratic control with every *majority's* having control in this sense? What is so special about

majorities? I comment below on the reasons for this move. First, let me point out the connection between this characterization of popular control and a family of models of democracy that have gotten extensive use in political science and economics. Doing so should generate more intuition for what it could mean for multiple majorities to have simultaneous control.

Median voter models of elections and legislatures identify conditions under which one would expect the shared preferences of every majority to constrain collective decisions. These models deserve more attention in philosophical treatments of democracy than they have received, although not because they are good accounts of actual democratic practice. Whatever their value as descriptions or explanations of actual democracies, they are useful aids to thinking through conceptual and normative questions about the power and control that majorities – *all* majorities – might have in the case of a hypothetical, majoritarian democracy in which nothing stood in their way. Actual democracies do not conform to this ideal type, but if one wishes to determine, say, whether constraints on the power of majorities are justified in actual democracies, or whether one can cite the inadequacy of existing countermajoritarian constraints in the diagnosis of their pathologies, one needs to have a coherent account of what it would mean, hypothetically, to fully empower majorities to do as they like. Median voter models are a useful point of departure for this thought experiment.

The most famous is Anthony Downs's (1957) model of electoral competition.⁷ Suppose that two political parties must decide where to locate

⁷The economist Harold Hotelling, in a passing comment on the wide applicability of the ideas captured by his spatial model of duopolistic competi-

themselves on a Left-Right ideological spectrum. As they are choosing from a *continuum* of positions, it would be unreasonable to expect any of the possible positions to correspond to the “will of the majority.” Under any remotely plausible assumption about voters’ attitudes, none of the possible locations on the spectrum will be a location that a majority of the citizens “want” the parties to adopt – not if “wanting the parties to adopt a position” means preferring that they adopt it rather than one of the other positions they could adopt instead. The ideological continuum on which the parties locate contains infinitely many possibilities, because any possible position shades off, by almost imperceptible degrees, into slightly more left- and right-wing positions. With even minimal diversity in voters’ attitudes, one can be sure that no large group of voters, certainly no majority, will all agree about where on the spectrum the parties should locate themselves.

What might determine the parties’ positions if there is no “will of the majority” to which they can respond? Suppose each voter has her own ideological position, a position where she would most prefer for the parties to locate themselves, and if the parties take different positions, she prefers whichever party has adopted the position closest to her own. If voters’ ideological positions are distributed continuously along the spectrum, and the share of voters to the left of any given position strictly increases as one considers more and more conservative positions, then there will be a ¹tion, had already articulated the core idea of Downs’s model several decades earlier (Hotelling, 1929, p. 54). Accordingly, some writers refer to it as the Hotelling or the Hotelling-Downs model of electoral competition.

unique median ideological position: the position of the “median voter” who is ideologically to the left of half of the electorate and to the right of the other half. Take any position to its right and a majority will prefer the median voter’s ideological position to it; take any position to its left and a (different) majority will prefer the median voter’s ideological position to it as well. The median voter’s position will be the unique Condorcet winner among the infinitely many positions along the ideological spectrum.

Under various assumptions about the parties, what their goals are, what they know about the distribution of voters’ ideological positions, and the behavior of voters, one would expect the parties to locate themselves at the position favored by the median voter.⁸ Suppose, as Downs did, that each party’s only goal in adopting a platform is to win the upcoming election, and every voter always votes for whichever party has adopted a position closest to her own ideological position, if they are not equally close, and votes for each with some probability if they locate equally close to her position. Then it can be easily shown that if either party fails to adopt the median voter’s position, it is failing to choose the position that maximizes its probability of winning, given the other party’s position.

Whatever its value for social science,⁹ Downs’s model is useful as a

⁸The conclusion that in a genuinely majoritarian democracy electoral competition would tend to produce results favored by a median voter holds under a wider range of conditions than those assumed in Downs’s model. For some examples, see Gehlbach (2013, chs. 1, 2).

⁹The model’s “prediction” of convergence does not fare well empirically (Snyder and Ting, 2003; Lee, Moretti and Butler, 2004; Fowler and Hall,

thought experiment for political theorists. Abstract away from various circumstances of real-world democracies that might explain why they fail to empower all majorities – a lack of political competition, uninformed voters, the role of money in campaigns, and so on. Posit deliberately unrealistic conditions that would seem conducive to majorities’ empowerment: fully informed and rational voters and two parties that articulate clearly defined political platforms, slavishly cater to voters’ preferences, and will implement whatever platform they campaigned on. If elections are ever capable of empowering all majorities to shape collective decisions, these would appear to be the optimal conditions for them to do so. So it is fruitful to ask, of the fictional world described by the model, whether and in what sense majorities are empowered.

In this imagined democracy, the effect of electoral competition is that the position of the winning party conforms to majorities’ preferences: if a majority prefers one possible position along the ideological spectrum to a second, then neither party will adopt the second position. This observation confirms our idea, expressed in the candidate definition of control above, about how to understand what it means for all majorities to control democratic decisions, without invoking the dubious concept of the will of the majority. Each ²⁰¹⁶, but this “failure” in no way undercuts the model’s usefulness for social science. Just the opposite. The model grounds an inference from the observation of divergence to the conclusion that one or more claims about electoral competition – those claimed captured by the model’s assumptions – is false.

majority's power to determine which party wins ensures that the winning party respects its preferences over the possible platforms. No majority has a common will, but each majority has shared preferences over the possible platforms. As each party wishes to win the election, it is compelled to adopt a platform that respects every majority's shared preferences. The parties are compelled to adopt the Condorcet winner—the position such that no majority prefers any other position to it. In other words, each majority's shared preferences constrain the parties' choice of platform, in equilibrium. Thus, in Downs's model, all majorities have control over the location of the winning party's position according to the candidate definition of majority control.¹⁰

The connection between Condorcet winners and outcomes that respect all majorities' preferences applies more generally, even when outcomes are not positions on an ideological spectrum and no median voter exists. A unique Condorcet winner may exist even if no median voter exists.¹¹ In

¹⁰Some readers may disagree that the parties are *compelled* to adopt this position. They are in fact free to choose differently, albeit at the cost of losing the election with certainty. In what sense then are they “constrained” to decide in accordance with majorities' preferences? Should we deny that majorities control the location of the winning party's position if the parties are free to choose any position? See the discussion in section 4.4 on whether control must be robust to variation in the dispositions of other agents.

¹¹As an example of this possibility, interchange *expansion* and *contraction* in the preference ordering of the expansionists from the example involving federal funding for education. After this modification, preferences fail to be

such cases, the Condorcet winner will be the unique policy that respects every majority's preferences. So we may conclude that, in general and not just in the context of median voter models, the effect of all majorities' having control, in the sense of the candidate definition, would be that a Condorcet winner among the feasible policies is chosen, a policy such that no majority prefers a different policy to it.

A virtue of the definition is that it recognizes the plurality of majorities. It is a first attempt to articulate a sense in which each of them can exert control over policy at the same time. It avoids the error of positing a singular, monolithic majority with a common will. But even if the multiple majorities conception of popular control improves upon the familiar view that identifies it with the rule of "the" majority, why, in particular, favor a "majoritarian" conception of popular control, whether it is formulated with reference to the will of the majority or the shared preferences of multiple majorities?

One familiar strategy for answering this question would be to appeal to May's (1952) theorem or related characterizations of majority voting (Dasgupta and Maskin, 2008). These results are normally interpreted as statements about voting rules, but they arguably capture appealing features of majoritarianism more generally, not just as a quality of voting rules but of the democratic process conceived more holistically. In the case of binary decisions, political theorists argue that simple majority rule uniquely single-peaked and no median voter exists, but *contraction* is a Condorcet winner.

honors normative commitments to equal respect for individuals' judgments (Waldron, 1999; Scwhartzberg, 2014). However, there is a different kind of argument for a majoritarian interpretation of popular control. It appeals to minimal egalitarian commitments, but otherwise bears little resemblance to the usual arguments for majority voting.

The first assumption is that if a decision is under popular control, then there is some group of citizens who have control over it. Perhaps more than one group has control over it; perhaps each individual – a group containing one member – has control over it. But we must be able to identify at least one group of persons with control over the decision. If *no* group of citizens has any control over it, then it cannot plausibly be described as under popular control.

The second is that we are concerned with the idea of popular, *democratic* control, which is to say an *egalitarian* system of control. Political equality requires that “citizens be equal in the control they exercise over the decision-making process” (Christiano, 1996, p. 3). So if popular control implies that a certain subset of the citizen body has control over a decision, then presumably it must imply that every other subset with just as many or more members also has control over it. Thus, our candidate answers to the question of who has control over a decision, when it is under popular control, will be all answers of the form: *every group with k or more members*, and the different candidates will specify different values for k . Put differently, the candidate answers are (1) every sufficiently large minority, (2) every majority, or (3) every sufficiently large supermajority. Answers (1) and (3) are actually families of answers, varying in the interpretation of “sufficiently

large.” An extreme version of (1) is that a decision is under popular control if and only if every individual citizen (the smallest minority) has control over it. At the other end of the spectrum, an extreme version of (3) is that a decision is under popular control if and only if the entire citizen body (the largest supermajority) has control over it.

So why favor the middle option, identifying popular control with every majority’s having control? A comparison with the first option has to wait until chapter 5, once we have a better account of what it means for a group to have collective control over something. But let me explain why we should reject the third family of answers.

Answers in the third family result in conceptions of popular control that, much like the standard equation of popular control with the rule of the majority, impose almost trivial requirements. Imagine a hypothetical regime in which officeholders must win elections in order to hold their positions, but an extreme incumbency advantage is built into the electoral system: when a sitting officeholder competes in a bid for reelection, he wins the election unless every voter prefers his challenger. According to the candidate definition, the entire citizen body would have collective control over who holds office: if they all preferred the challenger to the incumbent, then the incumbent would be thrown out of office. But this collective control is, practically speaking, insignificant. It is extremely unlikely that all citizens will ever share the same preference over the incumbent and challenger. If unanimity rule were replaced with a supermajority rule, the problem would be less acute but still persist to a worrying degree. Generalizing from this example, we see that the third view about which groups must have control

suffers from a problem similar to the standard equation of popular control with the rule of the majority. When officeholders or collective decisions are merely under the control of supermajorities, they are less constrained than they would be if they were *also* under the control of simple majorities. (Of course, if all majorities have control, then supermajorities also have control.) If it is to generate meaningful constraints on officeholders and other variables, a conception of popular control must list majorities' control as one of the necessary conditions for popular control.

To see which features of the candidate definition of control are driving this conclusion, it may help to express the definition as a pair of assumptions. The first is that a group has control over a variable only if the variable is constrained to respect the group's preferences, where respecting the group's preferences means that if the group prefers one option to another, then the latter is not chosen. When are we entitled to say that a group prefers one option to another? The second assumption is that a group prefers one option to another if, but only if, every member of the group prefers the first to the second. We are, in other words, never allowed to impute a preference over two options to a group if any members of the group hold conflicting preferences over the pair. The two assumptions together lead to the conclusion that if the entire citizen body has collective control over an outcome but no smaller subset has control over it, then the outcome is in effect subject to no meaningful constraints at all. For the entire citizen body will rarely have any preferences – its members will rarely all have the same preference over any pair of options. And so respecting the shared preferences of the entire citizen body – its being the case that *if* everyone

shares a preference over two options, then the less preferred option is not chosen – becomes a trivial constraint.

Call the second of the two assumptions—that one can attribute a preference to a group if and only if every member of the group shares the preference—*Pareto extremism*. Note that it reflects an even stronger use of the unanimity criterion than the conditions with similar names (e.g., “the Pareto condition”) that show up in social choice treatments of preference aggregation. The usual conditions are offered as sufficient conditions for a group preference: if every member of a group prefers one option to another, then the group prefers the first to the second. What I am calling Pareto extremism puts this condition forward as not just sufficient but also necessary.

Pareto extremism is probably too strong a view. Following its strictures consistently would require drastic revisions to how we normally talk about groups. But it is worth seeing whether we can develop a satisfactory account of popular control that respects these strictures, and what such an account would look like. Skeptics of popular control, popular rule, and related ideas like to embarrass the standard formulations of these ideas by pointing out the diversity of individual opinions and preferences that standard formulations ignore or distort – “evidently the will of the majority is the will of the majority and not the will of ‘the people’ ” (Schumpeter, 1942, p. 272). If an account of popular control never requires attributing a preference to a group except when every member of the group holds that preference, then it will not be vulnerable to this criticism. That is one virtue of the multiple majorities conception of popular control.

By accepting Pareto extremism, we are, in a sense, already shifting the terms of the debate between Riker and his interlocutors. In a stylized version of that debate, one side starts with the view that popular control over a decision (or popular rule or popular sovereignty) requires that the decision respect “the will of the people.” If an option corresponds to the will of the people, then the people must prefer it to all others, or at least consider it no worse than any other option, so one is left trying to explain the sense in which the people can prefer one option to another even though not every citizen shares such a preference. Skeptics like Riker (1982) argue that no such explanation is convincing in light of Arrow’s impossibility theorem, while others, like Mackie (2003) or Saari (2003), argue that an aggregation rule like the Borda count could plausibly be taken to identify the will of the people, notwithstanding its violation of Arrow’s independence axiom. On the approach I am taking, by contrast, the goal is not to identify a plausible interpretation of the popular will or any basis for aggregating disparate individual preferences into “collective preferences.” Throughout the book, anytime a group is said to prefer one option to another, that means merely that every member of the group holds that preference.

So the provisional account of popular control to which this approach has led us does away with the concept of a popular will and, indeed, even with the concept of the will of the majority. Those are its virtues. Unfortunately, it is untenable, as provisionally stated. Social choice theory explains the source of its problems, even though its problems have nothing to do with the aggregation of individual preferences into collective preferences. That explanation can be found in the next chapter, but the argument so far

already brings the contours of the problem within sight.

To see the difficulties, first consider why one would not want to identify popular control with every minority's having control – not, at least, if one adopts the candidate definition of control. As a simple example, imagine a choice between two options, a and b . Assume they exhaust the possibilities, so that, as a matter of logic, one of them must be chosen. (Perhaps option a is “do nothing.”) Given the candidate definition of control, if two different minority groups each have control over the choice, and the members of one group prefer a and members of the other prefer b , then neither option a nor option b can be chosen, contradicting our suppositions. One of these minorities must lack control according to the candidate definition. Thus, if popular control means that every minority of a certain size has control over a decision, and the candidate definition of control is adopted, then popular control will typically be impossible.

It turns out that an analogous impossibility arises if popular control means that every majority has control per the candidate definition. Citizens' preferences might be so arranged that no Condorcet winner exists. That is to say, their preferences might be such that no matter which option is chosen, the choice will fail to respect the shared preferences of a majority. Given the candidate definition of collective control, we would then have to conclude that some majority lacks control over the decision. Thus, if popular control means that every majority has control, popular control will be impossible in such circumstances.

To appreciate this ‘impossibility’ result, we need to familiarize ourselves with insights from social choice theory, and get clear on just the kind of

problem they raise for the idea that majorities can have control over democratic politics. The problem I am describing differs from other putative problems that social choice theory is often thought to raise. It has nothing to do with the notion of a popular will, an idea that we have in effect already abandoned by embracing Pareto extremism, as explained. The problem also has nothing to do with the instability or “chaos” that cyclical majorities are sometimes thought to threaten. The next chapter explains just what the problem is, and why it forces us to rethink the provisional account of popular control sketched in this chapter. After revising the candidate definition of control, in chapter 4, we will be able to put the multiple majorities view on solid footing, saving it from this impossibility result. The revised view, defended in chapter 5, will be that a variable is under popular control if every majority has control over it, where having control means that a group’s shared preferences constrain the variable if its members attach sufficient importance to the variable.

Chapter 3

An objection from social choice theory

The candidate definition from the previous chapter says that a majority has control over policy if its shared preferences constrain policy, that is, if it can prevent the choice of a particular policy x whenever its members all prefer a different policy y to x . The implication is that if *every* majority has control over policy, then the chosen policy has to be a Condorcet winner, a policy such that no majority prefers an alternative to it.

But there may not be a Condorcet winner. It might be impossible to satisfy the preferences of every majority. In such a situation, it would be impossible to give every majority control, according to the candidate definition. Some writers contend that such situations occur in practice, even frequently. The candidate conception of majority control must therefore address an objection that what it characterizes as popular control is frequently impossible. In this chapter, I present the objection and explain why political theorists' typical responses to social choice theory do not address it.

§3.1 Condorcet's paradox

Consider an example in which the government is considering whether to reduce, maintain, or increase funding for a program. To keep the example simple, we may imagine that two alternatives to the status quo are in discussion: eliminating funding entirely or doubling funding. Citizens divide into three groups, with the following preferences over the three alternatives, expressed as percentages of current funding.

Liberals : $0\% \succ 100\% \succ 200\%$

Conservatives : $200\% \succ 100\% \succ 0\%$

Centrists : $100\% \succ 0\% \succ 200\%$

In this example, the status quo (100 percent of the current level), satisfies the preferences of every majority; no majority prefers a different level to this one. Among the three options, this option is a Condorcet winner. If every majority has control over policy in the sense of the candidate definition, then the result will be that neither alternative to the status quo is adopted.

Now imagine citizens' preferences are different. Liberals become less sanguine about the efficacy of the program as it is currently run, and conclude that, while eliminating the program would be best, a doubling of current

funding would be better than maintaining the status quo:

$$\begin{aligned}
 \text{Liberals : } & 0\% \succ 200\% \succ 100\% \\
 \text{Conservatives : } & 200\% \succ 100\% \succ 0\% \\
 \text{Centrists : } & 100\% \succ 0\% \succ 200\%
 \end{aligned}
 \tag{3.1}$$

(The liberals are now like the ‘status quo critics’ from the example in the previous chapter; their preferences are “double-peaked,” a phenomenon that Egan (2014) verifies empirically with survey data.)

A problem now arises if we suppose that the determination of the funding level is subject to majority control. No matter which of the three alternatives is chosen, a majority will prefer one of the other options to the chosen alternative. No Condorcet winner exists. Thus, the choice will fail to be under the control of one of the three majorities. If every majority has control—in the sense of the provisional definition from the previous chapter—the level of funding cannot be 100 percent, because liberals and conservatives both prefer 200 percent. It cannot be 200 percent, because liberals and centrists both prefer 0 percent. And it cannot be 0 percent, because conservatives and centrists both prefer 100 percent.

This configuration of preferences is an example of *Condorcet’s paradox*. As I use the term here, its defining feature is the absence of a Condorcet winner: an option no worse than any other when the comparisons are made according to majorities’ preferences. (Some writers use the term *Condorcet’s paradox* to refer merely to cases like the given example, in which majorities have “intransitive” or “cyclic” preferences over three options. See the

discussion of cycles below.) When no Condorcet winner exists, then, by definition, it is true of every option that a majority prefers a different option to it.

It is worth recalling at this point that we are using the term *preferences* simply to designate an agent's dispositions to choose from hypothetical pairs of options. The statement that an agent prefers one option to another just means that she would choose the former over the second, given a choice between just the two of them. I reiterate this terminological point because some commentators adopt an overly narrow interpretation of the term *preferences* when they write about social choice theory—contrasting preferences with, say, judgments about the common good (Cohen, 1986). As I use the term, an agent's preferences could reflect her judgments about the common good or other ethical commitments; the term describes her disposition to choose from hypothetical pairs of options, not the reasons that would guide her choices. This observation is important because what makes Condorcet's paradox possible, and what produces other difficulties in "aggregating" individuals' preferences, is the fact that preferences have the structure of a *ranking* of the alternatives. Social choice paradoxes do not derive from the connotations of the word *preference* in ordinary language, where it often suggests self-regarding concerns or mere tastes and predilections, in contrast to weighty ethical priorities. Use any language you like to express facts about how each individual would choose from different pairs of options, and you can construct a scenario like Condorcet's paradox.

If Condorcet's paradox occurs on a regular basis, then popular control—according to the provisional, majoritarian conception from the previous

chapter—is often impossible. It will be useful to state the objection in the form of an argument.

1. If policy is under popular control, then policy respects majorities' preferences.
2. When citizens' preferences over policies have the structure of Condorcet's paradox, every policy fails to respect the preference of some majority.
3. Citizens' preferences over policies frequently have the structure of Condorcet's paradox.
4. Thus, policy is frequently not under popular control.

The second premise is true by definition—*Condorcet's paradox*, as just defined, refers to a situation in which, among all the feasible options, none is a Condorcet winner. A Condorcet winner is an option that is no worse than any other in the eyes of a majority. If citizens' preferences among the government's available policy options are such that no option is a Condorcet winner, then it follows that no matter which policy is adopted, it will fail to respect the preferences of some majority. There will be some majority whose members prefer a different option to it.

Thus, when citizens' preferences have the structure of Condorcet's paradox, there will be no possible institutional means of establishing popular control, as the majoritarian conception construes it. No matter how well informed and vigilant the citizens; no matter how well the media monitor and report on the misdeeds of elected officials; no matter how strongly elected

officials desire re-election and no matter how willing they are to cater to public opinion; and no matter the character of political institutions—the policy will fail to respect the preferences of some majority. According to the candidate majoritarian conception of popular control, it follows that policy will not be under popular control. Popular control will be impossible so long as citizens’ preferences have this structure. Call this the *impossibility objection*.

Since premise 2 is true by definition, there are only two strategies for defending the idea of popular control against the impossibility objection. Either we reject premise 3 and argue that citizens’ preferences do not frequently have the structure of Condorcet’s paradox. Or we reject premise 1 and explain why—contrary to the provisional majoritarian conception of popular control set out in chapter 2—popular control over policy does not imply that policy respects the preferences of majorities.

I claim the first strategy is unpromising. In the context of the typical policy problem, the set of policy options is infinitely large, for all practical purposes. As a consequence, there is no theory-free method of determining citizens’ preferences over all feasible policy options; the only reasonable basis on which to infer these preferences involves *a priori* assumptions about their structure. From reasonable assumptions about their structure, one can prove that Condorcet’s paradox is likely to occur. Thus, premise 3 is hard to refute. I elaborate this argument below.

The second strategy is to explain why the conceptual claim about popular control, in premise 1, should be rejected. This is the strategy I pursue in the next two chapters. I give a more rigorous definition of control in chapter

4, and I then show, in chapter 5, how it allows us to reconcile the possibility of majority control with Condorcet's paradox. Under the definition given in chapter 4, a majority may be said to have control over policy even when its members all prefer a different policy to the actually chosen policy. What is necessary for control is merely that majorities' preferences would constrain the chosen policy if they cared enough about policy relative to other things they care about. Since this account of control compromises the simplicity of the provisional account of majority control, it is important to first see why the alternative strategy of refuting premise 3 does not work. This is the task of the rest of the chapter.

§3.2 Two kinds of "majority cycle"

Condorcet's paradox is sometimes described as an instance of a "majority cycle." Riker uses this term when presenting his critique of populism. Gerry Mackie sets out to refute Riker's critique by "show[ing] that the cycles that are alleged to make democracy meaningless are rare" (Mackie, 2003, p. 17). Anthony McGann argues that "rather than cycling being a problem for democratic theory"—as Riker contends—"it may be that cycling is what makes democracy as we know it possible" (McGann, 2006, p. 26). It is useful to distinguish between two phenomena, each of which might be called an example of a majority "cycle," and which the literature does not always explicitly keep apart.

The first is a situation in which a majority votes for an alternative a in a choice between two alternatives a and c , then another majority votes for

b in a choice between a and b , and then another majority votes for c in a choice between b and c . Call this a *majority voting cycle*. It describes a process (“cycling”) that unfolds in time, and whether it occurs depends on how people vote.

The second phenomenon is a situation in which a majority prefers a to c , another majority prefers b to a , and another majority prefers c to b . Call this a *majority preference cycle*.¹ Whether a majority preference cycle occurs does *not* depend on how anyone votes or otherwise behaves. It depends only on individuals’ preferences. The term describes a property that citizens’ preferences might exhibit in the aggregate.

Recall that a Condorcet winner is an option such that no majority prefers an alternative to it, and Condorcet’s paradox is a configuration of preferences relative to which no Condorcet winner exists. The impossibility objection rests on the premise that citizens’ preferences regularly exhibit the structure of Condorcet’s paradox.

There is no logical connection between majority *voting* cycles and the existence of a Condorcet winner. By this, I mean simply that one cannot deduce whether a voting cycle will occur merely from knowledge of whether a Condorcet winner exists. Whether a Condorcet winner exists depends *only* on citizens’ preferences, not on how anyone votes or otherwise behaves; whether a voting cycle occurs depends on how people vote. But one cannot

¹More generally, a majority preference cycle over k alternatives is a situation where a majority prefers x_1 to x_2 , another majority prefers x_2 to x_3 , . . . , another prefers x_{k-1} to x_k and another prefers x_k to x_1 .

derive conclusions about how people will behave solely from assumptions about their preferences. (That is one lesson from game theory.)

There is a logical connection between the existence of a Condorcet winner and majority *preference* cycles. If no Condorcet winner exists, then there must exist a majority preference cycle.² But the converse is false: there may be a “best” alternative (a Condorcet winner) even if there are majority preference cycles involving other non-best alternatives. To illustrate the possibility, we can modify the example of Condorcet’s paradox given at the start of this chapter by stipulating the availability of a fourth policy option:

$$\begin{aligned}
 \text{Liberals : } & 0\% \succ 50\% \succ 200\% \succ 100\% \\
 \text{Conservatives : } & 200\% \succ 100\% \succ 50\% \succ 0\% \\
 \text{Centrists : } & 50\% \succ 100\% \succ 0\% \succ 200\%
 \end{aligned}
 \tag{3.2}$$

Now cutting the funding level by half is a Condorcet winner: it is preferred by a majority to each of the other options. But there is still a majority preference cycle over the other three options.

People who write about “majority cycles” and the questions the phenomenon raises for democratic theory often seem to have majority voting cycles in mind, because they casually interchange references to “cycles” and instability and dynamic change. For example, Shapiro writes:

The possibility of cycles gives those who lose in any given elec-

²The statement in the text is true when the set of alternatives is finite. With infinite sets, a qualification is needed. See the discussion of maximal elements and acyclicity in Austin-Smith and Banks (1999, chs. 1, 5).

tion an incentive to remain committed to the system in hopes of prevailing in the future, but the fact that cycles are actually rare means that government policies are not perpetually being reversed. (Shapiro, 2003, p. 15)

I should therefore stress that the challenge to the idea of popular control I am addressing has *nothing* to do with majority voting cycles. What underwrites the impossibility objection against the majoritarian conception of popular control is the claim that Condorcet's paradox arises on a regular basis. Given the definition of Condorcet's paradox, that is equivalent to the claim that Condorcet winners regularly fail to exist. From this (premise 3) it follows that, whatever the institutional arrangements in place and whatever the actions taken by citizens and public officials, enacted policies will regularly frustrate the preferences of majorities. Thus, there is no way to establish popular control, given the assumption (premise 1) that policy is only under popular control if the chosen policy respects the preferences of majorities.

Whether voting cycles occur is *irrelevant* to this objection. Suppose that voting cycles never occur. The premises, and thus conclusion, of the impossibility objection might still be true. Suppose voting cycles occur all the time. Any one of the premises, and thus the conclusion, might still be false. In particular, the truth of premise 3—the claim that Condorcet winners regularly fail to exist—does not depend on whether voting cycles occur. Voting cycles might never occur even if these configurations of preferences happen all the time. Thus, establishing that voting cycles never occur does *nothing* to blunt the impossibility objection, because it does nothing

to discredit premise 3. To defend the idea of popular control against the impossibility objection, we must either reinterpret popular control so that it does not imply respect for majorities' preferences, or we must show that these problematic configurations of preferences, albeit logically possible, do not regularly arise in politics.

A more subtle point is that majority *preference* cycles are also not what raises the puzzle about majority control. Majority preference cycles are not per se the problem because they do not preclude the existence of Condorcet winners. As example (3.2) illustrates, Condorcet winners may normally exist even if majority preference cycles happen all the time, in which case the impossibility objection fails.

But since the nonexistence of a Condorcet winner *does* imply a majority preference cycle, evidence that these cycles are rare would be evidence that Condorcet winners generally exist.³ Thus, if one can produce evidence that majority preference cycles are rare, then one will have knocked out the third premise in the argument against the possibility of majorities having control.

Alas, it is difficult to show that Condorcet winners generally exist and *a fortiori* difficult to show that majority preference cycles are rare. If there are only a handful of options under consideration, then we might determine existence by gathering evidence of each person's preference over each pair of options, using a survey or votes or other behavior that might be taken to reveal preferences. One obstacle here is that election ballots and opinion

³The caveat, as mentioned earlier, is that the set of alternatives is finite. See footnote 2.

surveys usually do not give voters or respondents the opportunity to report their preferences over each pair of candidates or options, so the analyst must infer full preferences from the “incomplete” data afforded by election returns, opinion polling, or votes in a legislature.⁴ When Mackie sets out to debunk Riker’s claims of historical instances of cycles, he uses available evidence to impute preferences over each pair of alternatives. For example, Riker claims that the preferences of delegates to the 1846 U.S. House exhibited a majority preference cycle over three alternatives: the status quo; an appropriations bill to raise funds for negotiations with Mexico at the end of the war; and the Wilmot Proviso, which was an amendment to the appropriations bill that would have prohibited slavery in the conquered territories. Riker’s claim is based on conjectures about legislators’ preferences. Even under the assumption that a legislator’s vote always indicates his sincere preference, one can do no more than make reasonable conjectures about the distribution of preferences because the House did not vote over each pair of alternatives. Mackie takes issue with Riker’s conjectures, claiming that he misreads the Congressional record. He maintains that the actual record supports conjectures about preferences which, if correct, imply that there

⁴A noteworthy exception is an opinion poll from 2015, in which each of four Republican presidential candidates was judged worse than another candidate by a plurality of the respondents (self-registered voters who identified as Republican or Republican-leaning). See Kurrild-Klitgaard (2017) for discussion of this and other voting paradoxes from the 2016 Republican presidential primaries.

was no majority preference cycle.⁵

Setting aside the difficulty of making these conjectures, there is a deeper problem with this approach if its goal is to refute the objection to the possibility of popular control, presented above. When considering a decision-making group like the House, even if there was no majority preference cycle over the three alternatives its members voted on, there could have been a cycle over the entire set of feasible bills that the House could have enacted. More to the point—because it is the nonexistence of a Condorcet winner, and not the presence of a preference cycle per se, that supports the objection against popular control—there may have been no Condorcet winner among the entire set of feasible bills, even if, among the three alternatives on the agenda—the status quo, the original appropriations bill, and the amended appropriations bill—one of the alternatives was preferred by majorities to the other two.

We can illustrate the point with another version of our running example, where the question is over the level of funding for a government initiative. Suppose that a legislature is debating changes to the current level of funding for the government program, but the only proposals brought up and voted upon are a proposal to double the current level of funding and one to triple the current level of funding. Within the legislature, liberal, conservative, and centrist legislators have the following preferences over these two proposals,

⁵See Mackie (2003, pp. 241–243); Riker (1982, pp. 223–227).

the status quo (100%), and a fourth option of eliminating the funding (0%).

$$\begin{aligned}
 \text{Liberals : } & 0\% \succ 200\% \succ 100\% \succ 300\% \\
 \text{Conservatives : } & 300\% \succ 200\% \succ 100\% \succ 0\% \\
 \text{Centrists : } & 100\% \succ 0\% \succ 200\% \succ 300\%
 \end{aligned}
 \tag{3.3}$$

Suppose that the option of eliminating the funding (0%) is never formally proposed. It is not the subject of any debate or vote. If an observer asked the question that Riker's and Mackie's empirical inquiries address—was there a majority preference cycle over the status quo and the alternatives to it that were formally proposed and voted upon?—and she had means of reliably answering the question, she would conclude, with Mackie, that there was no evidence of a preference cycle. A majority prefers a level of 200 percent to 300 percent and another majority prefers a level of 200 percent to 100 percent. Thus, doubling the funding level is a Condorcet winner *among the restricted set of options* that the legislators vote on. This set excludes the feasible option of eliminating funding. Within the *entire* set of four options, there is no Condorcet winner: for every alternative, there is another alternative that a majority prefers to it.

Riker's findings of alleged cycles, even if they withstood Mackie's criticisms, would require the same qualification. Observing a cycle over a handful of alternatives that appear on a ballot or legislative agenda does not preclude the existence of a Condorcet winner among the set of feasible alternatives. In example 3.2 (page 66), an observer who only had access to the individuals' preferences over the levels 0 percent, 100 percent, and 200 percent

would conclude that there is a majority preference cycle. But the existence of this cycle is not a reason to conclude that none of the feasible alternatives respects the wishes of majorities: there is a fourth alternative, reducing funding to 50 percent, that majorities prefer to every other alternative.

Should we take a wide or narrow view of the set of relevant options? It depends on the question of ultimate interest. For our purposes, the wide view is appropriate. The impossibility objection derives its force from the assumption (premise 1) that popular control implies respect for majorities' preferences over the *feasible* alternatives, not merely respect for their preferences over the handful of alternatives that elected legislators choose to bring up for a vote or the handful of candidates that political parties choose to nominate. If a large majority opposes an existing practice of racial apartheid, but the officials with the authority to propose new policies never put abolitionist or reformist policies on the formal agenda, then—according to the intuition that the objection exploits—the government ignores the wishes of a majority in just the way that popular control rules out. What intuitively offends against the ideal of popular control here is the presence of a *feasible* policy—abolishing apartheid—which, whether or not it finds its way onto any formal policy agenda, a majority prefers to the government's actual course of action.

The relevant question is therefore whether a Condorcet winner typically exists among the set of *feasible* policies. Establishing that it is rare to observe majority preference cycles over the candidates appearing on a ballot, or over the bills debated on the floor of the legislature—which is the most that Mackie can claim for his survey of the empirical evidence—does not answer

this question. The same goes for other empirical studies of the frequency of majority preference cycles (Regenwetter et al., 2006; Tideman, 2006).

As I explain in the next section, we cannot hope to answer the question unless we are willing to make *a priori* assumptions about the shape of citizens' preferences. The most plausible assumptions support the view that in many contexts, Condorcet winners do not exist and majorities' preferences exhibit cycles.

§3.3 The spatial model of politics

Suppose one wanted to administer a survey to the general population in order to determine whether any level of the minimum wage is a Condorcet winner, relative to citizens' preferences over *all* possible alternatives. It would be a long survey that asked each respondent her preference over each pair of alternatives. Even if one restricted attention to levels within plus or minus \$2 of the current minimum wage and measured wages to the penny, there would be 400 different levels and $\frac{400 \times 399}{2} = 79800$ different pairs of levels.

Plausible assumptions about a person's preferences would allow one to reduce the number of questions on the survey. For example, if one knows that a survey respondent considers a minimum wage of \$7.25 to be better than a minimum wage of \$7.24, and considers a minimum wage of \$7.24 to be better than \$7.23, then one can infer that the respondent prefers a minimum wage of \$7.25 to one of \$7.23, so long as one is willing to assume that preferences are transitive. The assumption is plausible enough. It

reduces the amount of information one needs in order to determine whether a Condorcet winner exists.

One is arguably safe with an even stronger assumption. A person's preferences over levels of the minimum wage are plausibly assumed to be *single-peaked*. This would mean that

- i) there is some level of the minimum wage, call it w^* , which she prefers to every other possible level;
- ii) among any two possible levels less than w^* , she prefers the greater of the two; and
- iii) among any two possible levels greater than w^* , she prefers the lesser of the two.

The assumption of single-peakedness takes its name from the shape of the function that would represent these preferences, if the minimum wage were measured on the horizontal axis and one plotted the desirability—or “utility”—of the different levels on the vertical axis. In Figure 3.1, the function plotted on the left is single-peaked, in contrast to the “double-peaked” function plotted on the right:

The stipulated preferences of the liberals in examples 3.1, 3.2, and 3.3 (pp. 60, 66, 71) do not satisfy this assumption. They were assumed to prefer a funding level of 0 percent of current funding to all other levels, but for two possible levels greater than 0 percent (namely, 200 percent and 100 percent), they were assumed to prefer the greater of the two (violating condition (iii)).

The assumption of single-peaked preferences does wonders for the program of determining whether a Condorcet winner exists. From the assump-

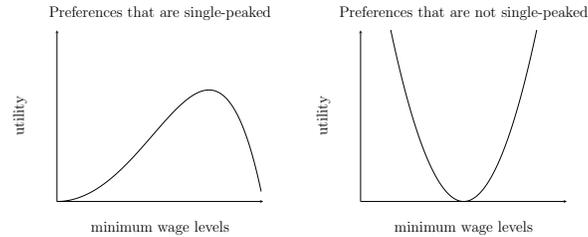


Figure 3.1: The figure on the left displays the graph of a single-peaked utility function; the figure on the right displays the graph of a utility function that is not single-peaked.

tion that each citizen’s preferences are single-peaked, it follows that a Condorcet winner exists. In the case of an odd number of citizens, the Condorcet winner will be the most preferred policy of the “median voter,” whose most preferred policy lies in the middle of the distribution of most preferred policies, with half of her peers preferring minimum wage levels less than this and half preferring levels greater than this. This proposition is an application of the median voter theorem, first proved by Black (1958).

An assumption like transitivity reduces the amount of information that must be gathered in order to determine whether a Condorcet winner exists. The assumption of single-peaked preferences obviates the need for empirical data altogether.⁶ It may feel almost like cheating. But we will have to “cheat” like this if we want to come to any conclusion about whether, among a collection of several hundred policy alternatives, there is one that

⁶Determining *which* policy is the Condorcet winner would necessitate some empirical inquiry, as it would require identifying each citizen’s most preferred policy.

no majority would judge worse than any of the hundreds of other options.

The different levels of the minimum wage or different levels at which to fund a program can be represented by points on a line; they belong to a “policy space” that has a single dimension. But political communities often make choices that cannot be described this way. Consider the choice of an appropriations bill specifying funding levels for multiple government programs. It can make sense to link the determinations of the various funding levels together as part of a single choice. Viewed this way, the different possible bills that could be passed differ along multiple dimensions, one dimension for each program; if there are k programs, then each bill can be thought of as the specification of k numbers, i.e., as a point in a k -dimensional space.

The multidimensional spatial model, as this setup is called, has been the object of extensive study (Le Breton, 1987; McKelvey, 1979; McKelvey and Schofield, 1987; Plott, 1967; Saari, 1997). It turns out that if we represent policies as points in a multidimensional space and make minimal assumptions about the structure of individuals’ preferences over these policies, then we can prove that *Condorcet winners almost never exist*.⁷ The sense in which they will almost never exist is this: any configuration of preferences supporting a Condorcet winner is almost identical to infinitely many similar configurations that fail to support a Condorcet winner. A configuration of preferences supporting a Condorcet winner will be like a pencil balanced on

⁷This implication is often stated with a different piece of jargon: the majority preference relation is said to have an “empty core,” which just means that no Condorcet winner, as defined here, exists.

its point, vulnerable to the slightest perturbation.

The theorems to this effect are mirror images of the median voter theorem. Their *a priori* assumptions about preferences are no more objectionable than the *a priori* assumption of single-peaked preferences in the minimum wage example.

The proofs of the theorems require advanced mathematical tools, but the intuition behind them is fairly simple. To say that no Condorcet winner exists among a set of feasible policy bundles is just to say that, no matter the status quo policy bundle, there is always some majority that could in principle come together and compromise on an alternative that its members would all consider superior to the status quo. Thus, in order for the status quo policy to be a Condorcet winner, there must be no majority whose members can find enough common ground to identify a mutually agreeable alternative. *Thus, Condorcet winners should be unlikely in settings that facilitate compromise.* The multidimensional spatial model is just such a setting. Since policies differ along multiple dimensions, people can compromise by yielding ground on the dimensions they care relatively less about in exchange for concessions on the dimensions they care more about. A liberal who prefers a higher minimum wage and more generous unemployment benefits and a conservative who prefers lower levels of each might both view a compromise that reduces one and increases the other as an improvement over the status quo. The more dimensions there are, the more opportunities there are for “ideological trade,” where people compromise on some dimensions in exchange for improvements on others. And since policies vary continuously, they can look for agreement on incremental changes,

which further facilitates compromise. It should be easier to agree on an incremental reduction in the minimum wage and an incremental increase in unemployment benefits than to agree on drastic changes to each.

To provide further intuition, we can use figure 3.2 to visualize policies in a two-dimensional plane and graphically display individuals' preferences. Suppose the problem is to choose spending levels for two programs, so that each policy bundle, or budget, can be represented as a pair of positive numbers, *i.e.*, as points in the northeast quadrant of the two-dimensional plane. We will make an assumption about preferences that is stronger than what the theorems require but helps to illustrate the basic intuition behind them. The assumption is that individuals have “Euclidean preferences,” which means that each has a most preferred budget and, for any two alternative budgets, prefers whichever is closest to her ideal budget.⁸ For example, if the point z_1 in the figure is the ideal budget of someone with Euclidean preferences, then she prefers all of the points (all of the possible budgets) in the interior of the indicated circle to the point x (which represents another budget), because they are all closer to her ideal point than x . The points z_2 and z_3 in the figure represent the ideal budgets of two other people.

The shaded regions represent the budgets that at least two of the three people—a majority—prefer to x . Clearly, x is not a Condorcet winner: every point in the shaded region is a budget that a majority prefers to x .

With a moment's reflection, it should be intuitive that there is nothing

⁸Notice that in the case of a single-dimensional policy space, Euclidean preferences are examples of single-peaked preferences.

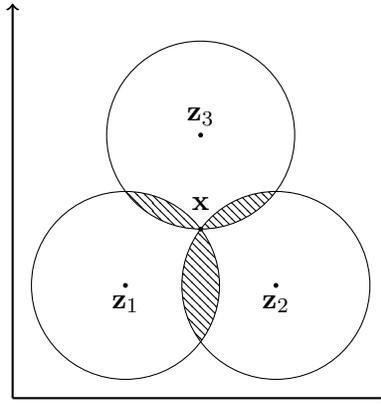


Figure 3.2: Every point in the plane represents a policy—such as the choice of a budget that specifies spending levels for two programs—and the points z_1 , z_2 , and z_3 represent the ideal policies of persons 1, 2, and 3, respectively.

special about the choice of x . Had we picked any other point in the plane and then drawn three circles through it, centered at z_1 , z_2 , and z_3 , we would have gotten the same result. At least two of the circles would intersect, indicating alternatives to this chosen point that a majority consider superior to it. With this configuration of preferences, no Condorcet winner exists.

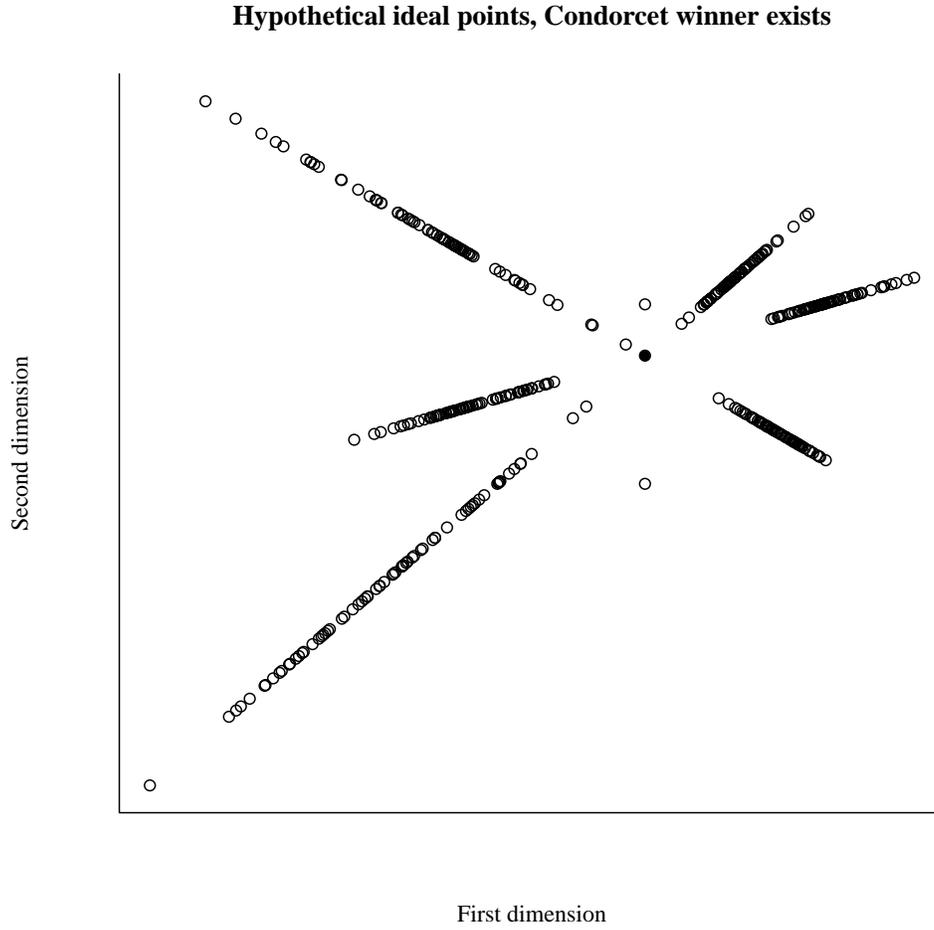


Figure 3.3: 435 hypothetical ideal points, arranged in a pattern that exhibits “radial symmetry.” In this context—two dimensional policy space, odd number of individuals with Euclidean preferences—radial symmetry is necessary and sufficient for the existence of a Condorcet winner.

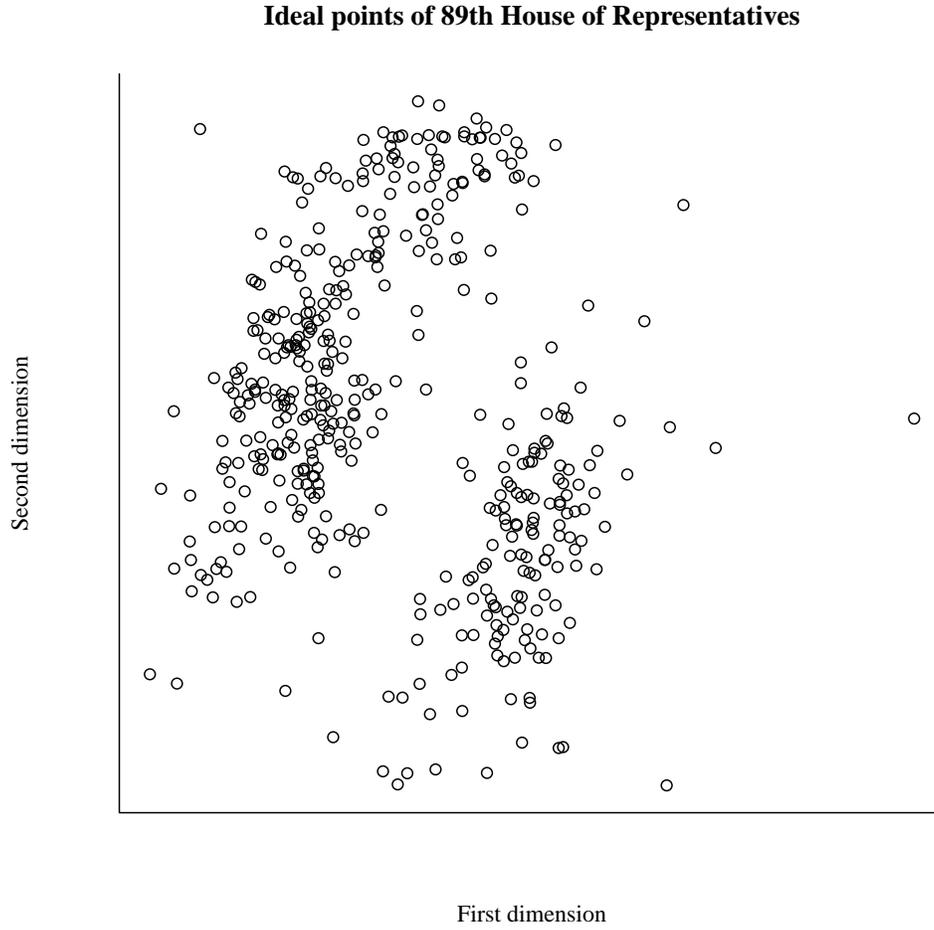


Figure 3.4: Poole and Rosenthal’s NOMINATE estimates of ideal points in the 89th U.S. House of Representatives. The ideal points do not exhibit radial symmetry, so no Condorcet winner exists.

A Condorcet winner could exist if individuals’ ideal points were different from those specified in the figure. But their ideal points—the locations of the most preferred budgets in this space—would have to be distributed in a special and highly improbable pattern. When there are just three persons,

their ideal points would all have to fall on a single line. That means that with the slightest perturbation to the pattern, one of the points will no longer lie on the same line as the other two and a Condorcet winner will no longer exist.

The necessary pattern is harder to describe when there are more than three persons, but a picture conveys the intuition. Figure 3.3 gives an example of a configuration of 435 ideal points that satisfy the necessary pattern. The solid dot in the center of the distribution, which would be some person's ideal point, is the Condorcet winner.⁹ The reader will notice the striking and improbable symmetry of the other points around this central ideal point. This form of "radial symmetry" is a necessary (and sufficient) condition for a Condorcet winner to exist when there are more than two dimensions and individuals prefer policies closest to their ideal points; if the number of in-

⁹Notice that every other ideal point falls on one of five lines that pass through this point. On each of these imaginary lines, the Condorcet winner is the median in the distribution of ideal points falling on the line. As a consequence, if you divide this space into two halves by drawing *any* line l through the Condorcet winner, no more than 217 points will lie in either half of the space (excluding those that lie on l itself). Thus, no more than 217 people, a minority of 435, will wish to move policy from the Condorcet winner in a direction that is orthogonal to the line l , given the assumption of Euclidean preferences. Since this is true of any line l drawn through the Condorcet winner, the result is that no majority wishes to move policy away from the Condorcet point.

dividuals is odd, then it is necessary and sufficient for a Condorcet winner to exist whenever there is more than one dimension (Plott, 1967).

The number 435 was chosen for this demonstration because it is the number of legislators in the U.S. House of Representatives. Political scientists have methods of estimating the ideal points of members of Congress, so we can compare this hypothetical example of how preferences would have to be configured in order to support a Condorcet winner with the actual preferences of House members, as estimated by these methods. Figure 3.4 displays Keith Poole and Howard Rosenthal's estimates of the ideal points of members of the 89th House of Representatives (1965–1967). Clearly, these points do not display the radial symmetry condition that is necessary for the existence of a Condorcet winner.

Of course, it is unsurprising that the distribution of estimated ideal points does not satisfy radial symmetry. For any configuration of preferences that satisfies radial symmetry, there are infinitely many alternative configurations that are nearly identical but do not satisfy radial symmetry. Nearly identical constellations of ideal points must be assumed to occur with nearly identical probabilities. The implication is that the probability of radially symmetric preferences is zero.

Thus, once we make some minimal *a priori* assumptions about citizens' preferences over budgets or other kinds of policies that differ along multiple dimensions, we can conclude that Condorcet winners exist with probability zero. Since there are infinitely many possible budgets, the only way to form any reasonable belief about the existence of a Condorcet winner among the set of feasible budgets is by making *a priori* assumptions about the shape of

individuals' preferences, just as an assumption like single-peaked preferences is necessary to form a reasonable belief about the existence of a Condorcet winner among the set of feasible levels of the minimum wage.

The nonexistence of a Condorcet winner implies the existence of a majority preference cycle, as explained above. How, then, does Mackie reconcile his claim that cycles are rare with the theoretical argument we have just examined? Writing of the theorems on the multidimensional spatial model, Mackie claims that “the theorems... are normatively tangential, their predictions are falsified, and standing alone they are empirically irrelevant for understanding the world of politics” (Mackie, 2003, p. 174). Mackie's subsequent discussion makes it clear that what he means by their “predictions” is a prediction of *voting cycles*—some dynamic, wandering sequence of actual decisions without rhyme or reason. For what he treats as falsification of their predictions are laboratory experiments in which participants play a majoritarian decision-making game, but “stability” and “equilibrium” are observed.

But the theorems themselves do not predict voting cycles. In general, mathematical theorems do not, literally speaking, *predict* anything; they are demonstrably true propositions about mathematical objects, not testable predictions about future or hypothetical events. Moreover, the theorems on the multidimensional spatial model are not, strictly speaking, statements about *voting cycles*, even though writers gloss them as such. They are not statements of the form *If such-and-such, then there will be voting cycles and no stable outcome to majority voting*. Rather, they are of the form *If such-and-such, then a Condorcet winner almost never exists* (or, to use

the jargon with which they are more often stated, the core of the majority preference relation will almost always be empty). They state assumptions about the space of alternatives and preferences under which, for almost all configurations of preferences, it will be true that no matter which alternative one considers, there will be another option that a majority prefer to it.

If this conclusion about the generic nonexistence of Condorcet winners (or the generic “emptiness of the core”) further implied *voting* cycles, then observing stability in experimental or real-world majority rule decision-making would constitute evidence that the theorems’ conclusion—and therefore one of their assumptions—fails to hold for the participants. But, as explained above, majority preference cycles and the nonexistence of a Condorcet winner do not logically imply voting cycles. Not only do they not logically entail voting cycles, but there is little reason to think that they cause voting cycles. If there is no reason to think that the absence of a Condorcet winner causes voting cycles, then failing to observe voting cycles in experimental or real-world settings does not “falsify” the conclusions of the theorems on the multidimensional spatial model.

§3.4 Does politics have just one dimension?

Many policy problems would appear to fit the multidimensional spatial model. Anytime policy, or policy bundles, differ in more than one quantifiable way that people care about, they can be represented as points in a multidimensional space. It is hard to make sense of familiar phenomena such as legislative logrolling without assuming that there are multiple di-

mensions to the policy space, enabling legislators to make concessions on one dimension for the sake of gains on others.¹⁰ But the results of empirical research like Poole and Rosenthal's and other efforts to measure legislator preferences empirically are often taken as evidence that, notwithstanding the *apparent* multiplicity of dimensions to many political problems, political conflict unfolds mainly along a single Left-Right ideological dimension (Mackie, 2003, pp. 90, 91). If policy preferences can be adequately represented with a single Left-Right dimension, then we can apply the median voter theorem and conclude that Condorcet winners always exist. To assess this claim let us review the methodology behind the empirical research it appeals to.

Suppose every bill that comes before the House of Representatives can be placed on a Left-Right continuum and each legislator has single-peaked preferences with respect to this continuum. If this assumption is right, and if legislators vote sincerely according to their preferences, then hypotheses about the locations of the legislators' ideal points along this ideological spectrum have observable implications. Extremely conservative and extremely liberal members will almost never vote the same way on a bill, but moderate conservatives will often vote with liberals, and moderate liberals will often vote with conservatives. Seeing who votes with whom provides evidence of where members of Congress are located along this Left-Right dimension. Given adequate data, we can try out different hypotheses about the ideo-

¹⁰For an argument that logrolling is evidence of majority preference cycles, see McGann (2006, p. 70–76) and the sources cited there.

logical locations of members of Congress and identify the hypothesis under which the observed voting record was most likely. This is the idea behind Keith Poole and Howard Rosenthal's NOMINATE scores of ideology (Poole and Rosenthal, 2007).

If Congress members' preferences were not completely described by a single point along a Left-Right dimension—*e.g.*, if one needed two or more coordinates to describe their preferences because political conflict unfolds along two or more dimensions—then the estimates of these points would not do a very good job of explaining their voting behavior. We would observe errors in the form of postulated left-wingers joining forces with postulated right-wingers in opposition to the postulated moderates.

American history provides examples of periods in which considerable variation in Congressional voting behavior is hard to explain with the hypothesis of a single ideological dimension. Ogden Reid, a Republican from New York, and John W. Davis, a Democrat from Georgia, were both delegates to the House of Representatives in the 89th Congress, 1965-67. Poole and Rosenthal's NOMINATE estimation method generates nearly identical estimates of their ideological positions under the assumption of a single ideological dimension. Yet they voted differently on several important pieces of legislation, such as the Voting Rights Act of 1965, which Davis opposed and Reid supported.

What explains this variation, if not their positions on a single Left-Right ideological spectrum? Given this unexplained variation, we might posit a second ideological dimension to account for why Southern Democrats like Davis voted differently from Northern Republicans like Reid. If we assume

that bills occupy locations in a *two*-dimensional space, and that members of Congress have ideal points in this space, want bills to be as close to their ideal points as possible, and vote sincerely according to these preferences, then we can use the voting record to carry out the same estimation strategy and estimate these ideal points. Figure 3.4 from above illustrates Poole and Rosenthal's results for the 89th House. On this second dimension, which picks up civil rights issues, Davis and Reid occupy, respectively, the far right and the far left ends of the spectrum, despite their similar positioning along the first ideological dimension.

It is important to recognize that the NOMINATE method of ideal point estimation and its offshoots are not methods of estimating the dimensionality of the issue space. Nor do they test whether legislators have Euclidean preferences over this issue space. The approach starts from the assumption that the spatial model is correct. It assumes, in particular, an ideological space of a particular dimension. On the basis of those assumptions, it then estimates legislators' ideal points in the assumed space. These *a priori* assumptions are the lever which enables reasonable conjectures about legislators' preferences over the set of possible legislation, a set which is infinitely large.

But even though this approach does not provide formal, statistical tests of hypotheses about the dimensionality of the issue space, it is sometimes suggested that its results count as indirect evidence. Outside of periods like the 1960s, when issues involving race divided one of the major parties, the hypothesis of a single dimension can explain most of the Congressional voting record. Positing a second dimension picks up "little" of the remaining

error.

But how much extra variation should we expect a second dimension to pick up if there is in fact a second dimension to ideological conflict? Why assume a 5 percent increase in explained variation counts as evidence *against* rather than *for* the hypothesis of a second dimension? What we would like to know is whether the observed data are more likely given the hypothesis of a single dimension or given the hypothesis of two (or more) dimensions. *The most likely explanation for the data may be the hypothesis of two dimensions, even if this hypothesis explains “only” an additional 5 percent of the observed variation.* There is no way to settle the question without a statistical model that predicts probabilities of various voting records as a function of the dimensionality of the issue space, which the methods for estimating ideal points do not provide. In the absence of a formal statistical test, one must rely on informal evidence for hypotheses about the dimensionality of the issue space.

Some of the evidence that researchers have produced gives us reason to doubt that all pieces of legislation can be placed on a single dimension relative to which legislators have single-peaked preferences. Agenda control allows the majority party to prevent votes on bills that would split the party, which, in turn, could cause the voting record to appear one-dimensional even if the space of possible legislation has in fact multiple dimensions that legislators care about. As one might therefore expect, the hypothesis of a single dimension does indeed perform best during periods when the majority party in the House has had stronger control over the legislative agenda (Dougherty, Lynch and Madonna, 2014).

In the case of the U.S. House, scholars have shown that whether the hypothesis of multiple dimensions accounts for “significantly” more variation than the hypothesis of a single dimension depends on whether the analysis is conducted at the level of the entire voting record or whether separate analyses are conducted for particular issue areas, such as those corresponding to the jurisdictions of thirteen different subcommittees of the appropriations committee. Multiple dimensions appear important to explaining voting within particular issue areas (Crespin and Rohde, 2010).

Such studies may be taken as indirect evidence for the view that the underlying space of feasible policies has multiple dimensions to it, even if the observed Congressional voting record is largely one-dimensional, as Poole and Rosenthal find. If so, then, in general, there is no Condorcet winner relative to legislators’ preferences over the space of feasible policies.

What matters for the argument about popular control over policy is whether there is typically a Condorcet winner relative to *citizens’* preferences, not legislators’ preferences. But if there were strong evidence that all feasible policies could be placed on a single Left-Right ideological spectrum, relative to which elected legislators’ preference were single-peaked, then one might expect the same to be true of citizens’ preferences, and this would knock out the critical, third premise in the critique of popular control presented at the beginning of this chapter. But the evidence for one-dimensional elected legislatures is mixed, at best.

§3.5 Deliberation and single-peaked preferences

Political theorists often point to deliberation as a potential solution to the problems of preference aggregation (Dryzek, 2000; Dryzek and List, 2003; Farrar et al., 2010; Knight and Johnson, 1994, 2011; List, 2002; Miller, 1992). For example, Knight and Johnson (2011) suggest that political argument might allow participants “to sort out, and hopefully reduce, the dimensions over which they disagree” (p. 149). More specifically, the thought is that deliberation might produce agreement that there is one relevant dimension of disagreement, relative to which participants’ preferences are single-peaked. If so, then a Condorcet winner will exist after deliberation. If deliberation has this effect, one would have grounds for rejecting the third premise of the impossibility objection.

It is worth emphasizing here that deliberation is relevant to the impossibility objection only if it changes preferences so that a Condorcet winner exists. That objection, recall, did not rest on any assumption about how people behave. Its critical assumption was just that Condorcet winners frequently fail to exist, so that whatever decision is taken, a majority will prefer a different decision to it. Thus, even if deliberation causes people to agree on the relevant dimensions of a problem and to address each dimension separately, it does not dislodge that critical assumption behind the impossibility objection unless it changes preferences over the set of feasible options—unless it causes a Condorcet winner to emerge. Relatedly, even if it generates an agreement about what the relevant dimensions are, it does not necessarily dislodge the critical assumption. If it generates an agree-

ment that there are *two* relevant dimensions, then we would not expect a Condorcet winner to exist, given the theorems on the multidimensional spatial model. If deliberation is to have effects that are relevant to the impossibility objection I am considering—as opposed to other concerns about stability and the like—the claim must be that it has a tendency to produce Condorcet winners, for example by generating agreement that there is exactly *one* dimension and causing preferences to be single-peaked with respect to it.

Some experimental evidence does suggest that deliberation can have this effect (Farrar et al., 2010; List et al., 2012). In several of the choices that groups confronted in these experiments, there was a “natural” ordering to the alternatives (List et al., 2012, p. 86). However, given the theory reviewed in this chapter, the ideal experiment would test the effects of deliberation in a setting where the policy options differ along multiple dimensions, dimensions that we might expect people to care about. Those are the situations where theory tells us that, under seemingly minimal assumptions about the shape of preferences, majority preference cycles are nearly guaranteed. In the experiments, there were a few choices that might be construed as multidimensional in nature (List et al., 2012, p. 86). But there was nothing as clearly multidimensional as the choice of how to allocate funds across multiple programs. Empirically, it remains an open question whether, in those settings, deliberation has any tendency to produce single-peaked preferences.

Theoretically, is there any reason to expect deliberation to induce single-peaked preferences in those settings? When it comes to a budget, *reasonable* people will not agree that just one dimension is relevant. For example, rea-

sonable people will not agree that all that matters in the drafting of the U.S. federal budget is the amount of money allocated to defense. When the possible objects of choice differ along multiple dimensions that any *reasonable* person would care about, no form of deliberation that tends to produce reasonable beliefs would induce single-peaked preferences, as McGann (2006, p. 66) and Achen and Bartels (2016, p. 29) point out. So, theoretically, there is an open question about the mechanism by which deliberation could induce single-peaked preferences in such settings.

Finally, it is worth pointing out that if deliberation had this effect, it would not be a panacea for all of the conceptual and practical problems that social choice theory has identified. Contrary to a widely held view, individuals' preferences being single-peaked preference does not eliminate the possibility for strategic misrepresentation of preferences (or “manipulation”) described by the Gibbard-Satterthwaite theorem. Individuals would still sometimes be able to achieve better outcomes, by their lights, if they acted as though their preferences were not single-peaked with respect to the same dimension as everyone else's (Penn, Patty and Gailmard, 2011).¹¹ So there would still be an open and non-trivial question about whether and how institutions could be designed so as to reliably *implement* the decision that respects every majority's preferences, even if deliberation reliably ensured that such a decision *exists*.

¹¹Penn, Patty and Gailmard (2011) show that the single-peaked domain restriction affords an “escape” from the Gibbard-Satterthwaite theorem only if the domain of “ballots” as well as true preferences is restricted.

§3.6 Conclusion

Disproving the premise that citizens' preferences often have the structure of Condorcet's paradox is an inherently hard problem when the number of feasible alternatives is large. When a choice must be made between a large number of alternatives—*e.g.*, the infinitely many rates at which to tax something, the thousands of levels at which the minimum wage could feasibly be set, etc.—collecting data on each individual's preferences over the alternatives is practically impossible. The only way forward is to make *a priori* assumptions about individuals' preferences. An assumption like transitivity reduces the empirical data one needs to acquire in order to determine whether a Condorcet winner exists; the assumption of single-peakedness eliminates the need for data altogether. Efforts at showing that Condorcet winners generally exist are doomed from the start unless one is willing to make these theoretically motivated assumptions. But in contexts where the alternatives differ along multiple relevant dimensions, such as the choice of a budget, the most theoretically defensible assumptions imply that Condorcet's paradox occurs with certainty.

Perhaps these assumptions are false. Perhaps Condorcet winners do generally exist and majority preference cycles are rare. But in the absence of empirical evidence, there is no reason for this to be the default expectation. Given the difficulty of disproving the premise about Condorcet's paradox, we should look into an alternative strategy for defending the idea of popular control against the impossibility objection. The only alternative is to re-examine the premise that popular control implies respecting the preferences

of majorities. The first step in executing this strategy is a more careful examination of the concept of control, which I undertake in the next chapter.

Chapter 4

The concept of control

The provisional account of popular control from chapter 2 is vulnerable to an objection. It makes the possibility of popular control depend on whether a decision can be taken that respects the preferences of every majority, but there may be no decision satisfying this condition. Indeed, there are no *general* theoretical reasons to expect any decisions to satisfy the condition. We should therefore re-examine the assumption that the condition is necessary for popular control. In this chapter, I assemble the materials for an account of popular control that allows us to reject that assumption. The materials are a definition of control and arguments about how the definition applies to controversial cases. In the next chapter, I will argue that it is possible for democracies to give all majorities simultaneous control over certain decisions, in the sense of this definition.

Let me reiterate what we are after—what would, and what would not, count as an adequate response to the impossibility objection to popular control. The goal is *not* to find a voting rule that performs well despite the presence of cycles. The impossibility objection is not an objection to the use of democratic voting rules. Its force derives, not from claims about voting

rules, but rather from a claim about the *concept* of popular control, namely that a necessary condition for a decision to be under popular control is that it respects majorities' preferences. The objection is that this condition is in general impossible to satisfy: no matter the voting rules and other institutions in place, a majority will prefer an alternative decision to whatever decision is taken. We have already seen why the factual claim about majorities' preferences is hard to refute. So, faced with this challenge, the question to ask is whether there are compelling reasons to reject the conceptual claim. Are there reasons to revise the assumed criteria for judging that something is under popular control? Suitably revised, can the criteria can be met in general, even when majorities' preferences form cycles and no Condorcet winner exists? Only after identifying these reasons and revising the criteria accordingly can we return to questions about voting rules, elections, and other institutions.

How should we assess rival accounts of the criteria for judging that a decision is under popular control? A concept of popular control should bear some connection to the concept of control as it is ordinarily used (List and Valentini, 2016). We should therefore start with a general analysis of the concept of control—general, in that it appears plausible even outside of the context of democratic theory and questions about popular control. To this end, we should consider non-political examples of control and construct a definition that matches our judgments in these examples. The goal is not to perfectly characterize what people ordinarily mean when they use the concept of control. But a conception of popular control should respect ordinary intuitions about the concept to some extent, otherwise it will be

misleading to advertise it as a conception of *popular control*, as opposed to something else (cf. Lovett, 2010, pp. 4, 5).¹

After using some non-political examples to arrive at a general account, I discuss its implications for several controversies relevant to a theory of democratic control, such as whether exercises of control must be intentional, and whether an agent's control must be robust to the dispositions of other agents.

§4.1 A preliminary statement

Suppose the temperature in *A*'s office is regulated by a thermostat. There is a range of temperatures such that, for any temperature in the range, if she programmed the thermostat to aim at that temperature, that would soon be the temperature in the office. Knowing nothing else about the situation, it will be plausible to judge that *A* has control over the temperature in the office.

The temperature can be understood as a *variable*, which has different possible *values* or *outcomes*. To capture the intuitive judgment about the case, we want a general definition of what it means for an agent to have control over a variable. Here is one possibility. We might say that an agent has control over a variable if there is a range of possible values such that, for each value in the range, she has actions that would cause the variable to take

¹Lovett (2010, pp. 15–21) also has a useful description of what he calls a “constructivist” approach to conceptual analysis, of which the approach I take here is also an example.

on that value. Control is defined in terms of the agent's *menu of actions* and their causal relationships to the variable. Within the large literature on the concept of power, originating with Dahl (1957), there is a family of positions that define an agent's power in terms of a menu condition like this one.²

There are two shortcomings to identifying control with the satisfaction of a menu condition. The first is that such a definition does not extend well to groups. The natural extension would be to say that a group has control over a variable if, for each value in a range of possible values, its members have actions that would, taken together, cause the variable to assume that value. But this definition leads to implausible judgments about groups whose members are unable to surmount collective action problems. It implies, for example, that the firms in a competitive market, taken as a group, have control over the market price: for any possible price p , it is possible for the firms to each choose quantities of output that, through the market-clearing mechanism, cause the market price to be p . Profit-maximizing firms would not choose to produce these quantities of output in a competitive equilibrium, but, according to the menu-based definition of group control, all that matters is that the firms' menus of actions include the means to determine the price, not whether it would ever be rational for the firms to choose these actions. That feature of the definition seems like a good reason to reject it. We do not want to say that the firms in a competitive market jointly control the market price.

²For example, Barry (2002); Dowding (2003); Lovett (2010). See also Sher (N.d.), who characterizes a measure of the freedom that an agent's menu offers in terms of the 'control' that the menu offers the agent.

A second limitation of the menu-based definition appears even in the case of a single agent. Suppose that A 's boss has instructed her to keep the temperature at a particular level, and she would be disciplined, perhaps fired, if she ignored his instructions and programmed the thermostat to achieve a different temperature. Presumably we would then wish to say that she has less control over the temperature than she would in the absence of the boss's intervention. Yet her menu of actions still includes the same means of determining the temperature. What has changed is merely that these actions, in addition to affecting the temperature, now affect other variables that she cares about, such as retaining her job. If control over a variable requires only that an agent's menu of actions includes the means of determining the variable's value, then we cannot distinguish between A 's control in the two versions of the example.³

To address both limitations, our definition should make judgments about agents' control depend on how their preferences affect the variable. When groups suffer from collective action problems, the results of the members' actions are outcomes that they all consider worse than some alternative. When

³Some readers may object that the agent's menus of actions are different in the two situations, and a menu-based definition of control could therefore support different judgments about the agent's control in the two situations. In reply, I conjecture that if one can explain the different judgments about control in the two situations using only properties of the agent's menus of actions, the characterization of the menus' relevant properties will be a round-about way of describing what is below called the preference condition.

the office worker's boss would discipline her for changing the temperature, the actual temperature may fail to reflect her preferences in the manner we associate with maximum control. A definition that accommodates these intuitive judgments might run: agents have control over a variable if, whenever they want the variable to take on a particular outcome—in the sense of preferring that outcome to all the other possible outcomes—then it does.

This condition is, however, too strong, in addition to being ill-suited for the analysis of a group's control. It is too strong because an agent may plausibly be described as having some control over a variable even if the condition is violated. Suppose that the office worker's boss has not given her any instructions regarding the temperature, but she believes there is a small chance he would punish her if she adjusted the thermostat so as to increase the temperature. She does not care much about the temperature relative to avoiding punishment, so she does not adjust the temperature to the level she most prefers. Should we then conclude that she has *no* control over the temperature? That conclusion seems wrong, but the proposed definition entails it.

Between the scenario in which there is no chance of being disciplined by her boss, and the scenario in which she knows for certain that she will be disciplined if she adjusts the temperature, there is a continuum of scenarios in which it is more or less costly, in expectation, for *A* to adjust the temperature according to her preferences. At each point along this continuum, it will be true that if *A* cared enough about the temperature relative to avoiding punishment, she would adjust the temperature to the level she prefers. But how much she has to care about the temperature—the “weight” she has

to attach to it, relative to the weight she attaches to avoiding punishment—in order for it to be rational to adjust the temperature according to her preferences, is more or less, depending on where she finds herself on the continuum. At one end of this spectrum, where she is certain to be fired for adjusting the temperature, she would have to attach a lot of importance to the temperature in order for it to be rational to set it to her preferred level; at the other end, where there is zero chance of reprisal, it is rational for her to adjust it to her preferred level no matter how little weight she attaches to the temperature relative to keeping her job. One might say she has more or less control over the temperature, or say her control is more or less significant, depending on how “intense” her preferences over the temperature have to be—relative to her desire to avoid being disciplined—in order for it to be worthwhile to adjust the temperature to the level she prefers. We might therefore say that what is necessary for an agent to have *some* control over a variable is not that the variable take on the value the agent most prefers, but merely that if the agent cared enough about the variable, then she would choose actions resulting in the variable taking on the value she prefers. And the agent can be described as having more or less control, depending on how much weight the agent has to attach to the variable, before she finds it worthwhile to act so as to make it take on the value she prefers.

This condition needs further revising if we want the definition of control to extend neatly to groups, however. It normally makes no sense at all to refer to the value of a variable preferred by a group’s members, because not all members of the group will agree on which possible value of the variable is best. They will share preferences over some pairs of possible values,

but not all pairs. What we might reasonably expect of a group that has control over a variable is that if they share any pairwise preferences, and they attach enough importance to the variable, then their shared preferences will *constrain* the variable. That is, if each member of the group preferred one possible value of the variable to another, then the variable would not assume the less preferred value, provided the group's members all cared enough about the variable. For example, two workers who share an office may not agree on the optimal temperature, but they may both want the temperature to be warmer than 65 degrees and cooler than 75 degrees Fahrenheit. These are the “Pareto optimal” temperatures—the values of the variable such that there is no other value that the workers both prefer.⁴ If the temperature is under their joint control, then one would expect their shared preferences to constrain the temperature, causing it to lie in the interval between 65 and 75 degrees, provided they care enough about the temperature relative to other things they care about (such as retaining their jobs) that could potentially come into conflict with achieving a Pareto optimal temperature.

So far, we have identified one dimension along which degrees of control vary—the threshold of importance agents must attach to a variable in order for their preferences to have a constraining effect. A second one is the range of possible values over which the agents' preferences have this constraining effect. A worker whose preferences prevent the less preferred of two temper-

⁴An alternative x is Pareto optimal (relative to the preferences of some group of individuals) unless there is another alternative y such that at least one person strictly prefers y to x and no one strictly prefers x to y .

atures from occurring, for any two numbers in the range 0-100 degrees, has more control than one whose preferences have such effects only for numbers in the range 60-80 degrees, all else being equal.⁵

The preceding observations suggest the following definition of control.

Definition of control. A group has *some control* over a variable if, for some pairs of values that the variable might take,⁶

- (*) if all members of the group preferred one value to the other, then the variable would not assume the less preferred value, provided the members of the group cared enough about the variable (relative to other variables they also care about).

⁵There are still other dimensions along which control may be thought to vary, intuitively speaking. For example, the fact that an agent prefers one outcome to another may merely reduce the probability of the latter, without driving its probability all the way to zero, and we might think the degree of her control depends on the magnitude of this effect. I ignore these other possible dimensions of control, for simplicity's sake.

⁶By values that it "might" take, I mean possible realizations of the variable that have some nonzero *ex ante* probability of occurring: they might occur, but for the fact that all members of the group prefer one realization to the other and hold their preference with sufficient intensity. I include this qualification so that the definition does not imply, absurdly, that a group has control over a variable merely in virtue of there being a pair of possible values that the variable surely will not take irrespective of the group members' preferences. Without the qualification, condition (*) would be trivially satisfiable.

All else being equal, the group has *more control* over the variable, the less weight its members have to attach to the variable in order for condition (*) to hold, and the greater the range of pairs of possible values for which condition (*) holds.

I refer to (*) as the *preference condition*.

The preference condition refers to counterfactuals. What would happen, counterfactually, if both office workers preferred a temperature of 70 degrees to 75 degrees? If the workers have control over the temperature, then, whatever happens, the temperature would not be 75 degrees, provided they care enough about the temperature relative to other things they care about. These counterfactual judgments can also be expressed with causal language, as when I refer to the *constraining effects* of agents' preferences on a variable under their control. The workers' shared preferences constrain the temperature to be a value in between 65 and 75 degrees, provided they both attach enough importance to the temperature.

Let me right away acknowledge and forestall a possible worry. Given how the preference condition is stated, it may seem like the definition makes it too easy for a group to have control. That their preferences would, counterfactually, make a difference to an outcome if only they cared enough, may seem like a rather trivial fact. Anticipating the application to popular control, the reader may worry that the definition will remove the concept of its bite. In response, it is important to recognize that control comes in degrees according to this definition. For a *minimal* degree of control, it suffices if a group's preferences would have some constraining effect on a variable if they cared enough about the variable. That is indeed a trivial accomplishment

in most contexts—but on any good definition, having a *minimal* degree of control should be a trivial accomplishment. Since control comes in degrees, democratic popular control will also come in degrees. I return to this point in chapters 5 and 7.

The definition of a group's control covers an individual agent's control as a special case. An individual agent is a group containing one member. The definition supports the judgment, from the original example involving a single individual, that the worker has some control over the temperature in her office. The preference condition is satisfied, because for any two possible temperatures, she will not program the thermostat to target the less preferred of the two if she cares enough about the temperature relative to other things she cares about, such as avoiding punishment from her boss. The definition also supports the judgment that she has more or less control, depending on how likely it is that adjusting the temperature will meet with reprisal from her boss. The more likely such punishment, the greater the weight she must attach to the temperature—relative to the weight she attaches to avoiding the punishment—before her preferences over the temperature have a constraining effect on it. Thus, the worker has less control over the temperature, the more likely she is to be disciplined for adjusting the temperature.

As another illustration of how the definition applies to groups containing more than one member, consider the following example. Suppose that two students, *A* and *B*, have a project they must complete as a team. Each student can choose either to contribute to its completion or to slack off. The project will be successful if and only if they both choose to contribute. Each

student finds contributing unpleasant in itself, but would find it worthwhile so long as it would result in a successful project. Neither student is sure how much importance the other student attaches to the project's completion and cannot rule out the possibility that the other student will fail to do his part and thereby make her own contribution a wasted effort.⁷

I expect the reader to share the intuition that the students have control over how their project turns out, given the stipulated facts. The outcome is in their hands. If they fail to complete the project and complain to the professor that they had no control over how it turned out, the professor would rightly scoff. The professor could reply that if they had both cared enough about the success of the project, then they would have each done their part and ensured its success.⁸ In making this reply, he would in effect

⁷One might express these assumptions by saying that student i incurs a cost $c > 0$ if she contributes, but attaches a value v_i to the project's successful completion, so that her payoff is $v_i - c$ if she contributes and the project is completed, 0 if she does not contribute and the project is not completed, and $-c$ if she does contribute and the project is not completed. Each student knows her own valuation, but of the other student's valuation knows only that it lies in some interval $[-\bar{v}, \bar{v}]$, where $\bar{v} > c$. (The lower bound to this interval is not a crucial element of the story, but its being less than zero allows for the possibility that a student prefers, all else being equal, for the project to fail.) She attaches (whatever her own valuation) a probability $F(v)$ to the other player's valuation being no greater than v , where F is a continuous, increasing function.

⁸What the professor might be thinking, if he is inclined to think game-

be invoking the given definition of control: his reply is just an assertion that theoretically, is that the students find themselves in a Bayesian Nash equilibrium to the game described in the previous footnote, an equilibrium in which each player i contributes if and only if $v_i \geq v^*$, where v^* satisfies $v^* - v^*F(v^*) = c$. From the failure of the project he infers that one or the other student did not value the project's successful completion enough.

The professor's incredulous reply to the students could not be expressed as the thought that the students find themselves in a Nash equilibrium in mixed strategies to a game of complete information. In the equilibrium to such a game, where the players know each other's valuation, the probability that both players contribute *declines* as their valuations increase. What fixes the probabilities of a player's actions in such an equilibrium is not her own valuation of the project but the assumption that the *other* player is indifferent among his strategies. For the other student to be indifferent between his strategies, it has to be true that when he cares *a lot* about completing the project relative to the cost of contributing, he believes it is very unlikely that his partner will contribute—otherwise he will strictly prefer contributing himself. Thus, in such an equilibrium, the probability of the students' both contributing declines as they attach greater and greater value to the project's successful completion. How and why such an equilibrium would arise in the real world is not at all obvious. But suppose, for the sake of argument, that such a game and its mixed strategy equilibria describe what would happen in counterfactual scenarios differing only in the relative weight the students attach to the project's success. Then the students would not have control over the project's success, according to the definition: the

the preference condition is satisfied. The satisfaction of that condition is, I submit, the intuitive reason we would judge that the students have control.

The preference condition implies a connection between a group's collective control and Pareto optimality. If the students have control, and both students prefer exerting effort and completing the project successfully to the outcome in which neither exerts effort and the project fails, then the latter, strictly suboptimal outcome will be avoided, provided they each attach enough importance to the project's successful completion relative to the cost of effort. Thus, while a group's control is compatible with a Pareto suboptimal outcome—because they may not care enough—it also implies that the particular suboptimal outcome would not occur if they did care enough. This observation anticipates the explanation, in chapter 5, for why the absence of a Condorcet winner does not preclude every majority's having control over a variable—a Condorcet winner is, by definition, just an outcome that is Pareto optimal relative to every majority. I comment further on the connection with Pareto optimality—as well as the sense in which having control in the “negative” sense captured by the definition sometimes, but not always, implies that an individual agent gets the outcome she most prefers—at the end of section 4.2.

In the remainder of this chapter, I develop the account of control more formally and point out some of its important implications for debates in the literature.

more strongly they felt about its success, relative to the cost of effort, the less likely the project would be completed, in equilibrium.

§4.2 A more formal statement

As a point of reference for the discussions to come, consider the following stylized example. A newspaper editor is considering whether to publish a story that would put the owner of the newspaper in a bad light. Before she makes her decision, the owner can issue a threat to fire her if she chooses to publish the piece. The owner can carry out the threat without cost, so it is credible. The editor would like to publish the story, but not at the cost of losing her job. The diagram in Figure 4.1 models their interaction. After each possible sequence of actions, there is a pair of numbers, which represent the two players' preferences over the eight possible sequences of actions (the first number represents the owner's preferences).

Suppose that the owner carries out his threats and never fires the editor without having first threatened to do so; the editor declines to publish the story if and only if the owner threatens to fire her; and the owner threatens her. These “strategies,” as they are called in game theory, form a *Nash equilibrium*: a pair of strategies such that each player's strategy is optimal, given her preferences over the outcomes of the game and correct expectations about the other player's strategy.⁹

⁹That is to say: holding fixed the other player's strategy, neither player can achieve a higher payoff by following a different strategy. The postulated strategies satisfy a more demanding (and compelling) equilibrium concept—subgame perfection—which requires that players' expectations about how other players would behave “off the equilibrium” path are not only correct but also reasonable, in that they do not believe that other players would behave irrationally off the equilibrium path.

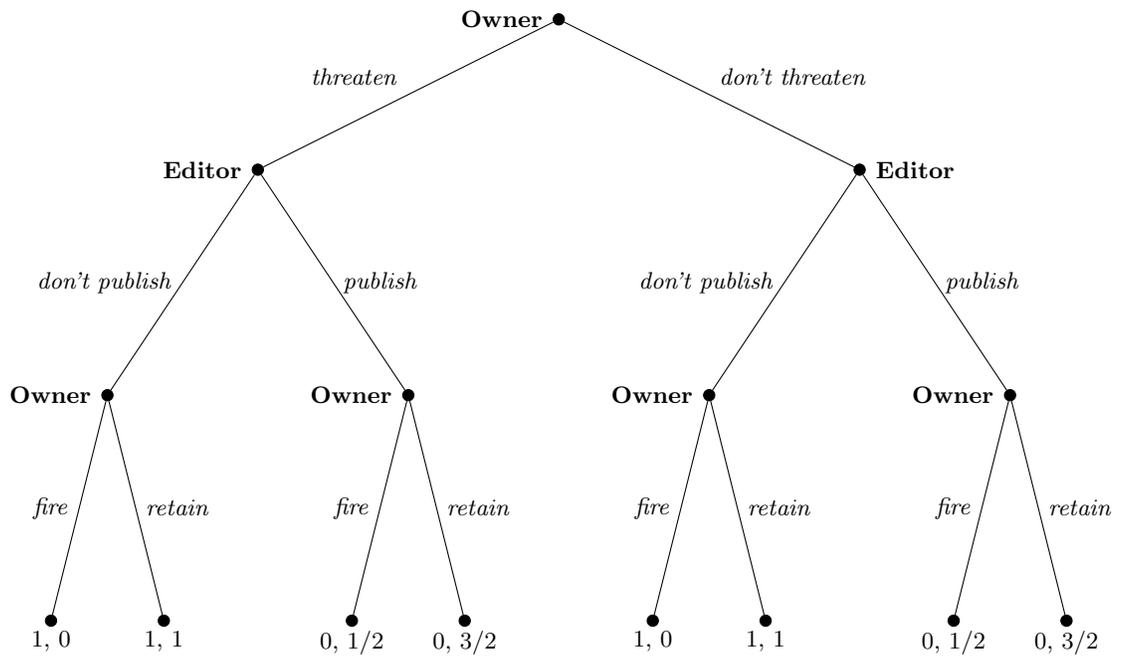


Figure 4.1: A diagram modeling the situation of the newspaper owner and editor.

I expect everyone will agree that the owner has some control over whether the story is published. The definition from the previous section captures this judgment. The example will illustrate how to interpret and apply the definition.

When we judge that the owner has control over whether the story is published, we are not referring to his ability to bring about one particular outcome, such as the second outcome, counting from left to right. We are instead referring to his ability to make different *events* occur, or, equivalently, to make a *variable* take on different values. We can distinguish between **outcomes**, **events**, and **variables** using definitions borrowed from probability theory. Label the eight branches of the game tree o_1, o_2, \dots, o_8 , from left to right. Each branch corresponds to one possible **comprehensively described outcome** of the owner and editor's interaction. The descriptions comprehend all of the outcomes' relevant properties. An **event** is a set of comprehensively described outcomes, namely those outcomes instantiating the event. For example, the event of the story's publication is equivalent to the set $\{o_3, o_4, o_7, o_8\}$. The event of the owner's issuing a threat is equivalent to the set $\{o_1, o_2, o_3, o_4\}$. The concept of a random variable can be used to describe events or, equivalently, the *features* of outcomes that an agent cares about. A **random variable** is a function $X : \{o_1, o_2, \dots, o_8\} \rightarrow \mathbf{R}$ that assigns a number to each of these outcomes. For example, the function X defined by

$$X(o) = \begin{cases} 1 & \text{if } o \in \{o_3, o_4, o_7, o_8\} \\ 0 & \text{otherwise} \end{cases}$$

is a dichotomous variable that indicates whether the story is published. Asking whether the owner has control over the story's publication amounts to asking whether the owner has control over this variable.

Note that asking whether the owner has control over the story's publication is, on its face, different from asking whether the owner has control over which of several comprehensively described outcomes obtains. The definition I have offered takes this difference at face value and treats an agent's control over variables as the concept in need of analysis.

The preference condition concerns the causal relationship between a variable and an agent's preferences over the variable. To clarify the condition, the first step is to explain the sense in which an agent can have *preferences over a variable*, as opposed to preferences over comprehensively described outcomes. In the newspaper example, there are two variables that the editor cares about: whether the story is published and whether she retains her job. The first variable is equivalent to X as just defined; the second is equivalent to the variable Z defined by $Z(o) = 1$ for each $o \in \{o_2, o_4, o_6, o_8\}$ and $Z(o) = 0$ for each $o \in \{o_1, o_3, o_5, o_7\}$. The editor's preferences over the comprehensively described outcomes are represented by the **value function**¹⁰ v :

$$v(o) = 1/2X(o) + Z(o)$$

The editor has **preferences over the variable** X in the following sense:

¹⁰These functions are commonly called utility functions, but for some readers, the common name has potentially misleading connotations I wish to avoid.

for any two outcomes o and o' , if $Z(o) = Z(o')$, the editor's value function assigns a higher value to o than o' if and only if $X(o) = 1$ and $X(o') = 0$. That is, for any two comprehensively described outcomes that are alike with respect to their other relevant properties, she prefers the one in which the story is published.

In this example, the editor's preferences over the story's publication are independent of the realized value of Z . Holding fixed whether she is going to keep her job or not, she prefers for the story to be published. (Note that I am not here saying that she optimally chooses to publish the story in equilibrium irrespective of whether it will result in losing her job. That would be false.) Whether $Z(o) = Z(o') = 0$, or instead $Z(o) = Z(o') = 1$, in either case she strictly prefers o to o' if and only if $X(o) = 1$ and $X(o') = 0$. In that sense, her preferences over X are independent of her preferences over Z . But in other contexts, an agent's preferences over a variable will depend on the values of other variables. A dinner guest prefers to drink red wine with his dinner if the host is serving red meat, but white wine if the host is serving fish. In such cases, judgments about an agent's preferences over a variable are conditioned on assumptions about the values of other relevant variables: the guest prefers red wine, conditional on the assumption that red meat is being served.

The preference condition refers to the importance, or weight, agents attach to variables ("...if the members of the group cared enough about the variable"). I sometimes express judgments about a variable's weight as judgments about the intensity of the agent's preferences over the variable. These judgments concern the *relative* importance of a variable: its importance rel-

ative to the importance of other variables the agents care about. What I have in mind, for example, are judgments like, “The editor cares so strongly about publishing the story that she is willing to sacrifice her job in order to publish it.” We are not referring to an unobservable psychic quantity; these judgments have observable implications.

For example, compare two versions of the newspaper example. In the first, the editor has the stipulated preferences, represented by the value function $v(o) = 1/2X(o) + Z(o)$. In the second version, the editor’s preferences are given by the value function $v'(o) = 2X(o) + Z(o)$. In this second version of the scenario, the editor attaches a greater relative weight to the story’s publication—relative to the weight attached to keeping her job—than in the first version. This judgment has observable implications, given the assumption that the editor’s behavior is a rational response to the owner’s strategy. In the first version of the scenario, the editor will choose not to publish the story if the owner adopts the postulated strategy (issuing the threat and carrying it out if disobeyed). In the second version of the scenario, the editor will choose to publish the story even if she knows the owner will fire her as a result.

Assessing the preference condition involves judgments about the effects of counterfactual variation in agents’ preferences. We have to ask what would happen, counterfactually, if the editor’s preferences over the story’s publication, and the importance she attaches to the variable, were different. To represent this variation, we can model the newspaper owner and editor’s interaction as a *Bayesian game*, as follows. Assume that before either player acts, “nature” first chooses the editor’s “type.” Specifically, the

editor's value function is of the form $v(o) = w_X X(o) + Z(o)$, where w_X can be a positive or negative number, and nature chooses the number before the owner chooses whether to issue a threat. Assume both players observe nature's choice. Figure 4.2 represents the game.

In game theory, type spaces are used to model players' uncertainty about others' preferences, so nature's choice of a player's type is normally not observed by all players in the game. I am using the same formal apparatus of a Bayesian game to represent judgments about counterfactuals, not the effects of uncertainty about players' preferences. We started with a model of the owner and editor's *actual* situation, and it was assumed that both players knew each other's preferences. The appropriate model of the counterfactual in which the editor's preferences are different is therefore a game in which the editor's preferences are different but, again, the owner has correct beliefs about the editor's preferences. If we are not justified in assuming the owner would form correct beliefs about the editor's preferences if they were different, then we should use a game of incomplete information, with uncertainty about the editor's type, to model the *actual* situation, too. In either case—with or without uncertainty about the editor's type—we assess the preference condition by considering the equilibrium to the Bayesian game in which nature chooses the editor's type, and we ask what would happen, in this equilibrium, if the editor's type were different from what it actually is.

A strategy for the owner now describes the action he takes after nature's choice of the editor's type as well as the actions he takes after each of the editor's possible actions. A strategy for the editor describes the actions she takes after nature's choice of her type and the owner's choice of whether to

issue a threat. Assume that the owner's strategy is to issue a threat and fire the editor if and only if she publishes the story, whatever the editor's type. Assume that the editor's strategy is to publish the story if both $w_X \geq 0$ (she prefers for the story to be published) and either the owner has not issued a threat or $|w_X| > 1$ (she cares more about the story's publication than retaining her job), and not to publish it otherwise. These strategies form a Bayesian Nash equilibrium. (They satisfy more stringent requirements, too. But what is most important here is the idea that in judging the editor's control, one considers a Bayesian game in which her preferences can vary, and one holds fixed the players' equilibrium strategies.)

We can now illustrate how the preference condition is to be interpreted and applied. Whether the editor has control depends on the effects of variation in her type, given the assumption about the players' equilibrium behavior. The preference condition requires that, provided the editor cared enough about the story's publication, relative to the other variable she cares about, the story would be published if and only if she preferred for it to be published. The condition is satisfied in the postulated equilibrium: provided $|w_X| > 1$ (i.e., the editor attaches more weight to the story's publication than retaining her job), the story is published if and only if $w_X \geq 0$ (i.e., the editor prefers for it to be published). Equivalently: the outcome of equilibrium play after nature chooses $w_X > 1$ is an outcome o with $X(o) = 1$, and the outcome of equilibrium play after nature chooses $w_X < -1$ is an outcome o with $X(o) = 0$. (Note that to show that the preference condition is satisfied, we do not need to say what happens when $w_X \in (-1, 1)$. We need only say what happens when the editor attaches "sufficient" weight to

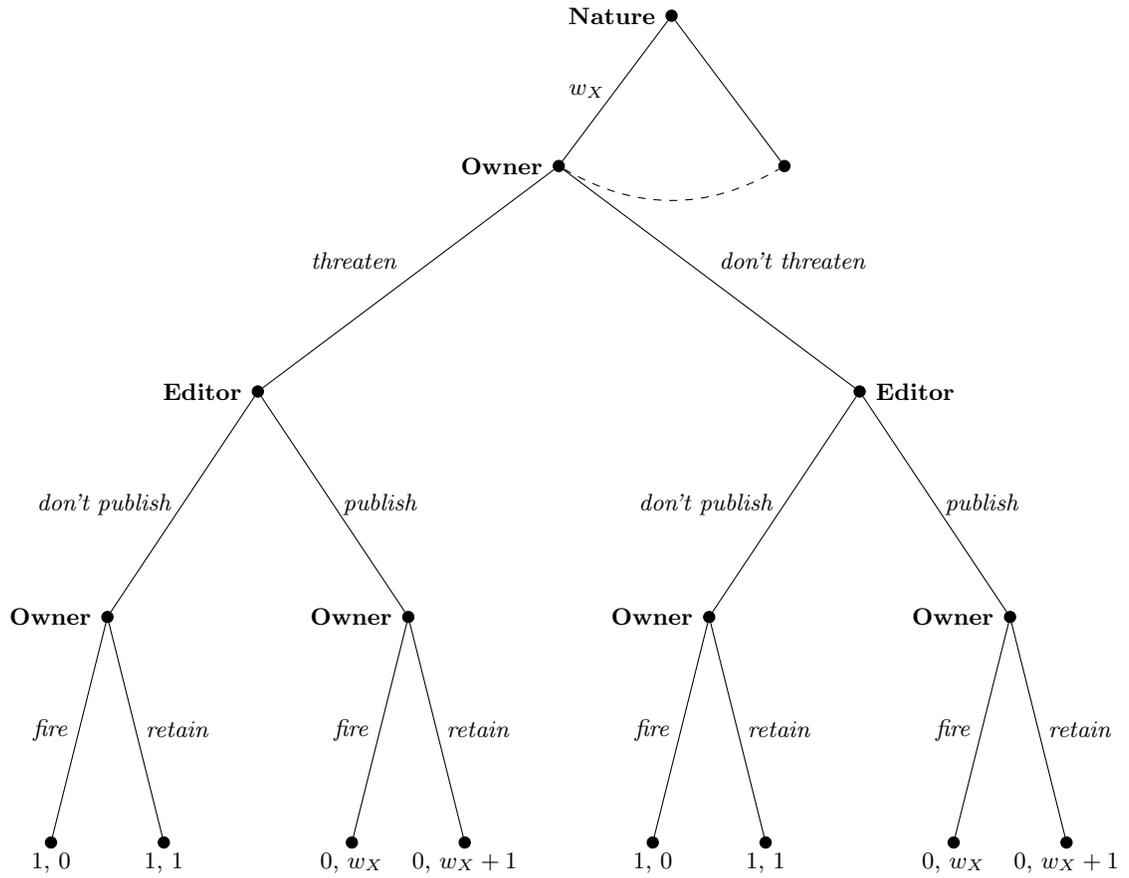


Figure 4.2: A diagram modeling the situation of the newspaper owner and editor. The determination of the editor's preferences is represented as nature's choice of w_X , prior to the owner's choice of action. The dotted arc indicates that there is a continuum of possible values of w_X , and the rest of the game tree depicts what can happen after nature's choice of any one of these values.

the variable.) The editor's preferences over the story's publication determine the value of this variable, provided she attaches enough importance to it.

It is important to keep in mind that the game tree depicted in Figure 4.2, with the fiction of "nature" as a player who chooses the editor's preferences, is only a device to clarify judgments about counterfactual scenarios. By this remark, I mean to emphasize not merely the fictitious status of "nature" as a player in the game, but also the fact that in the actual situation being imagined, the editor's preferences are fixed. The judgment that the editor has control is a judgment about the actual situation, as originally described, in which the editor's actual preferences are given by the value function $v(o) = 1/2X(o) + Z(o)$. That judgment is true, however, in virtue of what would happen in counterfactual scenarios in which the editor's preferences were different from her actual preferences. The more abstract game tree in Figure 4.2 is meant to represent these counterfactual scenarios and what would happen in them.¹¹

¹¹There is a subtle point about the preference condition that is worth recording here for future reference. In asking whether the story would be published if the editor not only preferred its publication but were also willing to sacrifice her job to have it published, I have restricted attention to counterfactual scenarios that differ only with respect to the parameter w_X , but all resemble the actual scenario with respect to the *form* of the editor's value function. Put differently, in assessing the preference condition, I considered only a *restricted domain* of possible preferences for the editor. While I make such domain restrictions in all of the examples throughout

In later applications, I will leave the modified game, with nature's choice of players' types, implicit, and refer to the equilibrium outcomes that arise after nature's choice of counterfactual types as the equilibrium outcomes to "counterfactual games," counterfactual games in which certain players' preferences, or their intensity, are different from what they actually are. That terminology better evokes the counterfactual judgments that the game-theoretic device is meant to represent.

Consider now the owner's situation. The preference condition holds for the owner in the original version of the example as well, given further assumptions. Assume that if, counterfactually, the owner preferred for the story to be published, then he would not issue the threat and would not fire the editor if she published the story, and that in the (actual) scenario with the preferences stipulated in Figure 4.1, he behaves according to the strategy originally postulated. (If we wished, we could model this assumption with the help of another extended game, in which nature chooses the owner's type, but I will proceed informally, having already explained how the judgments about counterfactuals can be formalized with this device, if desired.)

the book, the restrictions never eliminate the particular, problematic configurations of preferences that create problems for traditional accounts of popular control. The restrictions are made so as to pin down the meaning of, and render analytically tractable, the counterfactuals that figure in judgments about control. Unlike the domain restrictions that scholars defend in other contexts, they are not used to circumvent the puzzles about preference aggregation that social choice theory describes.

Assume also that, whatever the owner's preferences, the editor's strategy is to publish the story if and only if there is no threat issued. Then the preference condition holds for the owner: irrespective of how much weight he attaches to the story's publication, the story will be published if and only if the owner prefers for the story to be published. Thus, the owner has control over the story's publication.

It is essential to the arguments later in the book that to assess the preference condition (and thus control), we ask what happens "in equilibrium." Specifically, we hold fixed everything but an agent's preferences (and whatever they influence), and we ask what would happen if the agent's preferences over a variable, and the relative importance the agent attaches to it, were different. We model this counterfactual judgment as a judgment about the equilibrium outcome of the game when the agent's type is different from what it actually is. We hold fixed the preferences and strategies of the other agents whose control is not the object of the judgment. For example, when we assess whether the owner has control, the counterfactual judgments contained in the preference condition are assessed while holding fixed the assumption that the editor behaves according to the postulated equilibrium strategy. Under that background assumption, the counterfactual judgments contained in the preference condition are true. If the owner preferred for the story not to be published, then he would issue a threat and the editor would respond by not publishing the story; if the owner preferred for the story to be published, he would not issue a threat and the editor would respond by publishing the story.

Some readers might think that the preference condition is plausibly con-

sidered a necessary condition for control only under a different interpretation. One might take the view that an agent has control over a variable only if the effects of variation in her preferences are “robust” to other agents’ preferences. But that more stringent condition is hardly ever satisfied. If it is necessary for control, then no captain has control over his ship. Whether the captain is able to determine its direction depends on whether his sailors are disposed to follow his instructions or mutiny and throw him overboard. In most contexts of interest, the consequences of an agent’s actions depend on the preferences and behavior of other agents. The ability to determine the value of a variable is almost always socially conditioned. If control over a variable is identified with a socially *unconditioned* ability to determine the variable’s value, then, at least in the social and political contexts that interest us, almost no variables will ever be correctly described as under anyone’s control.¹² I discuss these issues at greater length in section 4.4.

We have already seen the interesting conclusion that follows from this interpretation of control. Both the owner and the editor have control over the story’s publication, at the same time. Both agents have control, at the same time, because the conditions for one agent’s control are assessed “in equilibrium,” holding fixed the other agent’s actual preferences and equilibrium strategy. It is therefore entirely intelligible and consistent to say, of each agent, that if the agent attached enough importance to it, then the variable would conform to his/her preferences (the preference condition). The preceding discussion explains how and why these judgments are simultaneously

¹²Cf. Lovett (2010, pp. 70, 71).

true.

That implication may seem like a defect. People sometimes talk as though one agent's control necessarily comes at the expense of other agents' control. The reason that the definition allows two agents to have control over a variable simultaneously is that it instructs us to hold fixed assumptions about other agents' preferences when we assess the conditions necessary for an agent's control. As explained, there are reasons to follow these instructions. These reasons push us toward the conclusion that two or more agents can each have control over a variable, at the same time. When we turn to the control of groups, the same reasons will enter into the explanation for why multiple majorities can each have control over a variable, at the same time.

Some readers may be uncomfortable with this implication of the definition, even though they recognize the force of the reasons pushing in this direction. It may bother readers because it may seem to erase the inequality of power in a hierarchical relationship like that between the owner and editor.

That would not be a good reason to resist the definition. Even though the definition of control implies that the owner and the editor each have control over the story's publication, it does not imply that the two actors are on a par, as far as their control goes. All else being equal, an agent has more control over a variable, the less weight the agent has to attach to the variable before the agent's preferences over the variable constrain it. The owner's preferences have this constraining effect irrespective of how much weight he attaches to the story's publication; the editor's preferences have a

constraining effect only if she cares more about the story's publication than retaining her job. It is possible for the editor to have more control over the story's publication than she actually has; the same is not true of the owner.

Other cases structurally similar to the owner and editor's situation illustrate the same asymmetry. When a slave's strategy is to obey his owner's commands, backed by a threat of deadly force, the owner will have control over which of two actions the slave performs, at the same time that the slave has control over which of the two actions he performs. But there is a kind of inequality in their control. The owner's preferences may constrain the slave even if the owner cares very little about the slave's actions, in relative terms. By contrast, the slave will act on his own preferences over his actions, rather than the owner's preferences, only if he is prepared to risk life and limb. (I discuss the issue of comparing degrees of control across persons or groups in chapter 7.)

The reader may have the different intuition that the editor has, in a sense, more control over the story's publication because she *chooses* whether to publish it, whereas the owner merely issues a threat. However, it would not be unnatural to describe the owner as *choosing* to prevent the story's publication by issuing the threat, given that he knows the editor will respond by not publishing it. The real difference between their situations lies in the counterfactual robustness of the editor's ability to publish the story. As the game is set up, the editor can choose whether to publish the story no matter what the editor does, whereas the owner only has an action that affects the story's publication conditional on the editor using a certain strategy. Readers who think that difference should affect our judgments about

whether the owner has control are directed to section 4.4, where I argue that the kind of counterfactual robustness enjoyed by the editor should not be a necessary condition for control.

The editor would have more control over the story's publication if there were constraints on the owner's power to fire her. Suppose that in our running example, the newspaper owner would incur significant costs if he threatened to fire the editor for publishing content he dislikes. The editor, we may suppose, has an employment contract that protects her against being fired, or intimidated with threats of being fired on these grounds, and the owner would risk a costly lawsuit if he tried to threaten her. In figure 4.3, the game tree is reproduced but with modified payoffs for the owner that reflect this stipulation. His payoffs are now represented as the sums of two terms: the value, B , of preventing the unflattering story from being published, and the costs, $c > 0$, of threatening to dismiss the editor. The editor's payoffs are also represented abstractly, using the variable w_X , so we can compare the effects of different counterfactual preferences.¹³

Now let us ask what difference these protections make to the owner's

¹³We are implicitly treating B and w_X , the values the owner and editor attach to the story's publication, as random variables whose realizations the owner and editor observe, prior to playing the game depicted in the figure. Thus, the appropriate game-theoretic model is a Bayesian game in which agents' preferences can vary, and their actions can vary as a function of their realized preferences. I omit these details for simplicity, having already indicated, in Figure 4.2, how they could be filled in.

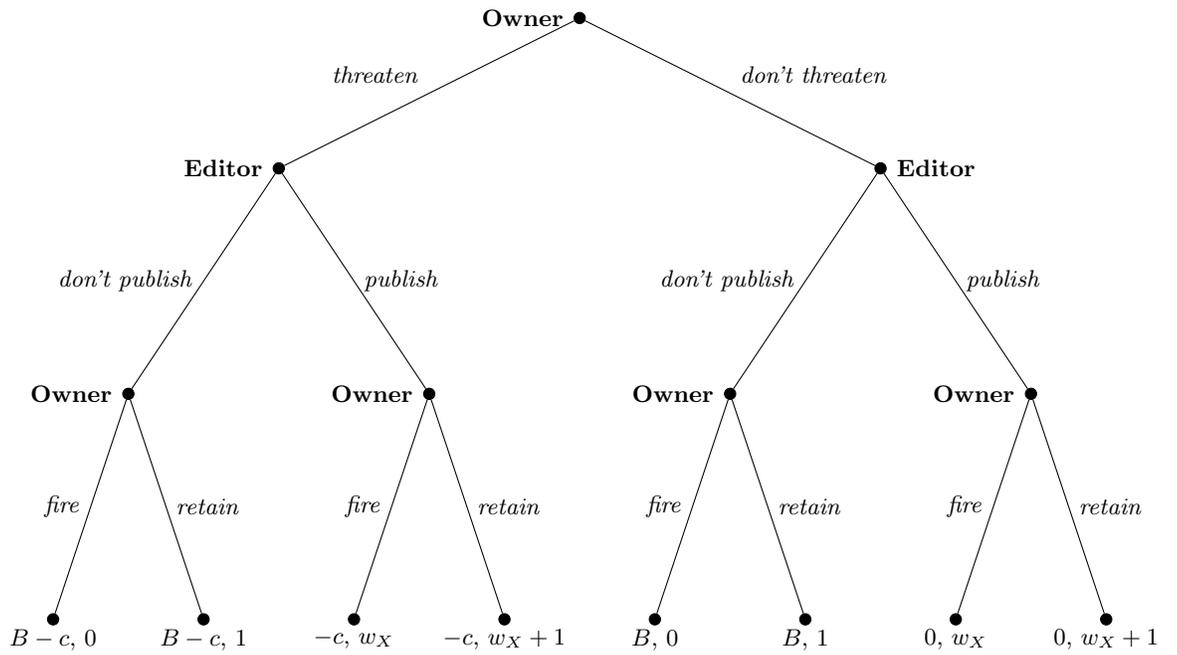


Figure 4.3: A diagram modeling the modified example in which it is costly for the newspaper owner to issue a threat.

degree of control over the newspaper. For the purposes of answering this question, hold fixed the assumption that the editor prefers for the story to be published, but places more weight on retaining her job, i.e., $w_X \in (0, 1)$. As noted above, in assessing the owner's control, in equilibrium, we hold fixed the editor's preferences. Suppose, as before, that the editor will publish the story if and only if the owner does not threaten her with dismissal. In equilibrium, the owner will choose to issue the threat if and only if $B > c$. Thus, the owner has *some* control over the story's publication. If he attached sufficient weight to this variable, relatively speaking, i.e., if $B > c$, then he would choose to issue the threat and the story would not be published. But he has *less* control than he would have if, as in the original version, it were costless for him to threaten the editor with dismissal. In the original version, his preferences determined whether the story was published, no matter how strongly they were held. Now, his preferences constrain the variable only if he attaches sufficient importance to it, that is, only if $B > c$. Thus, the editor's contractual protections against wrongful termination reduce the owner's control over what the editor chooses to publish.

Intuitively speaking, one would expect these protections to increase the editor's control over what the newspaper publishes, provided the costs of breaking them are large enough. The definition supports this judgment as well, given further assumptions about the situation. To assess the editor's control, hold fixed the assumption that $B < c$, and assume the owner's strategy is to refrain from issuing a threat and retain the editor whatever she does. Now consider the effects of variation in w_X , the parameter whose sign represents the editor's preferences over the story's publication and whose

magnitude represents the weight she attaches to this variable. Given the owner's strategy, the editor will choose, in equilibrium, to publish the story if and only if she prefers its publication to its non-publication, *even if* she cares more about keeping her job than publishing the story. That is, she will choose to publish it if and only if $w_X \geq 0$, and irrespective of the magnitude of w_X . Thus, her preferences over the story's publication now determine the value of this variable irrespective of how much importance she attaches to it. Contrast this state of affairs with the situation in the original version, where her preferences had this effect only if she attached more importance to the story's publication than to retaining her job. According to the definition, she now has *more* control than she originally did, because the minimum weight she must attach to the story's publication, before her preferences constrain it, is less than in the original version.

Consider two additional, more politically salient examples of how the definition captures intuitive judgments about control. Proponents of school vouchers, which provide parents with funds they can apply toward tuition at private schools of their choosing, often claim that they give parents greater control over where their children attend school. Advocates for women's reproductive rights often claim that access to affordable birth control gives women greater control over reproductive decisions. In each case, the judgment is that lowering the cost of a good (access to private school, birth control) increases an agent's control over some variable (where one's child attends school, whether one becomes pregnant) whose realized value can be affected by means of purchasing the good in question. The definition of control supports both of these claims. In general, reducing the cost of a

good increases an agent's control over a variable that she can affect only by acquiring the good, because it reduces the minimum amount of importance the agent has to attach to that variable, relative to conserving money, in order for the purchase of the good to be worthwhile and for her preferences to influence the variable, via the purchase.

We have so far concentrated on examples in which a single agent has control over a variable. Now let us revisit the example of two students who jointly control whether their project succeeds, so as to clarify how the formal apparatus just introduced above—variables and preferences over variables—applies to the example. Each student chooses whether to contribute or not; let $a_i \in \{0, 1\}$ represent student i 's choice. A comprehensively described outcome of the game is thus a pair $(a_1, a_2) \in \{0, 1\}^2$. The random variable, $Y : \{0, 1\}^2 \rightarrow \mathbf{R}$, defined by $Y(a_1, a_2) = 1$ if $a_1 = a_2 = 1$, $Y(a_1, a_2) = 0$ otherwise, represents whether the project is successfully completed. The random variable $E_i : \{0, 1\}^2 \rightarrow \mathbf{R}$, defined by $E_i(a_1, a_2) = a_i$, represents whether student i incurs the cost of effort. Student i 's value function is $u_i(a_1, a_2) = v_i Y(a_1, a_2) - c E_i(a_1, a_2)$, where v_i is the realization of a random variable representing the importance of the project's status, Y , and $c > 0$ is the cost of exerting effort. The reader can find the rest of the details of the game, and a Bayesian Nash equilibrium to the game, in footnotes 7 and 8. Given the definition of *preferences over a variable*, we can say that student i prefers for the project to be successfully completed if and only if $v_i \geq 0$, and each student prefers not to incur the cost of effort.

The preference condition for the students' control is satisfied, because, in the equilibrium, the students' shared preferences over Y constrain it,

provided they attach enough importance to the variable. That is, if $v_1, v_2 < 0$ (the students both prefer for the project to fail), then $Y = 0$ in equilibrium; and if $v_1, v_2 > 0$ (they both prefer for it to succeed) and, moreover, $v_1, v_2 \geq v^*$ (they attach sufficient importance to the project's success) where v^* is defined as in footnote 7, then $Y = 1$, in equilibrium. Thus, the students, as a group, have control over the project's success.

I conclude this technical section with observations on how the “negative” variety of control captured by the definition will sometimes, but not always, entail a “positive” ability to achieve particular outcomes. The definition of control refers to the constraining effects of agents' preferences—a group's shared preference for one outcome over another prevents the latter rather than causing the former. But if an agent has control over the entire range of pairs of a variable's possible values, then the constraining effects of her preferences will prevent any outcome other than her most preferred outcome, if she attaches enough importance to the variable.¹⁴ More generally, in the case of a group, if the group has control over the entire range of pairs of the variable's possible values, then the constraining effects of its shared preferences will prevent all outcomes that fail to be weakly Pareto optimal, provided its members attach enough importance to the variable.¹⁵

¹⁴A caveat to this claim is that it holds only for a variable that has finitely many possible values. See the argument in the next paragraph and footnote 16.

¹⁵We need the qualifier “weakly” because the definition of control only refers to the constraining effects of every member of group (strictly) preferring one outcome to another.

Suppose that a variable, X , can take on a finite number of values, $S = \{x_1, \dots, x_k\}$, and suppose that for some group of agents, the preference condition is satisfied for the entire range of pairs of the variable's possible values: for any two outcomes $x, y \in S$, if each member of the group preferred x to y and attached enough importance to the variable, then the variable would not take on the value y . Suppose that all outcomes in S except x_{k-1} and x_k are strictly Pareto suboptimal: for each $x_i \in \{x_1, x_2, \dots, x_{k-2}\}$, one can find an alternative in S , call it $y(x_i)$, that every member of the group strictly prefers to x_i . Since the preference condition holds, one can identify, for each $x_i \in \{x_1, \dots, x_{k-2}\}$, the threshold $t(x_i, y(x_i))$ such that if the weight each member of the group attached to the variable exceeded $t(x_i, y(x_i))$, then the variable would not take on the value x_i . Thus, we can conclude that if the weight each member of the group attached to the variable exceeded the maximum of $t(x_1, y(x_1)), \dots, t(x_{k-2}, y(x_{k-2}))$, then the variable would not take on a value in $S \setminus \{x_{k-1}, x_k\}$. And since, by hypothesis, S contains all of the variable's possible values, it would have to take on one of the weakly Pareto optimal outcomes, x_{k-1} or x_k .¹⁶

Note that a corollary of the preceding argument is that in the case of a single agent, if the agent has control over a variable that can take on finitely

¹⁶ Observe that the argument would fail if the variable can take on infinitely many values: if, say, \mathbf{R} is the set of possible values, then, even though $y(x)$ and the threshold $t(x, y(x))$ will be well-defined for each strictly Pareto suboptimal $x \in \mathbf{R}$, there may be no number corresponding to $\max\{t(x, y(x))\}_{x \in \mathbf{R}}$.

many values, and her control is maximal as far as the range of possible values goes—the preference condition holds for all pairs of the variable’s possible values—then the variable would take on her most preferred value if she attaches sufficient importance to it. In the case of a group containing just a single agent, the only outcomes that are not strictly Pareto dominated, relative to the “group’s” preferences, are those that are best according to the agent’s preferences.

Having explained how the definition of control is to be understood and applied, I turn now to several questions about the concept that are important for a theory of democratic control: whether an agent’s intentions when she exercises control are relevant; whether an agent’s control must be robust to counterfactual variation in other agents’ dispositions or strategies; and whether an agent who has no preferences over a variable can nonetheless be described as having control over it.

§4.3 Electoral control and the role of intentions

The power citizens have to remove elected officials from office—the dependence of elected representatives on the people—is often thought to give citizens power, or control, over elected officials. More generally, it is often thought that one agent will have power over another if the latter is dependent on the former for some benefit or the avoidance of some harm. Alexander Hamilton is describing the general mechanism when he writes that “in the main it will be found that a power over a man’s support is a power over his will.”¹⁷ He offers this claim in defense of the Constitution’s judicial com-

¹⁷*The Federalist* no. 73; see also no. 79.

pensation clause (article III, section 1), which prohibits reductions to federal judges' compensation during their tenures in office. If the legislature were to have a discretionary power to raise or reduce a judge's salary, then the fear was that the judge would come under the legislature's control. The same fear was part of the justification for giving members of the federal judiciary lifetime tenure. Any entity with the power to dismiss a judge at the end of a specified renewable term would have a controlling influence over judges who wish to retain their office. The claim about dependence and control is on display when Madison argues that, since the federal and state governments are all "substantially dependent on the great body of citizens," they should be viewed as "but different agents and trustees of the people," and it is an "error" to view them as "uncontrolled by any common superior." Elsewhere he tells us that the "federal legislature will. . . be *restrained* by its dependence on the people."¹⁸ In these examples, the inference being made is that if one agent (a judge, an elected legislator) is dependent on another agent (an executive or legislator, the people) for some benefit (maintaining a current salary or tenure of an office), then the latter agent can be expected to have power, or control, over him. (In this section, I use *power* and *control* interchangeably.)

One mechanism that could explain a connection between an agent's dependence on a principal and the principal's power over the agent is the principal's ability to issue instructions, backed by threats or promises, as in the example of the newspaper owner and editor. But that is clearly not the

¹⁸ *The Federalist* nos. 46, 52, emphasis added.

mechanism Madison was imagining, as he did not expect citizens to issue instructions to their elected officials. By what other mechanisms could one agent's dependence on another give the latter agent control over the former, if not the ability to issue threats and promises?

In the case of citizens and their elected representatives, the thought might be that citizens' voting behavior affects what elected representatives choose to do, because their voting behavior conveys information about what the elected officials need to do to win re-election. Elected legislators, observing, say, evidence that their predecessors' support for laissez-faire economic policies cost them votes, should be less inclined to support these policies in the future, if their objective is to win re-election. The result may be that voters' policy preferences, through their effects on voting behavior and the incentives that creates, constrain the elected official. The result may transpire without citizens intending to produce these incentives through their voting behavior.

In chapters 5 and 6, I examine more carefully the mechanisms by which elections might facilitate control. What I want to address here is a conceptual question that arises: must the actions by which an agent exercises control over another agent be undertaken with the *intention* of influencing the latter in order to count as an exercise of power, or can an agent exercise power over another unintentionally? How we answer this question matters to a theory of democratic control because it is debatable whether voters are intending to influence future elected officials when they vote to remove or retain an incumbent officeholder. The issue is not merely that it would be implausible to describe *individual* voters as intending to influence office-

holders. Even if a group of citizens all voting for the same candidate do, in some sense, have a “shared intention” as a group, their intention need not be to exercise control over the behavior of elected officials.¹⁹ They may merely intend to influence the type of person who holds office, not what officeholders will do. As we will see in chapter 6, even a principal with that intention can still exert influence over an agent in such a way that the preference condition for control holds. Thus, it is important, for this account of popular control, to defend the view that groups of citizens can exercise control over elected officials without intending to do so.

A controversy over the power of capitalists in democracies pivots around this conceptual question. Democratically elected governments are deterred from pursuing some policies—*e.g.*, heavy taxes on capital, certain labor and environmental protections, etc.—because they fear adverse responses from the owners of capital—*e.g.*, diminished investment, capital flight—that in turn may hurt the government’s prospects of re-election. According to Barry (2002), it follows that capitalists have power over democratically elected governments: they have an ability to bring it about that these governments enact the capital-friendly policies they want. Dowding (2003), who insists exercises of power must be deliberate efforts to influence, denies that capitalists have power over democratically elected governments if their decisions on how to allocate their capital are not made “*in order* to affect economic policy.” Capitalists are merely “lucky” that “in a global capitalist world,

¹⁹On the idea of democratic citizens’ having a shared intention when they vote, see Beerbohm (2012), especially pp. 45–50.

democratic governments' interests are tied to those of capitalists" (Dowding, 2003, pp. 317–318).

The definition of control I have given makes no reference to agents' intentions. On this view, voters (or capitalists, for that matter) could have control over elected officials even if they do not intend to influence elected officials' behavior by their actions. All that matters is that when voters and elected officials act in the manner one would reasonably expect, given their preferences and beliefs and the situation they are in—in game-theoretic jargon, when they act according to equilibrium strategies—voters' preferences constrain elected officials' behavior, when voters care strongly enough. On a narrower view of what counts as power, voters could be described as controlling their elected officials only if the actions by which this pattern of influence arises are undertaken with this intention. If voters' actions are undertaken with some other intention, then we could only say, with Dowding, that voters are "systematically lucky" that their voting behavior has the unintended consequence of making elected officials behave according to their preferences.

The narrow view clashes with ordinary intuition in significant cases. Congruence with ordinary intuition is not the ultimate standard for assessing a definition of control—a good definition *should* sometimes clash with ordinary intuition, if it has enough determinacy to generate unanticipated logical implications. But some degree of congruence is desirable, and it is important to see that, with respect to the role of intentions, ordinary intuition is not squarely on the side of the narrow view.

Take an example similar to what Hamilton has in mind when he speaks

of dependence and power: suppose that an autocrat has the unilateral power to remove judges from office, for any reason whatsoever. If a judge believes, on the basis of past observation, that ruling contrary to the autocrat's wishes is a sure path to losing office, then she will be deterred from acting contrary to his wishes. But suppose that the autocrat has pursued this strategy, not because he expects to shape the beliefs and incentives of future judges, but merely because an unfavorable ruling from a judge is evidence that the judge will reach unfavorable rulings in the future. The autocrat is trying to cull "bad types" from the ranks of the judiciary, which has the unintended effect of giving future judges incentives not to rule contrary to his wishes. While he knows his actions affect what future judges choose to do, this anticipated effect is not part of his intention. If we accept that an actor exercises power over another only if he intends to change the other's behavior, then we will be forced to deny that the autocrat is exercising power over the judges. But that seems absurd.

As another example, consider the claim that allowing corporations to spend without limit on behalf of elected officials' re-election campaigns will expand their power over elected officials. In his dissenting opinion to *Citizens United v. Federal Election Commission*, Justice Stevens warned that the decision "may well promote corporate power" and "corporate domination of the electoral process."²⁰ But the worry is not merely that it extends the abilities of corporations to undertake actions with the intended effect of

²⁰ *Citizens United v. Federal Election Commission* 558 U.S. , ___ (2010) (p. 85 in the Opinion of Stevens, J., slip opinion).

influencing the behavior of incumbents. The mechanism of corporate domination that *Citizens United* is thought to strengthen is not, or at least not only, *quid pro quo* arrangements or even subtle winks and nods by which a corporation might intentionally influence an officeholder. It may be that corporations fund campaign advertisements solely with the intention of influencing election outcomes and not with the intention of influencing the behavior of elected officials: they may spend on behalf of the candidates they wish to see elected, because their expenditures make the desired election outcome more likely. That is, their expenditures may be intended merely to influence who holds office, not how incumbent officeholders behave once elected. Yet, since an incumbent's actions in office will convey information to future corporate donors about how he is likely to behave if re-elected, his behavior is likely to influence corporations' preferences over the outcome of his future re-election campaign and, consequently, to influence whether they spend money on his behalf. He will therefore be deterred from acting contrary to corporations' preferences. In this case, it would be natural to conclude that corporations have power over the elected official's behavior, even if the power is not exercised with the intention of influencing his behavior.

The most plausible basis for retrospective voting in elections would be voting strategies analogous to the autocrat's and corporation's filtering strategies. In such a "selection" model of retrospective voting, voters are not attempting to influence the behavior of future incumbents. Instead, they are merely trying to filter out the incumbents who have betrayed themselves as the wrong type. It is a by-product of this filtering strategy that future

elected officials, contemplating actions that they know voters will treat as evidence of their type, are deterred from acting contrary to voters' wishes.

My point here is independent of the descriptive accuracy of the retrospective voting story. Whether it is true, it is a story about electoral *control* if we adopt the encompassing view according to which exercises of power and control can be unintended. And one should adopt the encompassing view if one wishes to say of the preceding examples that the autocrat has power over how the judges rule, and that the corporation has power over how the elected legislator votes.

Note that I am not denying that intentions must lie behind agents' actions. The claim is not that power or control can be exercised by means of non-intentional, involuntary behavior, but rather that power and control over a variable can be exercised by means of actions the intentions of which are not to influence the variable in question.

Apart from congruence with ordinary intuition, the narrow view would also be less useful for our purposes. We use concepts like control to organize and communicate our thoughts about the world, in particular to refer to shared properties of classes of situations. It is often a relevant fact about a situation that one agent has actions by which she can determine the value of some variable, such as how another agent behaves. It is also relevant when, as a consequence, the variable's realized value depends on the agent's preferences and how strongly she cares about the variable relative to other concerns. These facts are often significant, from predictive as well as normative points of view, even if the agent is unlikely to choose her actions with the intention of influencing the variable in question. The concept of

control has been defined so that we can talk easily about these facts and just these facts, not other facts about intentions. Even Dowding, in his debate with Barry over the power of capitalists, admits that in most cases there is no reason to care whether a government chooses capital-friendly policies because “it fears a deliberate response from international financiers” or instead “fears that its actions will merely elicit a response that will be bad for the economy” (Dowding, 2003, p. 318).

§4.4 Whether control must be robust

I return now to an issue briefly considered above, namely whether an agent’s control must be robust to changes in the dispositions of other agents. I have claimed that it need not be. In general, one agent enjoys control—including control over other agents—only in virtue of other agents’ dispositions. In this section I consider the alternative view, using Pettit’s (2012) recent characterization of control as an example.

Pettit characterizes control as “directed influence.” One has control over the result of a process if one has influence over the result and the influence imposes “a direction on the process, helping to ensure that a suitable result transpires.” The traffic cop, for example, not only influences but also controls the flow of traffic. Anyone can accomplish influence over its flow by planting himself in the center of an intersection, waving his hands, and causing chaos. The traffic cop also controls the flow of traffic, because his influence imposes a suitable direction, or pattern, on it. Control is always more than mere influence, whether we have in mind the control of an indi-

vidual over the flow of traffic, or popular control over government. Control is a particular variety of influence that leads to a particular kind of result (Pettit, 2012, p. 153).

On Pettit's account, control might be intentional, as it normally is with agents, but "purely mechanical, non-intentional mechanisms" can also have control. "An example might be the cooling-heating system that keeps the temperature in a room within a certain range"—it influences the temperature and the influence "makes a designed difference but without any intentionality on the part of the mechanism itself." Even agents may exercise non-intentional control, as in

the patient in a coma whose wishes, thanks to family or friends, control for the treatment provided by the surgeons at the hospital. Given the form taken by the patient's wishes, the surgeons provide one sort of treatment. And had the wishes taken a relatively different form. . . then the treatment would have been different too. . .

In this case, the agent's desires "are effective by virtue of the efforts of others, not his or her own" (Pettit, 2012, pp. 153–155).

In place of the idea of an agent's preferences constraining a variable in the manner described by the preference condition, Pettit's definition uses the more general notion of a variable (or the outcome of process) exhibiting a pattern or direction. When an agent's preferences constrain a variable, preventing it from assuming a particular value when the agent prefers a different value and attaches sufficient importance to the variable, this is

just one example of a pattern that the agent's influence may impose on the variable. But an agent's influence might impose other patterns, and these could also count as control according to Pettit's definition, but not the definition I have given. Pettit's concept of control therefore applies in instances not covered by the concept as I have defined it, such as in the case of non-agential systems like thermostats.²¹

Along another dimension, Pettit narrows the extension of his concept. Here the contrast with the definition I have given is more significant. Pettit claims that a necessary condition for one agent to have control over another is that "the influence exercised leads to the required result independently of the will of the controlled agent, or indeed of any third party." If *B* obeys *A*'s instructions but only because *B* wants to humor *A*, then *A* cannot be said to have control over *B*'s actions, according to Pettit. His control may depend on various contingent circumstances, such as his "greater physical strength or cultural clout," or on the fact that "*B* is in thrall to [his] charms." But "it cannot depend on *B*'s willingness to play the part" (Pettit, 2012, pp. 170, 171). It is not genuine control if it does.

By contrast, my definition allows one agent's control to be contingent

²¹This difference is arguably minor, because what Pettit's concept records as literal control, the concept of control as I have defined it can register as metaphorical control. For certain pairs of possible temperatures, if the thermostat "preferred" one temperature to the other, then it would take "actions" that prevent the air from stabilizing at the "less preferred" of the two temperatures.

on other agents' willingness to tolerate it. Suppose a factory owner allows his workers to take their allotted breaks whenever they please. If the worker wants to take his lunch break at noon, then he will; were he to want instead to take it later in the afternoon, then he would. His preferences over the timing of his break determine when he takes it (however much importance he attaches to the variable), so the preference condition is satisfied. Hence, my definition implies that he has control over the timing of his break, even though his control depends on the owner's willingness to give him control.

As far as congruence with ordinary language goes, it is a *virtue* and not a defect of the definition that it has this implication. A driving instructor lets his pupil take control of the car, but retains the ability to resume control at his discretion; the pupil has control over the car's movements—his preferences over its movement determine where it goes—but only on condition that his instructor tolerates his control. A parent gives his child control over the radio station during the ride to school; during the ride, the child's preferences determine the station, but only because and so long as the parent cares to let the child have control. The conductor of an orchestra, if he is any good, has some control over the performance of the symphony even if he enjoys his influence, not because the musicians are under duress or “in thrall to his charms,” but simply because they are willing to “play their parts.” There is no abuse of language here. We should concede that the worker has control over his schedule just as the pupil has control over the car and the conductor has control over the orchestra, even if the control is not robust to changes in the dispositions of the owner, instructor, and musicians, respectively.

To be sure, it *matters* whether one agent's control is independent of other agents' willingness to cede her control. The advocate for workplace democracy wants workers to have control over their workplace whether management likes the idea or not. This observation is not, however, a reason to fold a robustness requirement into the definition of control. Just the opposite. The thought that workers' control over their workplace should not depend on the indulgent dispositions of management would be an *ill-formed* thought if we accepted Pettit's account. It would be like the thought that it is important not to work more than 24 hours a day. That workers' control over aspects of the workplace ought not to be contingent on managements' goodwill is a meaningful thought only if it is possible for their control to be contingent on management's goodwill. Pettit's account rules this possibility out by stipulation.

In debates about who should have control in these contexts, it often goes without saying that the type of control at issue is control that would be independent of other agents' preferences. Advocates for workplace democracy praise workers' control over the workplace without adding that this control should not be conditional on their managers' dispositions to accommodate it. There is no reason for the additional stipulation because it is obvious from the context that they are advocating for robust worker control. This is perhaps what explains the errant intuition that there is no such thing as non-robust control: one mistakenly concludes that all control must be of this kind because it is so often the case, at least in political theory, that the *practically relevant* kind of control is control that is robust to changes in certain other agents' dispositions. But this mistaken conclusion must be rejected,

because it would force us to reject as nonsensical the perfectly intelligible thoughts about the pupil's control over the car, the child's control over the radio, or the conductor's control over the orchestra. Moreover, it would render the thought that workers' control over their workplace should be robust to changes in management's dispositions as ridiculous as the thought that no one should have to work more than 24 hours a day.

This point will be important when we turn to models of democratic control. For the control that groups of citizens exercise over their elected leaders, or over policy outcomes, will depend on the dispositions of their elected leaders to respond to electoral incentives. If elected officials were indifferent to the prospect of re-election, then citizens' preferences would be unlikely to influence their behavior. Yet the mere fact that citizens' preferences would be inconsequential if elected officials were indifferent to re-election incentives is not a good reason to conclude that citizens lack control over them in the actual world where they *do* respond to these incentives.

Rather than amend the definition of control with a robustness requirement, we should keep the given definition but bear in mind that for normative reasons it often matters whether one agent's control is robust to counterfactual changes in the dispositions of other agents.

§4.5 Indifferent and divided citizens

Control, as I have defined it, is "desire-independent": whether an agent has control does not depend on the preferences she happens to have.²² The

²²Lovett (2010, p. 68) makes the same observation in his analysis of power.

preference condition requires that counterfactual variation in the agent's preferences have the right kind of effects on the variable under the agent's control; it does not require that the agent have any particular preferences over the variable. In particular, the definition of control allows an agent to have control over a variable even if she does not care what value the variable assumes.

This is a virtue of the definition. Even the indifferent slave owner, who has no preferences concerning how the slave behaves, still has power and control over the slave, as neo-republican theorists of freedom as non-domination recognize (Pettit, 2012; Skinner, 1998). That is true because if, *counterfactually*, the master preferred the slave to do one thing rather than another, he would issue a command and the slave would obey. Thus, the preference condition for control is satisfied. The condition is satisfied for the master who is, in actual fact, indifferent, but satisfied in virtue of what would happen if the master were not indifferent. To borrow a useful analogy from Pettit's discussion of control: when an agent is indifferent towards a variable, but these counterfactual conditionals hold with respect to it, then the agent has control in the manner of a rider who is content to let his horse go where it will, but who could steer the horse this way or that if (counterfactually) he cared about the horse's direction (Pettit, 2012, pp. 156, 157).

Likewise, the people could have control over policy even if they have no policy preferences. Citizens of modern mass democracies, most of whom must work for a living and cannot devote significant time to public affairs, lack determinate opinions on many policy questions. Skeptics like Joseph Schumpeter conclude from such observations that actual democracies do

not involve popular control over policy, or anything resembling “popular rule” (Schumpeter, 1942). But this inference is invalid: the mere fact that voters lack determinate preferences over policy does not imply that they lack control over it, any more than the slave-owner’s indifference towards the slave’s actions implies that the slave is not under his control. In the case of popular control in a democracy, what matters is whether citizens’ preferences constrain policy if and when they have determinate preferences.²³

An analogous point is true of the collective control of groups whose members have conflicting preferences. A group has control over a variable, according to the definition, if for some pairs of possible values of the variable, *if* every member of the group preferred one value to the other and they all cared enough about the variable, then the variable would not take on the less preferred value. This condition may hold, and the group may therefore have control over the variable, even if, in actual fact, the group’s members do not share any preferences over the variable’s possible values. What matters is that if they did, and they attached enough importance to it, then their shared preferences would constrain the variable.

This observation affects how we think about the importance of heterogeneity in citizens’ preferences. The entire body of citizens, considered as a single group, may have control over their elected officials even if different citizens want their elected officials to do different things. Collective control over elected officials requires, not that all citizens have shared preferences,

²³For a closer examination of Schumpeter’s argument and a critique along the lines indicated here, see Ingham (2016*a*).

but rather that if they were to share preferences over a decision, and the decision mattered enough to them, then their shared preferences would constrain the decision.

Disagreement among citizens does not undermine the possibility of the entire citizen body having collective control, but it does diminish its practical significance. Imagine elections that were decided by an unanimity rule where the incumbent is re-elected unless everyone votes against him. It might then be true that *if* every citizen preferred a challenger to the incumbent, then citizens would turn out to vote and remove the incumbent from office by all voting for the challenger. The definition would then entail that the body of citizens has control over the election outcome. Yet their collective control would be practically worthless. The if-then proposition in the preference condition may be true, even though the counterfactual scenario described by its antecedent clause is wildly improbable. The citizen body would have control if the proposition is true, but its control would be insignificant, practically speaking, because the antecedent is so unlikely to hold. The incumbent could behave just as he would if he were not subject to their collective control.

What this observation demonstrates is not a problem with the definition of control, but rather the insufficiency, from a practical or normative point of view, of the entire citizen body having collective control. It is surely *necessary* for the entire citizen body to have control over government if government is to be considered subject to popular control. But it is *insufficient* for any meaningful notion of popular control. We can arrive at a more demanding and normatively significant conception of popular control

if we identify it with, not merely the control of the entire citizen body, but instead every majority's simultaneous control. I develop this view in the next chapter.

§4.6 Conclusion

A concept of popular control should be based on a general analysis of control, one that captures our intuitions about control outside of political contexts. According to the analysis given here, a group has some control over a variable if, for some range of pairs of its possible outcomes, if the group preferred one outcome to the other and held their preference with sufficient intensity, relative to their other concerns, then the less preferred outcome would not occur (the preference condition). All else being equal, the degree of a group's control increases with the range of possible outcomes for which this condition holds, and it is inversely proportional to the threshold of preference intensity at which the group's preferences would have the constraining effects described by the preference condition. This concept of control yields intuitively plausible judgments in a range of cases outside of politics, so it is not unreasonable to describe it as a concept *of control*.

Its most important implication for a theory of popular, democratic control is that two groups can have control over a variable at the same time, even if the variable is bound to frustrate the shared preferences of one of the groups. In particular, all majorities can have simultaneous control in a democracy, even if no Condorcet winner exists and decisions are bound to frustrate the preferences of some majority, as I explain in the next chapter.

Chapter 5

Rule by multiple majorities

The objection that popular control is impossible rested on two premises. The first was an assumption about the concept of popular control, namely that if a variable is under popular control, then it does not frustrate majorities' preferences. The second was an assumption about the world, namely that Condorcet winners often do not exist. I argued in chapter 3 that we should concede the latter premise, given the bleak prospects of refuting it empirically. This left us to examine the conceptual premise, which we are now in a position to do.

Having proposed a definition of what it means for a group to have control, we can now use this definition to give an account of popular control that is not vulnerable to the impossibility objection. In the next section, I describe the “multiple majorities” view of popular control, according to which a variable is under popular control if all majorities have control over it, at the same time. After presenting the view and contrasting it with alternative conceptions of popular control, I explain the reasons for accepting the multiple majorities view, its implications for voting rules, and its relationship to the concept of democracy. I then revisit the impossibility objection, from

Chapter 3, and discuss a model of electoral competition that illustrates how elections could, in theory, bring policy under popular control. I conclude by responding to possible objections to the argument.

§5.1 The multiple majorities view of popular control

According to the multiple majorities view, a variable (such as policy) is under popular control if every majority has control over it. In Chapter 7, I will suggest that in a fully democratic system of popular control, every majority must have a degree of control that is adequate and comparable to the degree of control enjoyed by other majorities of similar size. These further qualifications are needed to characterize a compelling normative ideal of democratic control, but they can be largely ignored for now. The main purpose of this chapter is to explain what it means for multiple majorities to have control simultaneously, why we have reasons to understand the concept of popular control in those terms, and why this understanding of popular control is compatible with the insights from social choice theory that cause difficulties for traditional views.

Here and throughout the book, when I refer to “a majority,” I just mean any subset whose members form a numerical majority of the whole. They may not be organized into a political party. They may not have any understanding of themselves as a political force. They may not even be aware of any interests or values they share. Whether they have control does not *by definition* depend on any of these considerations, although it could depend on them as a contingent matter, by virtue of facts about the world.

Since a majority is just a group of citizens, we can apply the definition from the previous chapter to judge whether a majority has control. A majority has control over a variable if, for some pairs of its possible outcomes, when the members of the majority share a preference, their shared preference constrains the realized outcome, provided they consider the variable sufficiently important. The majority has more control, the less importance they have to attach to the variable for the purposes of satisfying the preference condition.

A group's shared preferences "constrain" the realized outcome if, for some pairs of alternatives, they prevent the less preferred of the two from being the variable's realized outcome. For example, suppose that the minimum wage can be set at any multiple of \$.25, and a majority's control is such that the preference condition holds for any pair of outcomes that are between \$0 and \$20. The members of a majority may disagree about the optimal level of the minimum wage, but they may all prefer a wage of \$7 to any lower wage level, and they may all prefer a wage of \$12 to any higher level. Then the wage level will be constrained to lie in the interval between \$7 and \$12, provided its members care enough about the minimum wage relative to other variables.¹ This is the set of Pareto optimal wage levels (relative to the preferences of this group). If they do not share any further preferences over the levels within this interval, then their control is compati-

¹This claim depends on the assumption that there are only finitely many possibilities in the range of pairs for which the preference condition holds; see the end of section 4.2 for explanation.

ble with the variable taking any value in the interval. The majority's shared preferences thus act as a kind of *veto*, preventing less preferred options from being chosen, but without fully determining the variable's value.

It may be that a variable fails to respect a majority's shared preference over two possible values but only because its members hardly care about the variable. *The majority may still be said to have control so long as its shared preferences would act as a kind of veto, provided its members attached sufficient importance to the variable.* How much control a majority has depends on how much importance its members must attach to the variable before their shared preferences have these constraining effects. Do their shared preferences constrain the variable no matter how little they care about it? Then they have maximal control along this dimension.² Do the constraining effects materialize only when members of the majority are willing to give up everything for the sake of satisfying their shared preferences over the variable? Then they have minimal control.

As an analogy to multiple majorities' having control over a variable, imagine a bureaucrat who must design a policy that meets with the approval of multiple oversight committees. An oversight committee can veto a policy if its members can identify an alternative that they all prefer to it. But exercising this veto power is costly—they have to write up reports and cut

²Recall that another dimension along which control varies, according to the definition from chapter 4, is the size of the range of pairs of possible outcomes for which the preference condition holds—the greater the range, the greater the control.

through a lot of red tape—so a committee only exercises its veto power when its members feel sufficiently strongly about policy relative to the costs of cutting through the red tape. Intuitively speaking, the committees have a kind of control over the bureaucrat’s choice of policy. And a committee has less control, the more bureaucratic red tape they have to cut through in order to exercise their veto power. The more red tape, the stronger they have to feel about the policy variable before the constraining effects of their shared policy preferences come into play.

Like the oversight committees, different majorities can have control over policy, at the same time. When an elected official’s actions are under popular control, she is subject to the oversight of all majorities. These majorities are like committees with overlapping membership. They can each “veto” her actions when their members prefer a different course of action and attach sufficient weight to her actions, relative to other variables they care about.

One sometimes encounters a superficially similar idea in the literature on social choice theory. In commentary on the phenomenon of cyclical majorities, writers sometimes suggest that majority cycles are normatively unproblematic in so far as they allow “rival majorities each to have their way for a time” (Pildes and Anderson, 1990, p. 2185). McGann (2006) goes so far as to suggest that majority cycles are *necessary* for a well-functioning democracy, because, “in the absence of cycling, there may be permanent winners and permanent losers” (p. 25). Here the idea is that majority preference cycles allow for situations in which different majority coalitions alternate in power. We might call it the “rotating majorities” view of popular rule. The view I am putting forward is different. It is not the idea that different ma-

majorities take turns at the helm, first one majority implementing its agenda, then a different majority imposing a different agenda. The idea is rather that every majority subset of the citizen population can have *simultaneous* control over a variable.³

One might ask how it is possible for multiple majorities to have control at the same time, just as one might wonder how multiple oversight committees could all have control over a bureaucrat at the same time. What if every committee is willing to cut through the requisite red tape and each of the bureaucrat's possible choices—even choosing to “do nothing”—is ve-

³The present account also differs from my initial statement of a similar idea in Ingham (2016*b*). In that article, a group was said to have control if the outcome of its influence was “acceptable,” and to have more or less control depending on how acceptable the outcome was. Popular control was identified with majorities' having maximal control. As I now see it, a major shortcoming of that concept of control is that it wrongly directs attention to the relationship between an *actual* outcome and *actual* attitudes rather than the manner in which the outcome depends *counterfactually* on attitudes. To assess a group's control, we need to know how the variable would be different if their preferences or the intensity of their preferences over the variable were different, not whether the actual realization of the variable is sufficiently acceptable according to the group's actual preferences. Knowledge of the latter relationship is insufficient for attributing control even if we also know that the relationship is, in some unspecified sense, an effect of the group's influence.

toed by one of the committees, entailing the nonsensical conclusion that the bureaucrat doesn't choose anything? What if each majority cares enough about policy and their shared preferences cannot all be satisfied at the same time?

The analogy with the bureaucrat breaks down at this point because one naturally assumes that the amount of red tape a committee would have to cut through, in order to exercise its veto, is "exogenous" and fixed independently of the preferences and actions of the other oversight committees. In this respect, the analogy does not extend to the multiple majorities conception of popular control. To give an illuminating answer to our question, we will have to depart from the analogy with the bureaucrat and consider a concrete model of democratic decision-making, in which all majorities have control at the same time. In such a model, how much a majority has to care about policy in order for its shared preferences to have a constraining effect will be "endogenous"—it will depend, in part, on the preferences and behavior of other groups of agents, including the corresponding minority. I discuss such a model in this chapter. In preview, the answer to the question is that it can be true of each majority, at the same time, that if its members cared even *more* about policy than they actually do, then their shared preferences would constrain the chosen policy. To appreciate this possibility, recall from the previous chapter that when we assess this counterfactual conditional, we are to hold fixed the assumptions about other agents' preferences and strategies. Owing to that feature of the definition of control, both the owner and editor could have control over the story's publication, at the same time, even though the outcome cannot simultaneously satisfy their

actual preferences. For analogous reasons, every majority can have control over a variable at the same time, even if the variable cannot simultaneously satisfy every majority's *actual* shared preferences.

It is useful to compare the multiple majorities view with two alternatives: first, the “populism” that Riker (1982) took as his target, and, second, what I described as the standard view of popular rule in chapter 2. Riker's populist holds that in order to have popular rule, policy must reflect the will of the people. The standard view of popular rule is that it consists in the rule of the majority, which requires policy to reflect the will of the majority.

The multiple majorities view differs from these alternatives along two dimensions, giving different answers to the following questions. Which group(s) must have control over a decision for the decision to be under popular control? And what must be true of a group in order for it to have control?

Start with the question of who must have control over a decision in order for the decision to be under popular control. For Riker's populist, the answer is the People, acting as a kind of group agent. On the traditional majoritarian view, the answer is “the majority,” as if there were just one particular majority empowered under democracy. According to the multiple majorities view, the answer is that every subset of the citizen population that makes up a numerical majority must have control. Whether any of these majorities resembles a group agent is neither here nor there.

Now compare how each view characterizes what it means for a group to have control. Riker's populist answers that the people have control over policy—the “people rule”—only if policy respects the *popular will*. The

traditional majoritarian answers, in a similar vein, that the majority has control—the majority rules—only if policy respects the *will of the majority*. The multiple majorities view differs from these views in eschewing the concept of a group will and in introducing a measure of intensity of preference. It asks, not whether a decision respects a majority’s will, but instead whether a majority’s shared preferences constrain a decision—whether, when a majority prefers one alternative to another, that fact prevents the latter alternative from being chosen. And when a majority’s preference fails to constrain a decision—as necessarily happens, in the case of some majority, when no Condorcet winner exists—it instructs us not to conclude that the majority lacks control, but rather to ask whether the majority’s shared preferences *would* constrain the decision *if* the majority cared enough about the decision relative to other variables they care about. When this question about the counterfactual admits an affirmative answer, the conclusion is that the majority has control.

The multiple majorities view of popular control thus shows how one can divorce the idea that democracy requires popular control—“popular rule,” if one likes—from the populist’s appeals to “the people” and “the popular will.” Unburdening the idea of the populist’s conceptual and ideological baggage makes it less vulnerable to skeptical challenges to its internal coherence, in the spirit of Schumpeter’s (1942) and Riker’s (1982) critiques of populism. It also means one can defend the ideal of democratic popular rule, or democratic popular control, without resorting to the rhetorical devices that populists deploy—with increasing success in today’s democracies—for illiberal ends (Mounk, 2018; Müller, 2016). Indeed, as I will argue in

Chapter 7, if one accepts the multiple majorities view, then respect for the democratic rights of minorities, far from a limit on the power of “the majority” or “the people” as populist rhetoric may suggest, is actually a necessary condition for democratic popular control.”

§5.2 Why identify popular control with majorities’ control?

Several intuitive judgments about the concept of popular control lead to the multiple majorities view. The first judgment is that facts about popular control are just facts about which groups of citizens have control and how much and what kind of control they have: if two situations are alike as far as which groups of citizens have control and the quality and quantity of their control, then one’s judgments about popular control in the two situations must be the same. When we ask whether a variable is under popular control, we are just asking a question about which groups have control over it and the quantity and quality of their control. That is the question that a candidate conception of popular control answers.

The second judgment is that if a variable is under *popular, democratic* control, then the distribution of control across different groups must be minimally egalitarian: if one group has control over a variable, then any other group that is the same size or larger should also have control over it, all else being equal. Now, if one believed that two disjoint groups could not possibly both have control over a variable at the same time, then one would have an explanation for why democratic, popular control is compatible with asymmetries in control between equally sized groups. On our analysis of

control from chapter 4, however, disjoint groups can enjoy simultaneous control. Asymmetries between groups, where one group has control but an equally large or larger group does not, would therefore undercut the democratic credentials of the distribution of control.

The two judgments taken together entail that popular control consists in the right groups' having the right amounts and kind of control, where the collection of relevant groups has the form "all groups of size k or larger." I will postpone, until chapter 7, questions about the degree and kind of control that might be required; here I want to consider the different possible ways of filling in the first part of the concept of popular control—which groups that must have control, in order for something to be under popular control? Distinguish the possibilities as follows: (1) minorities of a sufficient size (and all larger groups) have control, (2) all majorities have control, and (3) supermajorities of a sufficient size (and all larger groups) have control.

We have to assess the plausibility of each possibility using some account of what it means for a group to have control. The previous chapter argued for one understanding of the concept. A group has control over a variable if, for some range of its possible values, their shared preferences (if they have any) constrain its realization, provided they all care enough about the variable relative to other concerns. Given this understanding of control, the middle option accommodates some intuitive judgments about popular control better than the other options.

Consider first an extreme version of the third possibility, in which one identifies popular control with the control possessed by the entire citizen body, as a single group (the entire population being the largest supermajor-

ity of itself). This view fails to accommodate the intuition that if a public official's choice of policy is under popular control to the greatest extent possible, then the public official will not normally enjoy unfettered discretion to choose whichever policy he likes. To see that the view fails to support that intuition, suppose all citizens care strongly about an official's choice of policy, but citizens do not share preferences over any pair of his actions: for each pair of actions, one can find at least one person who prefers the first to the second and at least one who prefers the second to the first. That supposition is very likely to be met in any large and diverse population. Then even if it is true that the citizens' shared preferences would constrain the official *if* they had any shared preferences, the official will enjoy total freedom to choose as he pleases. That will be the case even if the entire citizen body has *maximal* control over the official's choice of policy—even if their shared preferences would have these constraining effects, no matter how little they all cared about policy. Even though they have maximal control, the practical implications are trivial because they will almost never share preferences over any pair of the public official's possible actions. Popular control would be an objectionably lax standard if the concept could be applied in this way. Identifying popular control with the control of smaller supermajorities invites the same objection, although the problem becomes less acute as one considers smaller and smaller supermajorities.

Turn now to the other class of alternatives to the multiple majorities view, those that identify popular control with the simultaneous control of all minorities of sufficient size and all larger groups. One might associate this view of popular control with Robert Dahl and the pluralists:

Elections and political competition do not make for government by majorities in any very significant way, but they vastly increase the size, number, and variety of minorities whose preferences must be taken into account by leaders in making policy choices. I am inclined to think that it is in this characteristic of elections—not minority rule but minorities rule—that we must look for some of the essential differences between dictatorships and democracies. (Dahl, 1956, p. 132)

However, the conception of “rule by minorities” that I am contrasting with rule by majorities differs from what Dahl appeared to have in mind. His idea was that different minorities would deploy their resources to gain influence over the different, particular policy issues they cared about, but on any one particular policy issue, there would be a particular minority holding sway (Dahl, 2005). The view I am entertaining (and rejecting) is instead that in a system of popular control, all sufficiently large minorities would have *simultaneous* control over the same variable (the same decision, or the same elected official, as the case may be).

As with the first category of views, it is instructive to consider the most extreme version of the view. For this category, the limit case is the view according to which popular control requires every individual to have simultaneous control, a group containing one individual being the smallest minority possible. The plausibility and even the logical coherence of the view depends on what it means for an individual to have control. On the definition of control proposed in the previous chapter, the view is logically

coherent but implausibly demanding.

To see its logical possibility, imagine, as a thought experiment, that the right to decide the fate of a law is auctioned off to the highest bidder, and suppose the incentives of the auction are such that, in equilibrium, individuals submit bids reflecting their monetary valuation of the decision right, so that whoever values the right the most will win the right to make the decision. Then it will be true of every individual that his decisions would constrain (indeed, fully determine) the decision—the law would be enacted if and only if he preferred for it to be enacted—provided he attached enough importance to the decision. In assessing the truth of this counterfactual conditional, we have to hold fixed the actual preferences of the other players, as explained in chapter 4. Given their preferences, and thus their resulting bids, there is some amount of money that the individual could bid to win the decision right, and it is an amount he would bid if he valued the decision right more than this amount of money. Thus, in such an equilibrium, every individual would have some control over the decision, according to the analysis from the previous chapter.

It is logically possible for everyone to have simultaneous control over a decision, but the ideal of popular control would be extremely demanding if it required this. Decisions made by ballot initiative or some other kind of plebiscite are normally described as decisions under popular control. But those decisions are not reached through a mechanism, like the auction mechanism, that would ensure that each individual has control over the decisions. The right to make such decisions *could* be auctioned off, giving each individual a degree of control. Quadratic voting, discussed below, would also result

in each individual's having control over the decision. But that is of course not what happens with ballot initiatives, and we are happy to describe them as examples of popular control nonetheless. A concept of popular control must match at least some of our intuitive judgments about popular control, otherwise it would be misleading to label it as a concept *of* popular control, as opposed to something else (Lovett, 2010, pp. 4, 5). The intuitive judgment that ballot initiatives are means of bringing decisions under popular control—despite the fact that they do not give every individual simultaneous control over the decisions—is arguably one of the fixed intuitions that a compelling account of popular control should respect. (Whether it is desirable to use ballot initiatives to this end is an entirely different question.) An account that identified popular control with every individual's simultaneous control would not respect that intuition.

We are left with the view that popular control requires that every majority have control. It is important to see that the definition of control from the previous chapter does most of the heavy lifting in this argument. The definition implies that disjoint groups can have control at the same time. Together with the egalitarian commitments, and the assumption that popular control just consists in the relevant groups' having control (of the right amount and kind), that lets us narrow down the field of contenders to three possibilities: the multiple majorities view and two families of alternatives, all identifying popular control with groups of size k and larger having control, and differing only in their specification of the threshold k .

When we considered the idea that popular control requires merely that the entire body of citizens have control, our definition of control was again

responsible for eliminating versions of this view that have appealed to some people. Suppose one took the view that a group has control over a variable if it can make the variable respect its collective will, and the collective will of the group is defined by some preference aggregation rule like the Borda count (Saari, 2003). Then one could identify popular control with the control of the entire citizen population, and popular control would be a demanding constraint, requiring policy (or whatever it is that is under popular control) to match the policy that is best according to the preference aggregation rule.

But there is little to recommend *that* definition of group control. No one would be willing to use that definition in contexts outside of politics. No one would say, for example, that two parents have control over their toddler if and only if the toddler's actions are optimal according to the Borda count's aggregation of the parents' preferences. And it would be ad hoc to maintain that a preference aggregation rule like the Borda count is *by definition* implicated in a group's collective control just in the context of democratic politics, but not in other contexts where groups are said to have control. (The Borda count could of course be an *instrument* of a group's collective control in one context but not another, but that is a different matter. Here we are asking about the concept of control, not the instruments for achieving it.) To avoid the charge of being ad hoc, one should build a conception of popular control on top of a general analysis of the concept of control, one that is plausible outside of the democratic context. That foundation is what the arguments of the previous chapter supplied.

Pettit's (2012) account of popular control also has a foundation in a

general conceptual analysis of control. On his account, an individual is said to have control over a process if he has influence over it and his influence imposes an acceptable direction on the process. As Pettit characterizes popular control, the state is under popular control if the people have a certain kind of democratic influence over it and their influence imposes an “equally acceptable direction” on the state.

On this view of what it means for the entire body of citizens to have control over the state, it amounts to a demanding condition, perhaps impossibly so. For why think that the effect of popular influence will be a direction that *everyone* finds acceptable? Pettit stresses that the term *acceptable* is to have “a non-normative sense, implying that the object or policy or whatever is such that people are disposed to accept it; they find it acceptable, as we say” (p. 170). But some people have demanding standards for what they are “disposed to accept” in politics.⁴ It seems unlikely that a system of popular influence will push the state in a direction that they are all disposed to accept.

I have argued elsewhere that Pettit ends up conflating two distinct conceptions of popular control, substituting a less demanding ideal for the official conception when he discusses the model of democracy that might implement popular control (Ingham, 2015). The less demanding version

⁴This is true even of those who are “prepared to live on equal terms with each other,” the subset of citizens to whom, in the more careful statements of the view, the result of popular influence must be acceptable (Pettit, 2012, p. 170).

requires only that popular influence impose a direction on the state that satisfies norms of policy-making that everyone accepts. However, the direction resulting from popular influence could be unacceptable to some people, even if it satisfies all those norms that are accepted by everyone. A person might find the result unacceptable because it violates a norm that she but not others endorse, even though the result satisfies all *universally* endorsed norms.

Both versions of Pettit's view differ from the one I am putting forward in the underlying concepts of control on which they are built. The less demanding of the two versions of Pettit's view is reasonable enough, picking out a set of conditions whose satisfaction or violation by a regime we might have good reasons to want to record. And it resembles well enough our ordinary intuitions about popular control that it is not misleading to use that term in our record-keeping. The main purpose of the preceding observations is not criticism, but merely to bring out the contrasts between Pettit's account and my own.

The features of our concept of control, as defined in chapter 4, also explain the differences between the multiple majorities view and yet other understandings of popular control. Some political theorists may be inclined to define concepts like popular control, or popular rule, in terms of decision-making procedures.⁵ For example, a proceduralist definition of popular

⁵Tuck (2016), for example, appears to use the term *popular control* to refer to a kind of plebiscitary procedure (see pp. 197–199, 208). With most writers, it is hard to be absolutely sure that a proceduralist view underlies their use of the term, as it is usually also possible to read them as intending

control might be the following: a decision is under popular control if the decision is made by a majoritarian voting mechanism, where a majoritarian voting mechanism is one in which, for any decision d , and any majority M , there are ballots available to the members of M such that if they cast these ballots, the result will be that the decision d is taken. A proceduralist definition would be analogous to, or the extension of, the views of individual power that identify an agent's power with formal properties of the menu of actions available to the agent.

The account of popular control I am proposing is not proceduralist in this sense. The reasons to reject proceduralism, as I have defined it here, are analogous to the reasons, set out in chapter 4, to reject the view that satisfaction of a menu condition is sufficient for a group to have power. A majoritarian mechanism could be in place, but strategic considerations could give rise to collective action problems so that some majority's preferences fail to constrain the collective decision taken, irrespective of how much importance the majority attaches to the decision. We would then hesitate to describe the decision-making procedure as a perfect case of popular control.

Riker (1982) claims at one point that in an electoral democracy, voters would enjoy a "negative kind of control" (p. 243), notwithstanding all of the social choice paradoxes that, he argues, discredit the populist's theory of democracy. An electoral democracy does not establish "popular rule," but rather a "random, even perverse popular veto," and this limited accom-

to assert only a contingent fact, rather than a conceptual truth, when they write of a particular procedure as a form of popular control.

plishment would “sometimes” suffice to “restrain official tyranny” (p. 244). Riker’s critical project undermines the foundation on which one might develop this idea into a defensible account of popular control. As Cohen (1986) observes, if “the outcomes of elections are genuinely ‘random,’ then electoral ‘discipline’—the possibility of throwing the bums out—does not by itself provide rulers with any reason for not being bums since they would be equally likely to be thrown out if they were not bums.”⁶ The conventional explanation for why elections might “restrain official tyranny” would be that tyrannizing over the population is a good way to persuade most of the electorate that they would be better off with someone else in power. But claims that election outcomes reflect what “most voters” want are just the kind of claims that social choice theory exposes as problematic according to Riker.

Perhaps Riker’s idea is instead something like a proceduralist definition of popular control. The claim is perhaps just that citizens have actions by which they can remove elected leaders from office, and the claim that they have a “negative kind of control,” and are capable of sometimes “restraining” officeholders, means nothing more than this. But that will not make for a plausible account of popular control. We can imagine a dysfunctional “democracy” in which elections are free and fair, but in which the people elected to power always use state power to tyrannize over the population. That scenario is implausible under standard assumptions about democracy, but implausible or not, it would be perverse to describe the hypothetical

⁶Cohen (1986, p. 30). See also Mackie (2003, p. 35)

scenario as an instance of popular control—even a merely “negative kind of control” or “popular veto.” An account of popular control will only respect our intuitions about control if it includes something like the preference condition, which describes how outcomes depend on individuals’ preferences. To use game theorists’ jargon, the necessary conditions for popular control must do more than impose restrictions on the “game form” of democratic politics.

In summary, the reasons to adopt the multiple majorities view of popular control are just the reasons to accept the analysis of the concept of control, from chapter 4, together with the reasons to judge, first, that popular control is just equivalent to the right groups’ having control of the right kind and amount, and, second, that a distribution of control is only democratic and popular in character if one group’s having control over a variable implies that every other equally large or larger group also have control over it.

§5.3 Popular control and voting rules

Let me draw out further important implications of the fact that the multiple majorities view is not a proceduralist account of popular control. This fact explains why one can accept the view without taking a stand on the relative merits of different voting rules.

The multiple majorities view is not a proceduralist account of popular control because the preference condition must be satisfied for a majority to have control. Whether the condition is satisfied depends on how people are disposed to act, not just on the formal properties of the menus of actions

available to each individual. In the jargon of game theory, whether the preference condition is satisfied turns on features of the equilibrium to a decision-making game, not merely on formal properties of the game form.

To illustrate the point, let me briefly comment on how one would evaluate voting rules if one wished to choose a voting rule that would give all majorities control. Whether a voting rule gives all majorities control depends not merely on its formal properties, but also on how voters' preferences over the outcome, and the importance that voters attach to the outcome, affect their behavior. Suppose, for example, that plurality rule is used to select a candidate from a field of three or more contenders. We cannot deduce from the formal properties of the voting rule alone whether majorities have control over the outcome of the election. The voting rule fixes part of the menu of actions available to each voter (but only part, because there are still the further questions regarding whether voting is mandatory and whether and how voters can take actions to persuade others and otherwise influence the outcome). Even if a voting rule provides the members of a majority with actions by which they can determine the outcome of the vote, they fail to have control over the outcome if the preference condition is violated: if, given how they and others are disposed to behave, the outcome will frustrate their shared preference for an alternative outcome, no matter how important they regard the outcome of the vote. To determine whether the preference condition is satisfied for every majority, one would need a game-theoretic model of the election—or something that addresses the same questions as a game-theoretic model. One needs to consider what happens “in equilibrium,” as the joint effect of the voting rule and its formal properties, on the

one hand, and individual behavior, on the other.

A corollary is that “majority rule” is not guaranteed to be the voting rule that most reliably gives all majorities control over a decision. Whether a majoritarian voting rule or some non-majoritarian voting rule like the Borda count, range voting, or something else does a better job of giving all majorities control is an open empirical question. The answer depends on facts about equilibrium behavior, which cannot be deduced from the properties of the voting rule.

One might wonder whether voting rules like plurality rule or the Borda count are even capable of giving all majorities control over choices between three or more alternatives, for such choices must be sensitive to the relative strength of majorities’ preferences if all majorities are to have control at the same time. Yet voting rules like plurality rule or the Borda count are “ordinalist,” by which I mean that they solicit only individuals’ ordinal preference rankings over possible outcomes, not any information about the intensity of these preferences relative to their other concerns.

To describe them as ordinalist is merely to describe their formal properties, however. Whether they give majorities control depends on how voters and other actors are likely to behave once the voting rules are in place. It is possible that when an ordinalist voting rule is used to decide an election or some other issue, the importance that voters attach to the outcome of the vote, relative to other concerns, affects the outcome. It may affect the effort people devote towards persuading other voters of their views. Or perhaps the importance that some voters are seen to attach to the outcome affects the beliefs that other voters form about their voting intentions—perhaps

the more one cares about the outcome relative to maintaining one's own ideological purity, the more likely one is to vote strategically. Reflecting on such possibilities, we can see why even "ordinalist" voting rules could induce a causal relationship between the strength of voters' preferences—as measured by willingness to expend effort and time, or willingness to sacrifice ideological purity—and the outcome of the vote. This is of course just speculation; its point is to clarify the kinds of questions one would need to investigate in order to answer questions about which voting rules establish popular control.

One could design a voting rule specifically for the purpose of eliciting information about how much voters value their preferred outcome relative to other things they care about. Posner and Weyl (2015, 2018) propose "quadratic voting," where voters purchase votes for or against a measure, at a price equal to the square of total number of votes purchased. The measure is approved if there are more votes cast for it than against it, and the proceeds are then redistributed equally among voters. The pricing scheme is designed to ensure that a rational voter's purchase reveals how much she values the success (or defeat) of the measure, relative to the value of a dollar. If each voter's total number of purchased votes is proportional to the amount of money she would be willing to pay to achieve her preferred decision, then the measure will be approved if and only if it is efficient, in the sense that total benefits exceed its total costs, measured in voters' willingness to pay.⁷

⁷In their more recent book, Posner and Weyl propose a similar voting system in which citizens have annual budgets of "voice credits" that they can either use to purchase votes – at quadratic prices – in referenda during the current year or save for future use (Posner

It might seem as though a voting rule such as quadratic voting is the best means of giving majorities control, because the rule makes the decision sensitive to the strength of voters' preferences. But that is not right, at least not in the case of a binary decision. If we think that simple majority voting over a binary decision makes the outcome *insensitive* to the strength of voters' preferences, then we would expect quadratic voting to give majorities *less* control over binary decisions than simple majority voting. Recall that by our definition, all else being equal, a majority has more control over a decision, the lower the threshold of intensity at which its preferences must be held in order to have constraining effects on the decision. If simple majority voting makes the outcome insensitive to the strength of voters' preferences such that, in equilibrium, a proposed decision is taken if and only if a majority prefer for it to be taken, then majorities have maximal control. The threshold of relative intensity at which their preferences constrain the outcome of the vote is as low as possible. Under quadratic voting, by contrast, the shared preferences of a majority constrain the decision only if the members of the majority are willing to pay more, in the aggregate, for their preferred outcome than the members of the corresponding minority are willing to pay for the other outcome. The majority will still have some control—if each member of the majority cared enough about the outcome, she would be willing to purchase enough votes to ensure that the majority's preferred decision prevails. But the threshold of intensity at which the majority's shared preference would constrain the decision would be higher than

and Weyl, 2018).

if simple majority voting were used.⁸

This is not a criticism of quadratic voting, by the way. I am merely using quadratic voting to illustrate a broader point. In the case of binary decisions, voting rules that ensure that majorities' preferences prevail irrespective of their intensity give majorities maximal control. Whether that is desirable is another question.

§5.4 Popular control and democracy

Popular control, as I am characterizing it, may strike the reader as an underwhelming achievement. It may seem it is achieved too easily. After all, a majority can have control over policy even if its members' shared preferences constrain policy only when they are prepared to give up everything in order to satisfy their policy preferences. Indeed, the definition allows us to say that even a group of disenfranchised citizens, or the subjects of an autocrat, can have control over policy: it could be true that if they cared so much about policy that they were willing to risk their lives in mass protest or revolt, then policymakers would take their shared preferences into account and choose policies accordingly. This feature of the definition may appear to weaken the plausibility of the resulting view of popular control. We are

⁸Note that the argument in the text concerns majorities' control over a binary decision. When there are more than two possible outcomes and no Condorcet winner exists, it is impossible for every majority to have *maximal* control because there is guaranteed to be a majority whose preferences do not constrain the decision.

trying to characterize popular, *democratic* control, the kind of popular control we associate with a *democracy*, and democracy of course requires more than this.

The implication that citizens can have *some* control over policy and the behavior of their leaders even in undemocratic regimes or dysfunctional democracies is not at all implausible. Even in an authoritarian regime, the threat of collective action can induce government officials to be responsive to citizens' preferences. Chen, Pan and Xu (2016) report the results of a field experiment in which Chinese government officials were significantly more likely to respond to a (fictitious) request for help with obtaining government benefits when, as part of the randomized treatment, the researchers included an implicit threat of collective action as part of the fictitious request. Scholars of authoritarian regimes regularly use terms like *responsiveness* and *accountability* to describe such dynamics, while acknowledging, of course, how they differ from their democratic counterparts. It is not a stretch to say that such regimes exhibit *some* measure of popular control over government officials if, through threats of collective action or other means, every majority of the population can induce officials to heed their preferences when they feel sufficiently strongly about an issue. I am not hereby endorsing this empirical claim about authoritarian regimes, although it seems plausible. The point is simply that if this claim follows from the analysis of control, together with the empirical evidence, the implication is not absurd. It is no reason to reject the conceptual analysis or the resulting account of popular control.

What would be objectionable would be a conception of popular control

that provided no resources for distinguishing the degree of popular control one might plausibly find in an authoritarian regime from the degree of popular control that a well-functioning democracy is expected to exhibit. But that objection cannot be lodged against the multiple majorities view. It identifies two dimensions along which the degree of a group's control varies: the range of possible realizations of a variable over which the group's preferences can have constraining effects, and the minimum importance the group's members have to attach to a variable in order for their preferences to have these constraining effects. One would expect that by this measure, majorities would have less control over their leaders in authoritarian regimes than they have in well-functioning democracies. In authoritarian regimes, citizens have means of influencing government officials, but taking advantage of these opportunities exposes them to risks of political reprisal. Compared to their counterparts in democratic regimes, who can exercise political influence without the same costs, citizens of authoritarian regimes must therefore attach more importance to the decisions they are trying to influence in order for it to be worthwhile for them to use their opportunities for influence. Thus, citizens can be expected to enjoy less control in authoritarian regimes than democratic regimes.

Control comes in degrees, but we can draw a categorical distinction that will plausibly sort authoritarian and democratic regimes into two separate groups. Let us say that a majority's control over a variable is *adequate* if they do not need to be willing to jeopardize their physical safety or personal liberties, as the price for securing influence, in order for their shared preferences to have constraining effects. The idea here is to characterize the

amount of control that is adequate for a “minimal democracy,” a regime in which the degree of popular control is just enough that we are willing to entertain the proposition that the regime counts as a democracy; any less, and it would not make sense to describe it as a democracy at all. Of course, this is not the amount of control that we expect a “fully democratic” regime to establish. Democracy comes in degrees (Arneson, 2009). The thought is that popular control also comes in degrees, and, all else being equal, a regime is more or less democratic, depending on the degree of popular control it tends to establish, yet there is some minimum amount of popular control that must be present to qualify as a democracy at all. The characterization of the threshold for adequacy is admittedly vague and in need of refinement, but it should suffice to indicate how one could draw a distinction between authoritarian and minimally democratic regimes according to the degree of popular control they establish.

The claim that citizens of democracies have greater control over their leaders than citizens of authoritarian regimes is arguably not an empirically testable hypothesis, but rather a logical corollary of how the concepts of democratic and authoritarian regimes should be defined. Part of what makes a regime democratic is that it gives citizens an adequate degree of control over their leaders; that is not an empirically testable claim but a truth about the concept of democracy. To say that democracies tend to give their citizens greater control over their leaders than authoritarian regimes is like saying that even numbers tend to be divisible by two. What is empirically testable is the claim that citizens in those regimes with competitive multiparty elections, universal suffrage, free media, and the rule of law tend

to have greater control over their leaders than citizens in regimes lacking these features. If the claim is born out by the evidence, what we learn is that the regimes in the first category are more democratic than the regimes in the second, at least along the dimension of popular control.

Of course democracy is a contested concept, so it is unreasonable to refer to *the* definition of the concept, as if there were only one reasonable candidate. The claim I wish to make is rather that a concept of democracy that identifies popular control, as I have characterized it, as one of the criteria for judging a regime to be democratic, and identifies degrees of popular control as one of the dimensions along which regimes can be more or less democratic, will best balance two competing desiderata. In general, it is desirable if the definition of a concept puts it roughly in line with common intuitions about what the concept means (List and Valentini, 2016). For the concepts used in social science and political theory, however, fidelity to common intuitions is insufficient and potentially conflicts with a second, more important desideratum, namely that the concept be useful for the explanatory and evaluative purposes of these disciplines.

A common intuition about democracy is that democracy institutes popular rule, as the ancient Greek name promises. But since the actually existing regimes that are called democracies today appear radically different from the regimes originally designated as instances of *demokratia* in ancient Greece, and, to many observers, do not seem to establish anything resembling popular rule, we appear forced to choose between accommodating the intuition about popular rule and denying that any true democracy exists today, or doing violence to the intuition and “misdescribing” modern states

as democracies (Dunn, 2005). The worry is that trying to accommodate the intuition would make the concept of democracy useless for social science as well as a normative democratic theory, at least any normative democratic theory that is not utopian. To preserve the concept's utility, scholars of a "realist" bent define it minimally as a regime with competitive elections, the rule of law, free media, etc. The canonical example of this move is of course Schumpeter's (1942) minimalist definition of democracy, motivated ostensibly by the desire to have a concept that applies to actually existing regimes.⁹

If we say that a regime is more or less democratic, depending on the degree to which it brings political decisions, or the public officials who make decisions, under majorities' control, then we will preserve a connection between the concept of democracy and the intuition about popular rule. Democracy will not be popular rule in the sense of a collective agent, the People, issuing explicit commands and proclaiming laws. But democracy will institute a kind of popular rule—rule by majorities—in the same sense that an oligarchy institutes a kind of rule by the wealthy: in each case, there is no unified group agent, articulating commands and imposing an imaginary collective will, but instead different groups of individuals (different majorities, different groups of wealthy citizens) whose shared preferences, held with enough strength, constrain a range of political decisions, or constrain a range of behaviors of the public officials charged with making those decisions. To make this into a plausible concept of democracy, we will have

⁹For an alternative account of Schumpeter's motives, see Medearis (2001).

to add further criteria—the extent of equality in the distribution of control across different majorities; the scope of majorities’ control; its robustness to the whims of political elites (Pettit, 2012); not to mention the rule of law, protections for basic rights, and so on. (I discuss some of these requirements in chapter 7.) The claim here is not that whether majorities have control is the only relevant criterion when judging a regime’s democratic credentials, but rather that it is an important one.

This characterization of democracy is deliberately flexible with respect to the *object* of popular control in democracies—political decisions or the public officials charged with making those decisions. Majorities could have control over policy decisions if these are made through popular referendums and ballot initiatives, but, at least in theory, they could also have control over policy decisions in a representative democracy where elected public officials choose policies. I discuss an example later in the chapter, using a model of electoral competition. The possibility is analogous to the example, from the previous chapter, of the newspaper owner who has control over the editor’s choice of content. Alternatively, a representative democracy may not give majorities control over policy decisions, but may still give them control over the public officials who are authorized to choose policy. One possibility, which I discuss in the next chapter, is that majorities have control over the kinds of outcomes elected leaders aim to bring about with the policies they choose. This idea is similar to Christiano’s (1996) “choice of aims” model, in which democratic citizens are sovereign because they choose the aims that will guide the policymakers acting as agents on their behalf (Christiano, 1996, especially pp. 169–172). Either view of the object

of popular control in a democracy accommodates, albeit to different degrees, the common intuition that democracy is a system of popular rule.

A concept of democracy along these lines will fare well by the second desideratum, too. Defining it in terms of popular control will not render the concept of democracy useless as a tool for understanding, explaining, and evaluating modern political regimes. In those modern regimes normally called democracies, which make winning elections a condition for attaining and holding most of the highest political offices, there are plausible theoretical reasons, and some empirical evidence, supporting the view that they reliably produce situations in which majorities have control over elected officials. In the next chapter, I explain why retrospective voting is a plausible mechanism of popular control, even if one grants the kind of claims that Schumpeter's intellectual descendants offer as reasons for doubt. In chapter 7, I will explain why the concept of popular control is fit to play a role in a normative theory of democracy, one that identifies democracy's tendency to give majorities control over elected officials as a core element of its justification.

§5.5 Why the impossibility objection fails

The impossibility objection was the claim that in many circumstances, policy will necessarily fail to be under popular control—no matter the institutions in place and no matter how citizens behave—because Condorcet winners will often fail to exist. The key premise underlying the objection was its assumption about the concept of popular control, namely that if policy is

under popular control, then, by logical implication, policy respects majorities' preferences.

We now have a good reason to reject this assumption. The reason derives from a general analysis of the concept of control. The reason to reject the conceptual claim is not merely that it leads to a conclusion about democratic control that we find unpalatable. According to the definition of control from the previous chapter, a majority may have control over a variable even if its actual value frustrates the majority's preference for an alternative value. Having *some* control merely implies that, were the majority to care sufficiently strongly about the variable, then its shared preferences would constrain the variable's realized value. That no Condorcet winner exists among a variable's possible values—that it is bound to frustrate the preferences of a majority no matter what value it takes on—is therefore no reason to conclude that majorities lack control. Its actual value will frustrate a majority's preference for an alternative outcome, but the frustrated majority may still have control over it so long as, were the members of the majority to attach sufficient importance to the variable, then their shared preferences would constrain the variable, preventing the less preferred outcome from occurring. This claim follows from a general analysis of the concept of control, one justified by considerations independent of the ramifications in democratic theory, so it avoids the charge of being *ad hoc*. It gives us justified grounds for rejecting the first premise of the impossibility objection—that if a variable is under popular control, then it does not frustrate majorities' preferences.

I have referred to the impossibility objection as an “objection from social

choice theory,” because the premise that Condorcet winners frequently fail to exist gains plausibility from results in social choice theory, especially those on the multidimensional spatial model. However, it is important to emphasize that the objection differs from the results in social choice theory on which it is based, and it differs from other objections people have made on the basis of those results. In claiming that the impossibility objection fails, I am criticizing one particular inference from those results, not casting doubt on their importance or their implications for other questions.

In fact, I endorse some of the other skeptical conclusions scholars have drawn from them. I have not disputed Riker’s (1982) claim that Arrow’s theorem is a reason to reject the notion of a popular will.¹⁰ I claim it is a virtue of the multiple majorities view of popular control that it does not require us to invoke notions like “the will of the people.” It is consistent with “Pareto extremism,” the view that one can only say a group prefers one option to another if every member of the group prefers the one to the other.

Social choice theory gives us critical purchase on fundamental concepts in democratic theory, like popular rule and popular control, but it does not force us to the cynical conclusion that they are meaningless. In this respect, my argument about the implications of social choice theory for

¹⁰In fact, elsewhere I have argued, in the same spirit as Riker, that Arrow’s theorem undermines the plausibility of a populist theory of democratic legitimacy according to which decisions are legitimate only if they reflect the will of the people (Ingham, 2018).

the concept of democratic, popular control is analogous to the argument about its implications for the concept of political legitimacy in Patty and Penn (2014). Social choice theory is relevant to social science and political theory because it reveals logical constraints on how we understand some of the core concepts of the disciplines, concepts like popular control (my focus) or political legitimacy (theirs). It discredits some interpretations of these concepts—that popular rule or popular control requires satisfying the will of the people or the will of the majority, or that a collective choice is legitimate only if it is “best” according to the relevant criteria. But other interpretations remain; we need not conclude that popular control and legitimacy are meaningless or “impossible.”

So far, my claim has only been that it is logically possible for all majorities to have control in situations where no Condorcet winner exists. That no Condorcet winner exists among the feasible policies—that the actual policy will inevitably frustrate a majority’s preference for an alternative policy—is not, on its own, a reason to conclude that a majority lacks control over policy. Still, that claim leaves open the possibility that whenever a Condorcet winner fails to exist, one or more majorities lack control—not as a logical consequence but rather as a “strategic” consequence, that is, as a consequence of how decisions get made when no Condorcet winner exists. To address that worry, I now consider a mechanism by which democratic elections might give all majorities control even when no Condorcet winner exists. Later in the chapter, I return to the status of the impossibility objection and anticipate criticisms of my claim to have rebutted it.

§5.6 A model illustrating rule by multiple majorities

I take up two outstanding questions about the conception of popular control in this section. First, what does it *mean* for all majorities to have control over something at the same time? The terms have all been defined, so, at one level, an answer has already been given. But one is unlikely to fully understand what it means for all majorities to have control until one has a concrete model that illustrates the idea. Second, by what kinds of mechanisms could all majorities have control over, say, policy decisions in a democracy if no Condorcet winner exists among the policies, and majorities' preferences form cycles?

I answer these questions using a slightly adapted version of the models of electoral competition with probabilistic voting in Lindbeck and Weibull (1987), Dixit and Londregan (1996), and Persson and Tabellini (2000).¹¹ The model describes a hypothetical democracy in which majorities have control over policy decisions, even though no Condorcet winner exists among the feasible policies. The hypothetical scenario is unrealistic in various ways, but the mechanism by which majorities could come to have control in this democracy is not far-fetched. The model is a useful starting point for thinking about whether majorities might have control over policy in actual democracies.

To avoid confusion, let me emphasize that one can accept the multiple

¹¹I draw on the exposition of the model in Gehlbach (2013, pp. 28–31), where the reader can find discussion of more technical points as well as further literature references.

majorities view of popular control even if one does not like this particular model. The concept of popular control is not itself a model of democracy. It is instead an account of the criteria one should use to judge whether a democracy establishes popular control. In order to illustrate in concrete terms what those criteria mean and how they are to be applied, we need a model, and models of electoral competition with probabilistic voting serve this purpose well whatever their other shortcomings.

Imagine two political parties, L and R , that compete in elections by taking positions on various policy issues. I will refer to a set of issue positions as a *platform*, for short. The winning party can be expected to implement its platform, and voters have preferences over the possible platforms. Assume that their preferences over platforms are such that no Condorcet winner exists: no matter which platform a party has adopted, a majority of voters prefer a different platform to it.

Each voter votes for the party she prefers, and her preferences over the parties reflect two components: her preferences over their platforms together with idiosyncratic sources of affinity for one or the other party. Specifically, a voter i , in group g , is assumed to prefer party L , and to vote for L , if

$$\omega_g v_g(x_L) > \omega_g v_g(x_R) + \eta_i,$$

where x_L and x_R refer to the two parties' policy platforms, the function v_g represents the preferences over policy platforms of voters in group g ,¹² the

¹²That is, voters in group g prefer the platform x_L to x_R if $v_g(x_L) > v_g(x_R)$.

term η_i is voter i 's idiosyncratic, partisan affinity for party R (or, if $\eta_i < 0$, her affinity for L), and $\omega_g > 0$ represents the relative importance of policy for voters in group g .

The parties are assumed to know voters' preferences over policy but not their partisan affinities: the terms η_i appear to them as random variables, whose values are unknown and outside of their control. From a party's point of view, the probability that voter i votes for L , when the parties have adopted the platforms x_L and x_R , is the probability that the random variable η_i takes on a value less than $\omega_g[v_g(x_L) - v_g(x_R)]$. The parties choose their platforms with the goal of maximizing their expected vote shares.

In comparing two possible platforms, the party considers not only voters' preferences over the two platforms but also the *strength* with which voters hold their policy preferences. A voter in group g is more likely to vote for party L , the greater the strength of her preference for L 's platform, as measured by $\omega_g[v_g(x_L) - v_g(x_R)]$. As it increases in magnitude, it becomes less probable that the magnitude of an idiosyncratic partisan affinity for party R will outweigh it. As a result, a party does not necessarily wish to choose a platform that a majority of voters prefer to the opposing party's platform. It may be that the members of this majority do not care very much about policy relative to their partisan attachments, so that their voting decisions would likely be determined by partisan factors anyway, whereas the minority cares a lot about policy and would almost surely vote on the basis of its policy preferences.

In the equilibrium to the model, both parties converge to the policy platform preferred by the "mean" (not the median) voter, that is, the platform

that maximizes the average of voters' utilities. For example, if each group g makes up a fraction α_g of the electorate, then both parties choose the platform that maximizes $\sum_{g=1}^n \alpha_g \omega_g v_g(x)$. In general, the platform to which the parties converge would not be a Condorcet winner even if one existed. For the reasons explained, the mere fact that a majority of voters prefer a different platform is insufficient incentive for the parties to deviate from the equilibrium and adopt it.¹³

In this model, do majorities have control over policy, according to the given definition? A majority has control over the policy platform of the winning party if there is a range of pairs of possible platforms such that for any pair of platforms x and y in the range, if all members of the majority preferred x to y , then y would not be the platform of the winning party, provided each member of the majority cared enough about policy relative to other variables they care about. In this context, the other variable that each voter cares about is the identity of the winning party: voter i gets a payoff η_i when R wins, reflecting the idiosyncratic reasons, unrelated to

¹³I have left unstated various technical assumptions on which this conclusion depends. In particular, it depends on the assumption that the idiosyncratic affinities are all drawn from the same probability distribution. Otherwise the conclusion is that the parties choose a platform that maximizes a *weighted* average, where the weights reflect differences in how likely different groups of voters are to cast their ballots on the basis of their policy preferences, rather than their idiosyncratic affinities; see Gehlbach (2013, pp. 28–31).

policy, why she might care whether L or R wins the election, and the term ω_g represents the relative importance of her policy preferences. To simplify things, suppose there are just three equally sized groups (so that the size parameters α_g can be dropped), and pick the majority comprising groups 1 and 2. We will show that the preference condition holds for this majority.

Hold fixed the preferences of voters in group 3, the minority. Let \hat{v}_3 be group 3's actual policy preferences and $\hat{\omega}_3$ the weight they actually give to policy. (Recall the discussion, in chapter 4, about the need to hold fixed other agents' preferences and strategies when evaluating propositions about a group's control.) Take any two feasible platforms, x and y , and suppose groups 1 and 2 were to prefer x to y , meaning $v_g(x) > v_g(y)$ for $g = 1, 2$. Provided both groups attach enough weight to policy, the plan y will fail to maximize the average voter's utility: that is to say, one can choose ω_1 and ω_2 large enough so that

$$\omega_1 v_1(x) + \omega_2 v_2(x) + \hat{\omega}_3 v_3(x) > \omega_1 v_1(y) + \omega_2 v_2(y) + \hat{\omega}_3 v_3(y)$$

In any of these counterfactual scenarios, in equilibrium the parties would choose the platform that maximizes the average voter's utility. Thus, we conclude that for any platforms x and y , if groups 1 and 2 both strictly preferred x to y and attached enough weight to policy, then neither party would choose y in equilibrium. Assuming the parties implement the policies they propose during the election, the conclusion is that the majority comprising groups 1 and 2 has control over policy.

The same argument can be made about other majorities. When evaluat-

ing the preference condition for any majority, one holds fixed the preferences of the corresponding minority and asks what would happen if, counterfactually, the members of the majority cared even more about policy than they actually do. The condition can be satisfied for every majority at the same time because, in each case, it is evaluated while holding fixed the preferences of the corresponding minority and the strategies of the parties. Every majority can therefore have control at the same time.

The model is not meant to describe actual democracies, although the mechanism it describes is as plausible as any other mechanism capable of explaining how *policy* could come to be under popular control in an electoral democracy. The primary value of the model is to illustrate what it means to judge that all majorities have simultaneous control, and how one would go about checking the conditions for their control. I consider in the next chapter a more plausible mechanism by which actual democracies might establish some degree of popular control, not over policy but over elected leaders and their aims.

§5.7 Objections

Some readers may not be persuaded that I have rebutted the impossibility objection from social choice theory, as laid out in Chapter 3. In this section I address several possible reasons for skepticism.

The impossibility objection is an objection “from” social choice theory in the sense that one of its critical premises—that Condorcet winners regularly fail to exist—derives plausibility from results in social choice theory. But the

problem it raises for a theory of popular control differs from other problems social choice theory has been thought to raise for democratic theory. If one conflates the distinct challenges that social choice theory poses to democratic theory, one is unlikely to understand the book's contribution.

As I emphasized throughout chapter 3, the impossibility objection is not the claim that democratic processes will be marked by “disequilibrium” and fail to generate stable decisions. Proving the stability of democratic decisions would not refute the impossibility objection, and one can refute it without proving the stability of democratic decisions.

I emphasize those points in order to prevent misunderstanding of the impossibility objection and the sense in which I have met the objection, not to belittle the concern with stability. If the internal logic of democratic decision-making ensures that it could never produce stable decisions, that would be an important discovery. And, if it were true, it might diminish the value of democratic decision-making, although some political theorists have tried to give the instability interpretation of majority cycles a positive spin (Pildes and Anderson, 1990; Shapiro, 2003).

However, I see no reason to believe the instability thesis. What we learn from the theorems on the multidimensional spatial model—inappropriately and misleadingly referred to as “chaos” theorems—is that the stability of democratic decisions cannot generally be explained by the fact that no majority prefers an alternative decision to the one taken. In general, there is little reason to accept the premise of that explanation—that Condorcet winners exist, that no majority prefers an alternative to the outcome of a decision-making process—so it is not a convincing explanation of anything.

The stability of democratic decisions must instead be explained otherwise, but there are plenty of plausible explanations. Canonical examples of such explanations include the models of “structure-induced equilibrium” (Shepsle, 1979; Shepsle and Weingast, 1981): models in which institutional structures constrain majoritarian decision-making and thereby induce a stable equilibrium. Another example is the model of competitive elections with probabilistic voting, considered above. This model, as well as one of the models of electoral accountability through retrospective voting discussed in the next chapter, describe hypothetical democracies in which the democratic process produces a stable outcome even though no Condorcet winner exists.

One worry that democrats might have with some of these explanations is that they seem to explain the possibility of stable democratic decisions only by positing limits to what majorities are able to do. For example, in the early models of democratic legislatures in this vein, stability arose from institutional rules that forced legislatures to choose “issue by issue.” Stability arose only from curbs on popular democratic control, apparently.

The concept of rule by multiple majorities offers some guidance on how to think about this issue. Consider again the model of competitive elections with probabilistic voting. In the model, electoral competition has a “stable” outcome, in the sense that neither party has any incentive to act differently, given accurate beliefs about the platform the other party is adopting. In a sense, stability does come at the expense of majorities’ control: no majority has *maximal* control, owing to the same assumptions that produce stability. Some majority will prefer a different platform to the one that the parties adopt in the equilibrium, and to explain why neither party has an incentive

to deviate, we have made assumptions about their imperfect knowledge of voting behavior. These assumptions imply that the majority's preference for alternative policy platforms will constrain the parties, preventing them from choosing the less preferred platform they adopt in the actual equilibrium, only if the members of the majority care enough, and care more than they actually do in the equilibrium. That means the majority has less control than it theoretically could have. The majority would have more control if the parties were certain that its members would vote solely on the basis of their known policy preferences. This conclusion is no reason to despair of the level of popular control that the model describes, however. Not every majority can have *maximal* control at the same time. But control comes in degrees, and there is nothing in the model to justify the view that majorities have less than adequate control.

These comments are, I hope, enough to indicate how one would approach the traditional concerns about democratic stability from the perspective of the multiple majorities view of popular control. Let me now consider what many people would think of as “the” problem raised by social choice theory, namely, the aggregation problem that Arrow addressed with his impossibility theorem.

The formal aggregation problem that Arrow's framework describes can represent various substantive problems, and the substantive implications of the theorem will depend on the substantive interpretation one gives the framework (Ingham, 2018). For example, the preference aggregation rules the framework describes could model voting rules that take individuals' reported preferences as their input and produce collective rankings of candi-

dates or decisions as their output. Alternatively, they could model the manner in which judgments about social welfare depend on information about individuals' welfare. Neither of the substantive problems corresponding to those two interpretations—the problem of finding a reasonable democratic voting rule, and the problem of finding a defensible interpretation of social welfare—are the same as our problem, as discussed in chapter 3. Our problem was to find a defensible interpretation of the concept of popular control. As I emphasized in chapter 3, the problem is not solved by finding a reasonable democratic voting rule, because it is an open question whether reasonable democratic voting rules establish anything deserving to be called popular control. Still, one might ask whether Arrow's theorem has implications for our problem. It certainly has much broader implications than either of its two conventional interpretations would suggest (Patty and Penn, 2014, 2015; Ingham, 2018).

Our approach to the problem was to find a generally plausible concept of control, one that was applicable outside political contexts, and then use it to formulate a concept of popular, democratic control. One implication of that exercise is that the information one needs in order to assess a group's control cannot be represented by a preference aggregation rule of the kind Arrow studied. In the case of a group's control over a collective decision, one needs to know how the decision depends, not merely on citizens' preferences but also on the weight they attach to the outcome of the decision-making process, relative to their other concerns. If one were to formulate the preference condition for control as a constraint on a kind of "aggregation function," the function would not be a preference aggregation rule of the kind Arrow

described, but rather a function that takes as its arguments both the group members' preferences and measures of the importance they attach to the variable, relative to their other concerns. The preference condition would then specify not a particular aggregation function of this kind, but rather a constraint on the aggregation functions. What justifies the departure from an Arrovian framework is the conceptual analysis of control from chapter 4. Its independent plausibility—independent of its ramifications for a theory of popular control—is the reason the departure is not an ad hoc response to the conceptual difficulties raised by social choice theory.

These remarks might make one wonder whether the strategy for rebutting the impossibility objection relies on interpersonal comparisons of preference intensity. If so, one might say the strategy is unoriginal. When Arrow's theorem is viewed as a result about the possibility of social welfare functions, a standard response is to introduce a richer "informational base" for the social welfare function, one that permits interpersonal comparisons of utility (Sen, 1977a). From that point of view, the strategy for rebutting the impossibility objection may appear not only unoriginal, but vulnerable to the objection that interpersonal comparisons of utility are meaningless, unjustified, or merely inappropriate in this context.

The objection rests on a misunderstanding, however. The judgment that a group has control does not require interpersonal comparisons of preference intensity. They require only judgments about the effects of increasing the importance that each member of a majority attaches to a variable. Those are *intrapersonal* comparisons, which concern the effects of variation in a person's "type," specifically variation in the tradeoffs she would be willing

to make between the variables she cares about. For example, we ask whether the shared preferences of a group would constrain a variable, if each member were willing to pay more money, or exert more effort, or expose herself to a greater risk of imprisonment, or less averse to voting for members of the other party (as in the probabilistic voting model)—more generally, if each member felt more strongly about the variable in question relative to her other concerns. These are judgments with observable implications, as explained in section 4.2.

In chapter 7, I discuss interpersonal comparisons of the practical significance of different groups' control, comparisons that have a normative component and go beyond the intrapersonal comparisons described here. Whatever one makes of the arguments in chapter 7, those interpersonal comparisons are necessary if all one wishes to do is establish that all majorities have control. No interpersonal comparisons are unnecessary if one's goal is simply to refute the objection to the possibility of popular control.

Another objection is that I have not shown that popular control is *always* possible, and that in contexts where a collective decision or some other variable is sensitive only to *ordinal* properties of citizens' preferences over the variable, the impossibility objection from social choice theory applies with as much force as ever.

I grant what is alleged but deny that it is any objection. Take a situation in which citizens' preferences determine the realized value of a variable, but the importance they attach to this variable, by any plausible measure, has no effect on it, and suppose that no Condorcet winner exists. Then the shared preferences of some majority are frustrated. Since, by hypothesis,

the importance they attach to the variable is causally inert, their preferences would be frustrated even if they attached more importance to the variable. Thus, they have less than full control over the variable according to my analysis. (They may still have some control, because there may still be a more limited range of possible outcomes over which their preferences have constraining effects.) Does this concession detract from the significance or plausibility of the account of popular control? Not in the least. In particular, if there is no range of possible outcomes over which a majority's preferences constrain a variable, and that is true no matter how strongly the majority feels about the variable, then we should not say they have control. An analysis of control that supported any other judgment about such a case would be implausible. It would not be recognizable as an analysis of the concept *of control* as opposed to some other concept.

Even so, the possibility of such cases may seem embarrassing for my view. Does it not show that the skeptics are right, and that social choice theory undermines the coherence of popular control as a description of (or prescription for) democracies? Consider a revised version of the impossibility objection (treating policy as the variable whose subjection to popular control is in doubt):

1. If policy is under popular control, then each majority's preferences would constrain policy if its members attached enough importance to policy.
2. Whatever institutions are in place, it will frequently be both the case that citizens' preferences over policy have the structure of Condorcet's

paradox and the case that policy is unaffected by the relative strength of citizens' preferences.

3. Thus, whatever institutions are in place, policy will frequently not be under popular control.

The first premise follows from the account of popular control I have defended. The second premise implies, just by the definitions of the terms, that whatever institutions are in place, some majority's preferences will fail to constrain the choice of policy—because no Condorcet winner exists—and, moreover, would fail to do so no matter how strongly they are held—because policy is unaffected by the relative strength of preferences. The conclusion follows.

How should we judge the second, critical premise? I argued in chapter 3 that we have no reason to doubt that citizens' preferences often have the structure of Condorcet's paradox, which leaves only the second half of the premise to consider. Is there any reason to believe that no matter the institutions in place, the relative strength of citizens' preferences will have no effect on policy?

Just the opposite. There is good reason to believe that in actually existing democracies, the relative strength of citizens' preferences *does* affect policy and other variables of interest. It is one thing for a legislator to know that a life-long Democratic voter prefers pro-choice policies; it is another to know that her preference is so strong that she would vote for a pro-choice Republican over a pro-life Democrat. Kingdon (1989) provides qualitative evidence from interviews with members of Congress that their voting is in-

fluenced by the perceived intensity of constituents' preferences.¹⁴

Said one liberal congressman about the House Un-American Activities Committee vote, "The bulk of my constituents oppose me on this, but they don't care enough about it to make a difference." A Republican who voted for an increase in the debt limit described his constituents thus: "This isn't a big issue with them. They probably oppose raising it, but it's not crucial." (Kingdon, 1989, p. 35)

The model from the previous section describes one mechanism by which the relative strength of policy preferences could influence parties' (or candidates') policy platforms. Voters may have reasons unrelated to policy for preferring one party to another, reasons which they weigh against their policy preferences when choosing whom to vote for. Knowing this fact about voters, parties will give more weight to the policy preferences of those voters who attach greater weight, relatively speaking, to their policy preferences. While the model was unrealistic in other respects, it seems to be on firm ground when it assumes that the relative strength of citizens' policy preferences—how much they care about policy relative to partisan attachments or other non-policy-related motivations—affects voting behavior. There is certainly room for debate about how often the relative strength of citizens' preferences is causally efficacious, and whether the pattern of causal effects satisfies the conditions necessary for control, but these brief observations should make the reader skeptical of the revised impossibility

¹⁴I thank Dan Butler for the reference.

objection.

§5.8 Conclusion

Popular control is logically compatible with cyclical majority preferences and the absence of a Condorcet winner. The mere fact that a majority of citizens will prefer an alternative to the chosen policy, whatever the chosen policy is—as must be the case when there is no Condorcet winner among the policies—is compatible with popular control. For this reason, what I called the impossibility objection, from chapter 3, fails as an argument. Its key premise, a conceptual claim about popular control, is one we have good reason to reject. What is necessary for every majority to have control over policy is not that chosen policy satisfy the preferences of every majority, but rather that policy would respect a majority's shared policy preferences if its members all cared enough about policy relative to other variables they care about.

Now that we have formulated a conception of popular, democratic control, which does justice to ordinary intuitions about control but frees the idea of commitments that social choice theory shows to be untenable, we can pose well-defined questions about whether actual democracies establish popular control and when it would be desirable for them to do so. In the remaining chapters, I discuss how one could apply the concept of rule by multiple majorities in social science and normative democratic theory.

Chapter 6

Retrospective voting and control

The probabilistic voting model, used in the previous chapter to illustrate the concept of rule by multiple majorities, rests on debatable assumptions about voters. It assumes voters know the policy platforms that parties campaign on, trust the parties to carry out their promises, and vote “prospectively” on the basis of beliefs about which party will carry out the better package of policies. However, some empirical studies of voting behavior appear to show that the average voter is ill-informed about the parties’ positions and lacks stable, determinate preferences over most of these policy issues.¹

The theory of retrospective voting accommodates these studies. It recognizes that the typical voter is not a political news junkie. It assumes she votes on the basis of retrospective evaluations of how well things have gone during the incumbent’s time in office, assessments which depend on the *results* of the incumbent’s policies, not directly on the policies themselves. As Morris Fiorina put the idea,

¹Achen and Bartels (2016, pp. 30–45) review this literature. For an alternative perspective on the literature, see Gilens (2012, ch. 1).

the retrospective voter need not spend his life watching ‘Meet the Press’ and reading the *New York Times*. He can look at the evening news and observe the coffins being unloaded from Air Force transports, the increasing price of a basket of groceries between this month and last, and the police arresting demonstrators of one stripe or another. (Fiorina, 1981, p. 10)

Retrospective voters need not be aware of the parties’ issue positions, and they need not have stable policy opinions. They only need to take stock of how well things have gone under the incumbent official or party. They vote to re-elect the incumbent if and only if things have gone sufficiently well. The incumbent officeholder, in anticipation, has incentives to take measures to improve voters’ welfare, as they perceive it. In doing so, he need not cater slavishly to their policy preferences (if they even have any). He has incentives instead to exercise independent thought and leadership, identifying the policies that are more likely to produce good results for his constituents and earn him positive retrospective marks at the time of the next election. In this respect, the theory of retrospective voting gives elected leaders a more realistic, and arguably more desirable, job description (Fiorina, 1981, pp. 11, 201).

It has also been thought to describe a plausible mechanism of popular control. If the theory is correct and elected leaders respond to the incentives that retrospective voting creates, then, according to Manin (1997), “through their retrospective judgment, the people enjoy genuinely sovereign power,” even if elected leaders are in charge of choosing policy (p. 183). Its promise

to reconcile the ideal of popular control with the realities of representative democracy accounts for much of the theory's appeal, but it has also met with skepticism. In this chapter, I draw on the ideas from previous chapters to answer two skeptical questions, in particular.

The first is whether the theory is compatible with what we know about actual voting behavior. Achen and Bartels (2016) have given an extensive critique of the theory, arguing that voting behavior is explained largely by social identity and partisan attachments. But "if voting behavior primarily reflects and reinforces voters' social loyalties, it is a mistake to suppose that elections result in popular control of public policy" (p. 4). Moreover, when retrospective voting occurs, they argue it is largely blind, with voters allegedly punishing incumbent officeholders for shark attacks, natural disasters, and other events outside their control.

A second skeptical question is whether retrospective voting can produce popular control even when citizens' preferences do not aggregate into a collective will and different groups of voters assess the incumbent's performance according to different standards. One scholar speculates that if voters' retrospective judgments reflect their own personal, heterogeneous fortunes, "incumbents could pursue a divide and rule strategy by exploiting distributive conflicts in the electorate, and thereby mitigate, or perhaps avoid completely, the discipline of having to satisfy a minimal standard of macroeconomic performance augmenting aggregate welfare."² In a related

²Hibbs (2006, p. 570). Hibbs appears to have in mind here the conclusions of Ferejohn's (1986) model, which I discuss below. Achen and Bartels (2016, p. 99) echo his worry.

vein, one might doubt whether retrospective voting can be described as a form of popular rule if there is no “popular will,” or representative voter, to which elected leaders could be responsive (Green, 2010, p. 106). Most formal models of electoral accountability are not suitable devices for exploring these issues because they posit a representative voter. But in addition to that limitation, previous literature has not determined the implications of retrospective voting with heterogeneous preferences for *popular control*, because it has lacked an account of what the concept of popular control even means in contexts where citizens’ preferences do not aggregate into anything resembling a collective will. Using the analysis of the concept from the previous chapter, we are now in a position to address this source of skepticism.

I argue that both forms of skepticism are misplaced. Against Achen and Bartels and others, I argue that even if voting behavior primarily reflects social identity or partisan attachments, and even if voters irrationally treat random and meaningless events—events that are beyond the incumbent’s control and convey no relevant information—as evidence of the incumbent’s quality and vote accordingly, it does not follow that voters lack control over the incumbent. What matters, for popular control, is whether elected leaders’ incentives are such that a majority’s preferences over policy outcomes constrain the outcomes leaders pursue, provided the majority places enough importance on these outcomes. Whether retrospective voting behavior results from myopic or rational inferences is, in itself, neither here nor there. While the psychological mechanisms that skeptics allege may diminish voter

welfare, they are not good grounds for rejecting the claim that an important and realistic function of elections, and democratic institutions more broadly, is to establish popular control over government officials.

I make this point using an adaptation of a model from the literature on formal models of political agency. The model is used to describe a kind of thought experiment in which we posit the psychological claims made by skeptics and show how popular control over the incumbent emerges through retrospective voting nonetheless. The goal is not to describe actual democracies, but rather to expose a non sequitur in the skeptics' argument. Granting the premises of that argument does not undermine the plausibility of retrospective voting as a mechanism of popular control.

I address the second skeptical question with a further adaptation of the model. The second version of the thought experiment resembles the first, except now different voters care about different kinds of policy outcomes, and majorities' preferences over the incumbent's possible actions, as well as the policy outcomes he influences with his actions, are cyclical. There is no representative voter, no intelligible "will of the majority" to which the incumbent officeholder could be responsive. Nonetheless, the elected official has incentives to perform well along each of the dimensions that (different) voters care about, and the strength of his incentives reflects the amount of importance that voters attach to the policy outcomes relative to their concerns with the officeholder's partisan identity. Majorities end up with control over the elected official and the policy outcomes he targets with his actions, notwithstanding the absence of anything resembling a popular will.

Beyond addressing two sources of skepticism about retrospective voting

as a mechanism of popular control, the arguments in this chapter have a second, no less important function. The best way to understand the meaning of a concept is by learning how it applies to concrete examples. The arguments in this chapter illustrate how the concept of popular control applies in particular models of retrospective voting.

§6.1 Evidence of retrospective voting

I start with a brief review of the empirical literature on retrospective voting and electoral accountability, so as to frame both the skeptical challenges and my responses to them.³

The hypothesis that voters react to measures of elected leaders' performance receives considerable support from the literature on economic voting. If voters treat measures of the economy's performance as noisy signals of the quality of the incumbent public officials presiding over the economy, then one would expect an incumbent party's electoral fortunes to be tied to economic performance on its watch. In their influential comparative study, Duch and Stevenson (2008) conclude that economic voting is "pervasive" (p. 338), confirming Lewis-Beck and Stegmaier's (2007) claim, summarizing previous literature, that the pattern of voters "rewarding" or "punishing" incumbent public officials for economic performance "can be counted on, election after election, country after country" (p. 530).

A more recent literature looks for evidence of retrospective voting using

³I draw here on the helpful and more comprehensive reviews in Ashworth (2012) and Healy and Malhotra (2013).

non-economic measures of performance. Bechtel and Hainmueller (2011) exploit a natural experiment created by the 2002 flooding of the Elbe in Germany to study the effects of the incumbent government's flood relief spending on its electoral fortunes.⁴ The authors find that the massive spending on flood relief in affected districts increased the incumbent party's vote share by 7.1 percentage points on average, an 18% increase. They also find that the electoral return to this spending carried over into the following election, where the incumbent party's vote share increased by about 2 percentage points, on average, in affected districts. That the electoral rewards to beneficial policies can persist for more than one election cycle provides some evidence against strong versions of the 'myopic' voter hypothesis, according to which voters only reward incumbents for beneficial actions taken in the immediate past.

Healy and Malhotra (2009) also find evidence, from the U.S., that voters reward incumbents for spending on natural disaster relief (although not for spending on disaster preparedness). Gasper and Reeves (2011) report evidence, again from the U.S., that voters reward state governors for request-

⁴The plausible assumption permitting them to identify this causal effect is that, in the absence of the flood, the incumbent party's (counterfactual) vote share in electoral districts affected by the flood would have followed a similar trend to the (actual) vote share in electoral districts unaffected by the flood. This assumption of "parallel trends" has *prima facie* plausibility because one would not expect nature to direct the flooding of the Elbe according to trends in districts' electoral outcomes.

ing the president to declare a natural disaster emergency—which clears the way for relief spending—and punish presidents for declining these requests. Other evidence of retrospective voting comes from studies of how war casualties influence incumbent vote share. Using states, media markets, or counties as the unit of analysis, Karol and Miguel (2007) find that the per capita number of Iraq war casualties in the unit had a negative effect on the change in President Bush’s vote share from 2000 to 2004. Similarly, Kriner and Shen (2007) find that for Republican senators, local casualties in the Iraq war were negatively associated with changes in vote share from 2000 to 2006.

Formal models of political agency help us think about the possible mechanisms by which natural disasters, war casualties, and economic conditions during an elected officeholder’s term might affect his future electoral prospects. In the predominant approach to modeling the mechanism, voters are assumed to treat outcomes like economic growth, provision of natural disaster relief, or war casualties as noisy but informative signals of the elected official’s competence or other relevant qualities. Voters replace the incumbent when the measure of past performance is evidence that the incumbent is worse than the typical challenger, producing the observed pattern of retrospective voting.

The rational voter interpretation of retrospective voting has come under greater scrutiny in a more recent body of empirical work. Most notably, Achen and Bartels (2016) concede that retrospective voting happens, but claim that “voters’ retrospections are blind,” in that “incumbents will pay at the polls for bad times, whether or not objective observers can find a

rational basis for blame” (p. 118). Punishing incumbent officials for bad economic conditions when no objective observer could rationally believe the incumbent was responsible for the bad economic conditions is “no more sensible than kicking the dog after a hard day at work” (p. 93). Irrespective of its rationality, “the ‘blindness’ of retrospection poses a significant challenge to democratic accountability” (p. 114).

I argue below that even what is described as “blind” retrospective voting can be a mechanism of popular control, but let me first describe the kind of empirical findings that are taken as evidence of voters’ myopia. A now famous example is Achen and Bartels’s finding that shark attacks in New Jersey coastal towns in the summer of 1916 influenced the U.S. presidential election the following fall. They estimate that the shark attacks depressed President Wilson’s vote share by about 3 percentage points. No reasonable person would treat the shark attacks as evidence of Wilson’s competence or other attributes. Voters appeared to have been “blindly” retrospective, mechanically translating pain and suffering and negative feelings into negative assessments of incumbent performance, even though they had no reason to attribute their negative experiences to the incumbent’s performance. To provide further support for their theory of blind retrospective voting, Achen and Bartels assemble data on droughts and floods in the 48 contiguous U.S. states, from 1897 to 2000, and look at the association between a drought index and support for incumbent party candidates. Across several model specifications and versions of the drought index, they find that extreme droughts or wet spells cost incumbent party candidates significant numbers of votes.

Shark attacks and natural disasters may give elected officials opportunities to demonstrate their ability to deliver aid to their constituents, and the resulting press coverage may reveal elected officials' past negligence in making preparations for these events. If so, then the findings are not evidence against the rational choice interpretation of retrospective voting. As Ashworth, Bueno de Mesquita and Friedenbergr (2018) show, it is possible that these events systematically hurt incumbent officials' re-election prospects because they affect how much voters are able to learn about the quality of their incumbent officeholders.⁵ A cleaner test requires events that are truly irrelevant and cannot possibly convey, directly or indirectly, information about incumbents' attributes.

To this end, Healy, Malhotra and Mo (2010) examine the effects of sporting events on the incumbent party's electoral fortunes. Their hypothesis is

⁵Achen and Bartels (2016) argue that "a competent electorate must recognize that incumbents' preparations for, and reactions to, the substantial physical and social dislocations resulting from major disasters are, by definition, better than average about half the time," and they conclude that "we should expect to find rational electorates *rewarding* incumbents for better-than-average responses to natural disasters as often as they *punish* incumbents for worse-than-average responses" (p. 137). This argument is faulty for reasons that Ashworth, Bueno de Mesquita and Friedenbergr (2018) explain. Consult the motivating example, from their introduction, in which incumbents' re-election rate is lower in locations that suffered a natural disaster as a consequence of rational updating by the voter.

that, first, sporting events will affect mood, and voters will use their mood on Election Day as a “signal of the incumbent party’s success,” and, second, if they are feeling good, they will be more likely to access positive memories of what the incumbent party has done and to interpret its past actions favorably (p. 12804). The authors find that for local college football games no more than 10 days before the election, a victory by the home team increases the incumbent party’s vote share in presidential, Senate, and gubernatorial elections by 1.13 percentage points. Busby, Druckman and Fredendall (2017) replicate their finding using surveys of college students before and after a college football game. Students who were surveyed immediately after their football team lost were less approving of President Obama than students surveyed immediately before the game, while, at the opposing school, students surveyed immediately after their team’s victory were more approving than their counterparts surveyed immediately before the game. The authors also find that the game affected students’ mood, lending plausibility to the hypothesized mechanism for myopic voting, where shifts in one’s mood, whatever their causes, influence one’s assessments of unrelated aspects of the status quo.⁶

Huber, Hill and Lenz (2012) designed a laboratory experiment that tests

⁶Interestingly, the authors find that the effects of the football game on presidential approval completely disappear after a week. As they note, the disappearance is unsurprising if mood is the causal mechanism at work: over the course of a week, countless mood-altering events might occur, and any single event is unlikely to have a lasting impact.

the hypothesized mechanism. Experimental subjects received a sequence of 32 random monetary payouts from a computer “allocator.” The allocator was one of two types, differing in the expected average of the 32 payments. Subjects were not told their allocator’s type. After observing 16 payments, subjects could keep their allocator or replace it with a new one, whose type would also be unknown and would generate the remaining 16 payments. The first 16 payments provided information about the allocator’s type, analogous to the information voters receive about an incumbent officeholder’s competence over the course of his term. Rational experimental subjects would use this information to update their beliefs about their allocator’s type and choose to “re-elect” their allocator only after receiving evidence that it was the good type. In one version of the experiment, experimental subjects also received a payout from a separate random lottery—boosting their mood, presumably. Even after being informed that its outcome was unrelated to their allocator’s type, subjects who received positive lottery outcomes were more likely to retain their incumbent allocators than subjects who received negative lottery outcomes. The finding supports the mood-affiliation model of retrospective voting.

In light of such findings, it may seem implausible to put retrospective voting forward as a mechanism by which democracies subject leaders to meaningful degrees of popular control and electoral accountability.⁷ And

⁷One reason to await further research before putting too much stock in the validity of findings on myopic voting is that they have a mixed replication record. Fowler and Hall (2018), using a larger dataset comprising every recorded fatal shark between 1872 and 2012, find that the estimated

there are still other reasons one might harbor doubts. Achen and Bartels (2016) argue that voters choose whether to support the incumbent officials largely on the basis of partisan attachments, which are in turn expressions of social identities rather than the products of rational assessments of what the parties are likely to do. If that is right, then it seems that elected officials will have weak incentives to pay attention to voters' preferences over policy outcomes. Finally, even if we set these considerations aside, the instrument by which voters could possibly structure officeholders' incentives appears to be too crude to be effective: voters can remove or retain the incumbent, but they cannot scale their sanction to more fine-grained judgments of his performance. Some theorists express doubts that voters, even rational and effect on presidential elections is small and statistically indistinguishable from zero, and conclude that Achen and Bartels's finding on shark attacks is likely spurious. Fowler and Montagnes (2015*a*) conclude that the original finding of Healy, Malhotra and Mo (2010) was largely a false positive after conducting the same test but using NFL games. Even though NFL games are more popular than college games, the authors find no effect on election outcomes (for a reply from the original authors and a reply to their reply, see Healy, Malhotra and Mo (2015) and Fowler and Montagnes (2015*b*)). As Fowler and Hall (2018) explain, there are reasons to worry that the research on myopic voting is especially prone to false positives resulting from publication bias. It seems unlikely that prior to Achen and Bartels's study, an article titled "Shark Attacks Have No Effects on Election Outcomes" could have appeared in a top political science journal.

well-informed voters, could exercise control over the elected official with such a blunt tool (Dunn, 1999; Green, 2010).

§6.2 Popular control with myopic and partisan voters

To get a handle on the question, I now turn to a stylized model of retrospective voting. The model is not intended as a realistic description of democracy. Like models in the social sciences and political theory more generally, it is rather a tool for “conceptual exploration” (Hausman, 1992; Johnson, 2014): it serves as a kind of thought experiment whose purpose is to help clarify the meaning of concepts—of *retrospective voting*, *partisan bias*, *popular control*—and what is implied, or not, by claims that employ the concepts. Suppose the claims about myopic retrospective voting are true. What do they imply about popular control? It is easiest to see their implications if we consider a hypothetical scenario in which those claims hold but other extraneous facts about democracies are ignored. The model describes such a hypothetical scenario, and with it we can explore additional hypothetical scenarios that are “nearby” in logical space, by adjusting this or that feature of the model. I will use the model to argue that the various empirical claims that skeptics have treated as reasons for doubting that retrospective voting establishes popular control are in fact bad reasons for that conclusion.

I start with a model that is a minor adaptation to a “career concerns” model from Gehlbach (2013, ch. 7), in turn adapted from models in Persson and Tabellini (2000). In addition to helping us think through the implica-

tions of the empirical research on voting behavior, it is also the basis for an extension I consider later in the chapter, when I turn to the questions about retrospective voting with heterogeneous voters who use different performance standards.

The game has three players. In the first period, an incumbent officeholder chooses an action with consequences that a voter observes; the voter then chooses whether to re-elect the incumbent; and in a second period, either the re-elected incumbent from the first period or a newly elected challenger faces the same choice the incumbent faced in the first period.⁸ Let I and C refer to the incumbent and challenger, respectively. The politician holding office in period t receives a wage, w , and chooses a level of “effort,” $e^t \in [0, \infty)$, at a cost $-\frac{c}{2}(e^t)^2$, where $w, c > 0$.

One could adopt different interpretations of what “effort” and its cost represents, but, given my rhetorical purposes, here is the one I ask the reader to entertain. Suppose the politician can search for and, once identified, work to implement policies that would have an effect on some outcome of interest to the voter (the “policy outcome”). There are special constituencies and interest groups that oppose positive action by the politician to increase the value of the policy outcome, and $\frac{c}{2}(e^t)^2$ quantifies the value of favors and other benefits from these constituencies that the politician will sacrifice if he exerts effort e^t towards increasing the outcome of interest. That the

⁸In the extension to the model I consider below, I do without the simplifying assumption that the electorate can be treated as though it were a single voter.

cost is proportional to the square of the effort level captures the idea that the greater the effort he is exerting, the greater the marginal amount of contributions from the special interests that would be lost from exerting even further effort.

The policy outcome is assumed to reflect not just the actions taken by the officeholder but also the realization of a random variable outside the officeholder's control. Following the formal modeling literature, we will provisionally adopt the interpretation that this variable is the officeholder's "competence." Specifically, the officeholder's choice of effort, e^t , together with his level of "competence," determines a policy outcome g^t as follows:

$$g^t = \theta^P + e^t$$

where θ^P is the competence of the politician $P \in \{I, C\}$ holding office in period t . The officeholder's competence is not indexed to the time period because it is assumed to be fixed and to carry over from one period to the next. I return to questions about the possible interpretations of this variable below.

The voter does not observe the incumbent's effort, and neither the voter nor the incumbent observes his competence, θ^I , before the incumbent chooses his level of effort.⁹ Players know only that θ^P is drawn (in-

⁹One virtue of this assumption is that it allows one to bracket the complications and distinct problems that would arise from positing asymmetric information about the incumbent's type. But in this context, it serves a different rhetorical purpose: since the politician has no control over or

dependently for the incumbent and challenger) from a strictly increasing, differentiable cumulative distribution function F with density f . Assume that the distribution has a mean of zero, and f is strictly increasing below and strictly decreasing above its mean.

The policy outcome g could be a measure of economic growth, unemployment, or inflation; the number of casualties from a war; a measure of social justice, such as the degree of inequality, protection for basic liberties, or fairness in the distribution of economic opportunities; or some other outcome that the voter cares about. The key assumption is that the outcome is observable to the voter, and its realized value reflects both the incumbent's action—which we're calling "effort"—and one of his qualities that carries over from one term of office to the next—what we're calling "competence." Under these assumptions, the voter will draw inferences about the incumbent's "competence" from the observed outcome, and the incumbent, anticipating these inferences, will have incentives to take actions that improve the policy outcome.

To accommodate skeptics' claims about the influence of mood, partisan bias, or other non-instrumental motives on voting behavior, we assume that the identity of the politician holding office in a given period affects the voter's payoff irrespective of the politician's competence and effort level. Specifically, in period t , if politician $P \in \{I, C\}$ holds office and the outcome knowledge of the realization of θ^P , we can compare its role in the model, and the implications for popular control, to the role that theorists of myopic voting attribute to random events like natural disasters.

is g^t , then the voter's payoff is

$$vg^t + b^P$$

where b^P is the voter's bias in favor (if $b^P > 0$) or against (if $b^P < 0$) politician P , and $v > 0$ captures the importance the voter attaches to the policy outcome g^t relative to the identity of the officeholder.

To summarize, the game unfolds as follows:

- the incumbent chooses e^1 ,
- the incumbent's competence, θ^I , is drawn from the distribution F ,
- the voter observes g^1 and then chooses to vote for the incumbent or the challenger,
- the winner of the election chooses e^2 ,
- the challenger's competence, θ^C , is drawn from the distribution F , and all players' payoffs are realized.

We assume the players' behavior conform to the requirements of a perfect Bayesian equilibrium. Officeholders in the second period always choose $e^2 = 0$ because they no longer have any re-election incentive, so the voter optimally votes for the incumbent if and only if she is sufficiently confident in his superior competence to outweigh any bias she has for the challenger. In a perfect Bayesian equilibrium, the voter updates her beliefs about the incumbent's competence on the basis of the observed outcome and an accurate belief about the effort he exerts. The incumbent, in response, has

an incentive to exert the amount of effort the voter expects, lest the voter attribute the worse-than-expected performance to the incumbent's worse-than-expected competence. Further details of the equilibrium and its explanation can be found in the appendix. Here I take the conclusions of the analysis as given and consider the lessons one can draw from it.

To reduce notation, let $b := b^C - b^I$. As I explain in the appendix, in the game's unique equilibrium, the incumbent chooses the following level of effort in the first period:

$$e^1 = \frac{w}{c} f(b/v) \quad (6.1)$$

and the voter votes for the incumbent if and only if, conditional on g^1 , she believes the incumbent's competence is equal to or greater than b/v .

To understand the equation, note that the term $f(b/v)$ measures the incumbent's uncertainty about whether his competence will be close to b/v , which is his re-election threshold in the equilibrium. (Recall that f is the probability density function, so $f(x)\Delta$ is roughly approximate to the probability that θ^I lies in the interval $[x, x + \Delta]$, for small values of Δ . Higher values of $f(x)$ correspond to comparatively greater certainty that the variable will be close to x .) In other words, it measures his confidence that increased effort will make a difference to the election outcome: if he is confident his realized competence will not be near the threshold, then incremental increases or decreases to the threshold, resulting from marginal changes in effort, are unlikely to affect the election outcome.

Inspecting (6.1), we see that the incumbent's incentives to exert effort are increasing in the value of holding office, w , decreasing in the cost of

exerting effort, c , and increasing in his certainty, $f(b/v)$, that his effort will be pivotal to the election outcome.

We are now in a position to consider the implications of partisan bias and myopic voting. I first consider their implications for the voter's *welfare*; I then consider their implications for the voter's *control*.

The extent of the voter's partisanship can be measured as $|b|/v$, which is the amount of the policy outcome the voter would be willing to "pay" in order to get an officeholder of the preferred party.¹⁰ Recall that we assumed f is strictly increasing below zero and strictly decreasing above zero. Thus, as partisanship, $|b|/v$, approaches zero, the incumbent exerts greater effort. The voter's welfare is of course strictly increasing in his choice of effort.

Thus, news that partisanship plays a greater role in voting behavior than we previously realized should arguably make us less sanguine about elections as a means of promoting voter welfare. But that does not mean we should despair of elections as a means of promoting voter welfare. In the equilibrium, so long as partisanship is not the *only* thing that matters to the voter—so long as $v > 0$ —the incumbent still exerts more effort than he would in the absence of the re-election incentive. Moreover, there are tools for improving the welfare properties of the election mechanism. Increase the

¹⁰If $b = b^C - b^I > 0$, then the voter would be willing to give up $b/v = |b|/v$ units of the outcome in order to have the challenger in office: $vg + b^I = v(g - b/v) + b^C$. If $b = b^C - b^I < 0$, then the voter would be willing to pay $-b/v = |b|/v$ units of the outcome in order to have the incumbent in office: $v(g - (-b/v)) + b^I = vg + b^C$.

value of holding office, relative to the cost of exerting effort, and the voter's welfare improves.

So if we learn that voters are more partisan than we previously thought, the right conclusion is not necessarily to lower our confidence in elections, or retrospective voting specifically, as instruments of voters' welfare, as Achen and Bartels and others suggest. The right conclusion may instead be to increase elected officials' salaries (w) above what we previously thought was optimal. Or to adopt more extreme measures, beyond what we previously thought was optimal, to reduce elected officials' opportunity costs (c) from enacting good policies—the foregone favors and retirement opportunities that an official might jeopardize by acting contrary to the preferences of special interests, for example.

Now let us consider the implications of myopic voting. As I have told the story, the voter is not myopic: she forms an accurate belief about the effort that the incumbent exerts, and she rationally updates her beliefs about his competence after observing the policy outcome. But notice that she resembles the myopic voter in this respect: she lets a random variable that is outside the incumbent's control influence her retrospective assessments. We interpreted this random variable as a quality of the incumbent (“competence”) that positively affects the policy outcome and carries over from one period to the next. But the incumbent would have similar incentives to exert effort if the random variable θ^I were an irrelevant event, mere noise in the causal link between the incumbent's choice of effort and the policy outcome, so long as the incumbent knew that the voter would draw inferences (rationally or myopically) about the incumbent's competence from the

observed policy outcome.

Here is one way to make the point. Imagine a modified game, in which everything is the same as the game just presented, except that there are two states of the world, $s = 1$ and $s = 2$. If $s = 1$, then $g^t = \theta^P + e^t$, where $P \in \{I, C\}$ is the politician in office in period t , and θ^P is drawn from F , as before. If $s = 2$, then $g^t = \varepsilon^t + e^t$, where ε^1 and ε^2 are drawn independently from a distribution G , and neither the politicians nor the voter observe their realized values. In one state, the policy outcome is an informative signal of how the incumbent would perform in the future; in the other state of the world, it contains no information about future performance. Assume G is differentiable with density g and has full support on $(-\infty, \infty)$. The first-period politician, I , observes the value of s but the voter does not. Let $p := Pr(s = 1)$ be the voter's belief that the state is $s = 1$. In all other respects—payoff functions, timing—the game is identical to the baseline version.

If p is close to 1 (the voter is confident that g^1 is an informative signal of the incumbent's competence), then the voter will behave essentially as she did before, using an almost identical re-election rule. And if $s = 2$, then this behavior will look myopic to the incumbent—as well as to any social scientist who also knows the true state of the world—because the voter will vote as if the first-period outcome signaled something relevant about the incumbent, when in fact it does not.

But the incumbent does not care about the rationality of the voter's re-election rule; he cares about the prospects of his re-election. He will still have incentives to exert effort on the voter's behalf, lest the voter (wrongly)

draw critical inferences about the incumbent's competence. As I explain in the appendix, his equilibrium choice of effort will be

$$e^1 = \frac{w}{c} f(b/pv)$$

if $s = 1$, and

$$e^1 = \frac{w}{c} g(b/pv)$$

if $s = 2$. Note that in the first state, when the policy outcome is an informative signal of his competence, he behaves essentially as he does in the baseline version of the model, provided the voter is sufficiently confident about the state.

Importantly, the incumbent takes actions to improve the policy outcome *in either case*. Indeed, the voter could even be *better off* in state 2, the scenario where she looks like a myopic voter. This would happen if $g(b/pv) > f(b/pv)$. Imagine the voter is highly partisan and heavily biased in favor of the challenger (so that b/pv is a large positive number). And imagine the incumbent is extremely uncertain about the noisy influences on the policy outcome, putting significant weight on the possibility of extreme realizations in either direction (the distribution g has fat tails), while he is highly confident in his own mediocrity, sure that his competence is about average (the distribution f is concentrated around zero, its mean). Then we could have $g(b/pv) > f(b/pv)$, and the incumbent would exert more effort when he knows that the noisy, meaningless variable ε^1 influences the policy outcome but his competence does not, compared to when he knows that his

competence influences the outcome but the noisy variable does not.¹¹

In conclusion, neither empirical findings that voters are more partisan than we thought nor findings that their retrospective voting is myopic are good reasons for doubting that officeholders have incentives to try to improve the policy outcomes voters care about. These incentives are obviously not perfect, but incentives rarely are.

We can use (6.1) to draw conclusions about the voter's *control* over the incumbent, too. In this model, the voter can be described as having control over the policy outcome that the incumbent aims to achieve. I emphasize this interpretation because, among other reasons, it formalizes Christiano's "choice of aims" model of popular rule (Christiano, 1996). If the incumbent chooses an effort level e^1 , he expects the policy outcome to be $E(g^1|e^1) = E(\theta^I) + e^1 = e^1$. This is the policy outcome he can be said to be "aiming" at when he chooses e^1 . Thus, the set of policy outcomes that the incumbent could choose to aim at is just the set $[0, \infty)$, and we can reinterpret the incumbent's action variable not as a level of "effort" per se but rather as the policy outcome he aims to achieve. I will continue to refer to the action variable as "effort," for the sake of continuity with familiar models of political agency, but ask the reader to keep this interpretation in mind.

To evaluate the voter's control over the incumbent's effort, we consider how his equilibrium choice of effort responds to changes in the voters' preferences and the weight, v , that she attaches to effort relative to the office-

¹¹I thank David Wiens for pointing out this possibility.

holder's partisan identity. Specifically, we ask whether there is a range of effort levels such that for any two values in the range, if she preferred one to the other, and attached enough importance to the officeholder's effort relative to partisan identity, then he would not choose the less preferred effort level. To fix the meaning of the counterfactuals, let us restrict attention to the effects of variation in v across the domain $(-\infty, \infty)$. Given this restriction, when we assess the preference condition, we are asking about the incumbent's equilibrium actions in counterfactual versions of the game, corresponding to different values of the parameter v across this domain $(-\infty, \infty)$. When $v < 0$, the voter prefers smaller values of the policy outcome and thus prefers for the incumbent to exert no effort (or to "target" the policy outcome of zero, the smallest outcome he is capable of targeting).

Consider any effort levels $e, e' \in [0, \frac{w}{c}f(0)]$ such that the voter prefers e to e' . I claim that if the voter attaches enough weight to effort, then the incumbent will not choose e' in equilibrium.

Case (i): $e < e'$, implying $v < 0$. Then the voter would prefer for the incumbent to choose $e^1 = 0$ rather than any positive level of effort. Readers can verify for themselves the intuitive claim that in the unique equilibrium to that counterfactual game, the incumbent chooses $e^1 = 0$. Case (ii): $e > e'$, implying $v > 0$. From (6.1), we see that for any sequence of positive numbers v_n , the incumbent's equilibrium effort choice approaches $\frac{w}{c}f(0)$ as $v_n \rightarrow \infty$. Thus, if v is large enough, then e' will not be chosen, because with v large enough the incumbent's equilibrium behavior induces an outcome arbitrarily close to the upper bound $\frac{w}{c}f(0)$.

Thus, for any pair of effort levels in the interval $[0, \frac{w}{c}f(0))$, the voter's

preferences over the pair have a constraining effect, if the voter places enough importance on the incumbent's effort relative to the incumbent's partisan identity. The voter has some control over the incumbent's choice of effort (alternatively, over the policy outcomes he aims at), even though partisan biases influence her voting behavior.

Myopic voting is also compatible with popular control. Imagine, again, the modified version of the game described above, in which the voter mistakenly believes the observed outcome in the first period is predictive of the outcome in the second, while, in fact, better or worse than expected outcomes in the first period merely reflect meaningless noise. Then the incumbent will choose, in equilibrium, $e^1 = \frac{w}{c}g(b/pv)$, as explained in the appendix. This equilibrium behavior fixes the upper bound on the range of effort levels that the voter's preferences influence. The range of effort levels over which the voter now has control will be $[0, \frac{w}{c}g(0)]$. Depending on her preferences, and how much weight the voter attaches to effort relative to the officeholder's partisan identity, any choice of effort in the range can be supported in equilibrium. For any pair of effort levels, the voter's preferences over the pair will have a constraining effect, if she places enough importance on effort relative to the incumbent's partisan identity. The explanation is exactly the same as before, except with g taking the place of f . Thus, she has control, even when she appears to the incumbent (and social scientists) as a myopic voter.

Indeed, she may even have more control when she behaves like a myopic voter. It depends on whether $g(0)$ or $f(0)$ is larger, as these quantities fix the upper bound on the interval of effort levels over which the voter has control

in the two different circumstances. The voter would have more *more control* in the second state, where she appears myopic to the incumbent, if $g(0) > f(0)$. We do not have strong theoretical reasons to rule this possibility out. Imagine that in the first state of the world, where competence is the unknown variable, distributed according to the density f , the incumbent is not at all sure whether he is close to average in competence (so that $f(0)$ is small); while, in the second state of the world, where meaningless noise like football games or shark attacks are the unknown variable, the incumbent is fairly confident that these influences will have small effects on the observed measure of performance, close in magnitude to zero (so that $g(0)$ is large). Then the range of effort levels that the voter's preferences constrain is larger in the second state of the world, where the voter appears myopic, than in the first, where she acts just as she would if fully informed about the state.

§6.3 Popular control with heterogeneous voters

Like the model from the previous section, most formal models of political agency in the literature make the simplifying assumption that the electorate is a single voter. A skeptic might doubt whether the simplifying device is as innocuous as the literature assumes. Suppose that majorities have cyclical preferences over the possible policy outcomes that the incumbent can bring about, such that, no matter how he acts, he could produce superior combinations of policy outcomes in the eyes of a majority of voters if he acted differently. Would that be a reason to believe that retrospective voting will fail to produce some degree of popular control?

To address that question, I now consider an extension of the model from the previous section. It is worth repeating that the model is not intended as a realistic description of actual democracies. It serves as a thought experiment. We will imagine a hypothetical democracy in which the skeptics' claims about voters are true, and we will then check whether the skeptical conclusion about popular control follows. As with the baseline version from the previous section, the value of the model is that it helps us clarify the meaning of the concepts used to formulate both the skeptics' assumptions and their conclusion about popular control, so that the logical connection between them, or lack thereof, is easier to see. After presenting the model, I use it to assess, in particular, the argument that if voters' re-election standards reflect how well they have fared as individuals under the incumbent, rather than some measure of macroeconomic performance, then the elected official may be able to pursue a "divide-and-rule" strategy, pitting different groups of voters against each other and avoiding electoral control (Hibbs, 2006, p. 570; Achen and Bartels, 2016, p. 99).

We modify the previous model so that there are three voters, indexed $i = 1, 2, 3$, and three different policy outcomes, labeled g_1 , g_2 , and g_3 , and we assume voter i cares about outcome g_i but not the other two outcomes.¹² In each period t , whichever politician $P \in \{I, C\}$ holds office

¹²The model is similar in important respects to that in Persson and Tabellini (2000, pp. 228–230), although it incorporates the minor adaptation from the baseline version above, assuming that voters care about partisan identity and explicitly modeling the amount of relative importance they attach to partisan identity.

chooses $e_i^t \in [0, \infty)$, for $i = 1, 2, 3$. His choice of effort levels determines the three outcomes together with his outcome-specific competencies θ_i^P , for $i = 1, 2, 3$, as follows:

$$g_i^t = \theta_i^P + e_i^t$$

As before, we assume each voter's preference between the two politicians also reflects a "bias" for or against the incumbent, so that voter i 's payoff in period t , if politician $P \in \{I, C\}$ holds office and has chosen e_i^t , is

$$v_i g_i^t + b_i^P$$

where v_i and b_i^P have the same interpretations as above except they are now indexed to individual voters. To reduce the clutter of notation, we use the abbreviation $b_i := b_i^C - b_i^I$.

Each politician's competence θ_i^P in area i is an independent draw from a strictly increasing, differentiable distribution F_i with density f_i and $E(\theta_i^P) = 0$. Neither voters nor the politicians observe their competencies, as before.

A politician receives a payoff $w - \frac{c}{2} \sum_i (e_i^t)^2$ if he holds office in period t and a payoff of zero otherwise.

Let me comment briefly on this assumption about the aggregate cost of effort levels. In line with the interpretation suggested in the baseline version of the model, imagine that there are special constituencies and interest groups that oppose positive action in each of the policy areas, and $\frac{c}{2}(e_i^t)^2$ quantifies the value of favors and other benefits from these constituencies that the politician will sacrifice if he exerts effort e_i^t in order to increase the

outcome g^t (or, under the “choice of aims” interpretation of the model, if he chooses to aim at the policy outcome e_i^t). The politician would like to minimize the sum of these costs, all else being equal. The key assumption is that the greater the effort e_i^t devoted to an outcome g_i^t , the greater the marginal value of the benefits from the special interests—those concerned with that particular outcome—that the elected official forfeits by increasing his effort further still.

Interested readers can find the full analysis of the game in the appendix; here I take the conclusion as given and examine its implications. In the game’s unique perfect Bayesian equilibrium, the first-period incumbent chooses

$$e_i = \frac{w}{c} f(b_i/v_i) \underbrace{[F_j(b_j/v_j)(1 - F_k(b_k/v_k)) + (1 - F_j(b_j/v_j))F_k(b_k/v_k)]}_{\text{probability that voter } i \text{ is pivotal in equilibrium}} \tag{6.2}$$

for all $i, j, k \in \{1, 2, 3\}$ such that $i \neq j \neq k$.

The logic of the situation is similar to what we saw in the baseline version of the model. In equilibrium, the voter has a correct expectation about how much effort the incumbent chooses to exert, and the incumbent is deterred from exerting less effort on outcome g_i than expected by the danger that voter i will infer he is less competent, in that area, than the typical challenger. What distinguishes the extended model from the earlier version is that the incumbent now discounts this danger by the probability that voter i is pivotal. If the voter is unlikely to be pivotal, there is little reason to worry about the inferences she will draw.

We can use the equations in (6.2) to answer questions about popular control in the model. The set of feasible vectors of effort levels is $[0, \infty)^3$. (As before, we can also interpret this set as the set of vectors of policy outcomes that the incumbent could feasibly aim to achieve: if he chooses a vector e , he also expects e to be the vector of policy outcomes.) Majorities' preferences over the feasible set are cyclical, and no Condorcet winner exists. Nonetheless, the conception of popular control from chapter 5 specifies conditions for popular control that can, in principle, be satisfied even if no Condorcet winner exists. I claim those conditions are satisfied in the equilibrium.

Fix a majority—voters 1 and 2, say. What we want to show is that there is some range of effort vectors, such that for any two vectors in the range, if 1 and 2 both preferred one vector to the other and attached sufficient weight to effort relative to the partisan identity of the elected official, then their shared preference would constrain the incumbent's choice of effort vector, preventing the less preferred of the two from occurring. As in our analysis of the baseline version of the model, for the sake of simplicity and tractability I will assess the counterfactuals relative to counterfactual variation across a particular domain of preferences, namely all those payoff functions for the voters of the form $v_i g_i + b_i^P$ where $v_i \in (-\infty, \infty)$ for $i = 1, 2$.

From (6.2), we can see that as v_1 and v_2 become large, the probability that 1 is pivotal approaches $\pi_1 := F_2(0)(1 - F_3(b_3/v_3)) + (1 - F_2(0))F_3(b_3/v_3)$, and so the equilibrium effort e_1 approaches $\frac{\pi_1 w}{c} f(0)$. Likewise, as v_1 and v_2 become large, the probability that 2 is pivotal approaches $\pi_2 := F_1(0)(1 - F_3(b_3/v_3)) + (1 - F_1(0))F_3(b_3/v_3)$ and so the equilibrium effort e_2 approaches $\frac{\pi_2 w}{c} f(0)$. These two upper bounds on the equilibrium effort levels define a

range of the incumbent's actions over which 1 and 2 have control. For any two vectors of effort in that range, if 1 and 2 have the same preference over them, and $|v_1|$ and $|v_2|$ are large enough, then the incumbent will not choose the less preferred of the two. (The details of the argument can be found in the appendix.) There was of course nothing special about the majority comprising 1 and 2, so we conclude that every majority has some control over the vectors of policy outcomes.

The model describes a hypothetical scenario in which voters have heterogeneous preferences over policy outcomes that do not aggregate into anything resembling a popular will (majorities' preferences are cyclical, no Condorcet winner exists). They use different standards for assessing the incumbent's performance. The measures of performance the three voters use, the outcomes g_1^1 , g_2^1 , and g_3^1 , could be interpreted as measures of each voter's own welfare, or the welfare of some group of which she is a representative member. Hibbs (2006) claims that pure retrospective voting will produce electoral discipline only if voters "react to macroeconomic performance, rather than individual benefits" (p. 570). Achen and Bartels (2016) quote him in apparent agreement (p. 99). But that inference is not sound, absent further assumptions. Under one interpretation, the model describes a scenario in which voters react to the individual benefits rather than macroeconomic performance. Yet their retrospective voting behavior produces electoral discipline and indeed popular control.

In making this point, I am not denying that heterogeneity in voters' standards *could*, in conjunction with other circumstances, undermine popular control. The point is just that the conclusion does not follow straight

from the assumption about heterogeneity. Moreover, to my knowledge, no one has yet told a credible story about why heterogeneity in voters' standards, in conjunction with other circumstances, would have this effect. Hibbs's remark appears to be made on the basis of Ferejohn's (1986) canonical model of electoral accountability. In that model, the incumbent officeholder chooses how to divide resources among himself and self-interested voters. For every conceivable division, a majority of voters prefer a different division. In equilibrium, voters in effect compete to be included in the majority coalition that receives a share of the pie, with the unintended result that the incumbent is able to take everything for himself. But Ferejohn's model does not provide a robust justification for Hibbs's pessimistic conclusion, however. In the model the incumbent official and the challenger are identical in all relevant respects—they have the same preferences, knowledge, and abilities. The re-elected incumbent will behave no differently from the challenger if elected, so a rational voter has no reason to prefer one or the other. The voter's indifference between the candidates figures crucially in the explanation for why electoral accountability would break down if voters' retrospective assessments reflected heterogeneous standards, such as their own welfare. Voters can "compete" to be included in the incumbent's coalition only because they can credibly promise to lower their re-election standards, which in turn is possible only because at the time of the election, they are actually indifferent between the two candidates.

Models of political agency in which all agents have the same type, but the agents' actions are unobserved by the principal, are labeled 'pure moral

hazard’ models in the literature.¹³ The conclusions of moral hazard models are objectionably sensitive to the assumption that all elected officials are the same type (Fearon, 1999). Recent literature has therefore favored a different approach to modeling political agency. In the more widely used ‘selection’ models of political agency, voters’ standards reflect their beliefs about the incumbent’s qualities, updated on the basis of the evidence his performance gives them.¹⁴ The models in this chapter are examples. Voters cannot raise or lower their re-election standards in these models because

¹³Barro (1973) and Ferejohn (1986) are the canonical examples. See Ashworth (2012) for a review of the literature on formal models of political agency. In previous work, I have described a pure moral hazard model, in the same vein as Ferejohn’s, in which voters condition their votes on individual benefits received but their re-election rules still produce a kind of electoral discipline (Ingham, 2016*b*).

¹⁴Mansbridge (2009) uses the term ‘selection model’ to describe a different kind of principal-agent relationship, in which the agent has “self-motivated, exogenous reasons for doing what the principal wants” and the principal and agent “thus have similar objectives even in the absence of the principal’s sanctions” (p. 369). This usage differs from the meaning of the term here; in the ‘selection models’ I am describing, the elected official is only motivated to exert effort on the voters’ behalf because they fear being removed from office if they do not. What they have in common with the relationship Mansbridge describes, and what explains the common label, is that in each, the principal is concerned to select the right kind of agent. In the pure moral hazard models, by contrast, all agents are alike in relevant respects.

they do not *choose* what to believe about the incumbent. If one takes this approach, it is unclear whether one could tell plausible story about a “race to the bottom” dynamic like in Ferejohn (1986).

§6.4 Conclusion

One promise of theories of retrospective voting is to have identified a plausible mechanism by which voters, despite not being well-informed about the actions of their elected officials, might be able to exert some control over their elected leaders. The purpose of this chapter was to illustrate how one would formulate this thought using the multiple majorities conception of popular control as well as to address two sources of skepticism about the mechanism.

Skeptics have argued that retrospective voting is myopic. Voters allegedly treat random, irrelevant events—shark attacks, natural disasters, football games, and so on—as though they were informative of how well the incumbent is likely to govern if re-elected. And skeptics have pointed to the strong influence of partisan loyalties, allegedly rooted in a subjective sense of social identity rather than objective interests, as a further reason to doubt that retrospective voting provides meaningful incentives for elected officials to act in voters’ interests. Skeptics have also suggested that if voters’ retrospective assessments of an elected official reflect heterogeneous standards—perhaps each voter’s own welfare, or the welfare of different groups, or just different policy outcomes—then there is a reason to doubt that retrospective voting produces a meaningful degree of electoral accountability and popular

control.

The simple models of retrospective voting considered above suggest that each inference is too quick. In the first, the representative voter is influenced by partisanship, and the model permits an interpretation under which the voter mistakenly treats statistical noise as evidence of the incumbent’s qualities. Nonetheless, the incumbent has incentives to take the voter’s preferences into account and to try to improve the policy outcome, lest the voter draw inferences—perhaps with good reason, perhaps mistakenly—about the incumbent’s qualities. In virtue of these incentives, the voter has some control over the incumbent’s choice of effort. In the second model, we described a hypothetical scenario in which voters make different retrospective assessments, reflecting heterogeneous preferences over policy outcomes that fail to aggregate into a collective will. Nonetheless, majorities have control over the incumbent’s choice of effort levels, owing to the incentives that their retrospective voting gives elected officials. Or we might say—using the alternative interpretation, mentioned above, that brings out the model’s relevance for Christiano’s (1996) “choice of aims” account of popular rule—majorities have control over the incumbent official’s *goals*, the combinations of policy outcomes that he chooses to aim at.

§6.5 Appendix

§6.5.1 The model with one voter

The baseline version of the model (with one representative voter) is only a minor adaptation of the career concerns model in §7.2 of Gehlbach (2013),

and its unique equilibrium can also be found by the kind of argument Gehlbach gives. I include the argument here so that the reader does not need to consult an additional text.

In any perfect Bayesian equilibrium, the voter forms beliefs about the expected value of θ^I , conditional on the observed outcome g and a belief about the incumbent's strategy, using Bayes' rule. After observing an outcome g^1 , the voter believes that

$$E(\theta^I | g^1, \tilde{e}) = g^1 - \tilde{e}$$

where \tilde{e} is the level of effort that the voter believes (correctly, in an equilibrium) the incumbent chooses. Since the second-period officeholder chooses $e^2 = 0$ in any equilibrium, and $E(\theta^C) = 0$, the voter optimally votes to re-elect the incumbent if and only if

$$g^1 - \tilde{e} \geq b/v$$

where $b = b^C - b^I$, as in the text. The probability that the incumbent wins re-election, conditional on choosing e^1 in the first period, is the probability of a realization of θ^I such that the inequality holds. This probability is $1 - F(b/v + \tilde{e} - e^1)$, where we have used the fact that $g^1 = \theta^I + e^1$. The incumbent chooses his level of effort to solve

$$\max_{e \in [0, \infty)} w [1 - F(b/v + \tilde{e} - e)] - \frac{c}{2} e^2$$

Assume that c/w is large enough that the objective function is strictly con-

cave. Then a necessary and sufficient condition for a solution is

$$e = \frac{w}{c} f(b/v + \tilde{e} - e)$$

In an equilibrium, the voter's belief about the incumbent's strategy is correct. So, substituting $e = \tilde{e}$, the condition implies

$$e = \frac{w}{c} f(b/v)$$

The model of the myopic voter

The game resembles the baseline version in all respects, except that there are two states of the world, $s \in \{1, 2\}$. If $s = 1$, then $g^t = \theta^P + e^t$, where $P \in \{I, C\}$ is the politician in office in period t , and θ^P is drawn from F , as in the baseline version. If $s = 2$, then $g^t = \varepsilon^t + e^t$, where ε^1 and ε^2 are drawn independently from a distribution G , and neither the politicians nor the voter observe their realized values. Assume G is differentiable with density g and has full support on $(-\infty, \infty)$.

The first-period politician, I , observes the value of s but the voter does not. Let $p := Pr(s = 1)$ be the voter's belief that the state is $s = 1$. In all other respects—payoff functions, timing—the game is identical to the baseline version.

In a perfect Bayesian equilibrium, the second-period officeholder chooses

$e^2 = 0$. The voter, conditional on observing g^1 , forms the belief that

$$\begin{aligned} E(\theta^I | g^1, \tilde{e}) &= pE(\theta^I | g^1, \tilde{e}, s = 1) + (1 - p)E(\theta^I | g^1, \tilde{e}, s = 2) \\ &= p(g^1 - \tilde{e}) \end{aligned}$$

where we use the fact that $E(\theta^I | g^1, \tilde{e}, s = 2) = E(\theta^I) = 0$. Thus, the voter optimally re-elects the incumbent if and only if

$$p(g^1 - \tilde{e}) \geq b/v$$

where $b = b^C - b^I$, as before.

If $s = 1$, then the probability that the voter votes for the incumbent, as a function of e , is the probability that $\theta^I \geq \frac{b}{pv} + \tilde{e} - e$. Since the incumbent learns the state, we can use an argument similar to the one used in the baseline model to conclude that the incumbent's optimal action is

$$e^1 = \frac{w}{c} f\left(\frac{b}{pv}\right)$$

If $s = 2$, then the probability that the voter votes for the incumbent, as a function of e , is the probability that $\varepsilon^1 \geq \frac{b}{pv} + \tilde{e} - e$. Using an argument similar to the one used in the baseline model, but substituting the distribution G for the distribution F , we can conclude that the incumbent's optimal action is

$$e^1 = \frac{w}{c} g\left(\frac{b}{pv}\right)$$

Note that if p is close to 1, then the voter behaves essentially like the

voter in the baseline model. But if $s = 2$, then the voter looks to the incumbent (and the social scientist, if the social scientist knows that $s = 2$) like a myopic voter. For she votes as if the better or worse than expected realizations of the outcome g^1 were evidence that the incumbent is better or worse than expected, when in fact they reflect uninformative noise.

The model with three voters

I now turn to the analysis of the extended version of the model with three voters and three different policy outcomes.

We look for a perfect Bayesian equilibrium in which voters' strategies are weakly undominated.¹⁵ Similar to the baseline version of the game, the politician holding office in the second period chooses $e_i^2 = 0$ for each $i \in \{1, 2, 3\}$, and voters use Bayes' rule to form beliefs about θ_i^I conditional on the outcomes g_i^1 and the incumbent's strategy. After observing an outcome g_i^1 , voter i believes that $E(\theta_i^I | g_i^1, \tilde{e}_i) = g_i^1 - \tilde{e}_i$, where \tilde{e}_i refers to the voters' belief (correct, in equilibrium) about the incumbent's choice of e_i^1 . Thus, voting for the incumbent is consistent with voter i using a weakly undominated strategy if and only if

$$v_i E(\theta_i^I | g_i^1, \tilde{e}_i) \geq b_i$$

¹⁵In games in which three or more players vote in an election decided by simple majority rule, the restriction to weakly undominated voting strategies is needed to eliminate equilibria in which players' voting strategies are optimal only in virtue of their (self-fulfilling) beliefs that no player's vote will be pivotal.

where $b_i := b_i^C - b_i^I$. The probability that voter i casts a vote for the incumbent, conditional on having chosen $e = (e_1, e_2, e_3)$, is the probability of a realization of θ_i^I such that the above inequality holds. This probability is $1 - F_i(b_i/v_i + \tilde{e}_i - e_i)$, where we have used the fact that $g_i^1 = \theta_i^I + e_i^1$.

The incumbent chooses $e \in [0, \infty)^3$ to solve

$$\max_{e \in [0, \infty)^3} V(e) := w\Pi(e) - \frac{c}{2} \sum_i e_i^2$$

where $\Pi(e)$ is the probability that at least two of the three voters vote for the incumbent, conditional on the incumbent's having chosen $e = (e_1, e_2, e_3)$ and voters' using the re-election rules just described.

Let $\bar{e}_i := \frac{w}{c} \max\{f(x)\}_{x \in \mathbf{R}}$, and let $\bar{e} := \max\{\bar{e}_1, \bar{e}_2, \bar{e}_3\}$. Since V is continuous, and $[0, \bar{e}]^3$ is compact, we know V achieves a maximum on this set, which we will denote (e_1^*, e_2^*, e_3^*) . If $e_i > \bar{e}$, then

$$\frac{\partial V}{\partial e_i}(e_1, e_2, e_3) = wf(\tilde{e}_i - e_i + b_i/v_i)\pi_i(e) - ce_i < 0$$

where, to reduce notational clutter, we use $\pi_i(e)$ to refer to the probability that i is pivotal when the incumbent chooses e and voters use the re-election rules described. Thus, (e_1^*, e_2^*, e_3^*) is a global maximum. Since $f_i > 0$ and $\pi_i(e) > 0$ for all $e \in [0, \infty)^3$, the right partial derivative of V with respect to e_i is strictly positive when $e_i = 0$. Thus, (e_1^*, e_2^*, e_3^*) must lie in the interior of $[0, \infty)^3$ and satisfy the first-order condition

$$wf_i(\tilde{e}_i - e_i + b_i/v_i)\pi_i(e) - ce_i = 0, \quad i = 1, 2, 3$$

In an equilibrium, the voters' beliefs about the incumbent's strategy are correct, so $\tilde{e}_i = e_i$ for each $i \in \{1, 2, 3\}$. Making the substitution (and writing out the expressions for π_i), the equations reduce to

$$e_i = \frac{w}{c} f_i(b_i/v_i) \underbrace{[F_j(b_j/v_j)(1 - F_k(b_k/v_k)) + (1 - F_j(b_j/v_j))F_k(b_k/v_k)]}_{\text{probability that voter } i \text{ is pivotal in equilibrium}},$$

for all $i, j, k \in \{1, 2, 3\}$ such that $i \neq j \neq k$.

Let me now fill in the details of the argument for why majorities have control over the incumbent's choice of effort levels. As in the text, take the majority comprising 1 and 2, and consider variation in their preferences captured by variation in the parameters v_1 and v_2 across the domain $(-\infty, \infty)$. Take any two vectors of effort levels, e and e' , such that 1 and 2 each strictly prefer e to e' .

Case (i): $v_1 < 0$ or $v_2 < 0$, implying either $e_1 < e'_1$ or $e_2 < e'_2$, by the assumption 1 and 2 each prefer e to e' . If $v_i < 0$, then the incumbent optimally exerts zero effort on i in equilibrium (as the reader may verify) and so does not choose e'_i . Thus, e' is not chosen.

Case (ii): $v_1 > 0$ and $v_2 > 0$, implying $e_1 > e'_1$ and $e_2 > e'_2$, by the assumption 1 and 2 each prefer e to e' . As in the text, let

$$\pi_1 := F_2(0)(1 - F_3(b_3/v_3)) + (1 - F_2(0))F_3(b_3/v_3)$$

and

$$\pi_2 := F_1(0)(1 - F_3(b_3/v_3)) + (1 - F_1(0))F_3(b_3/v_3)$$

Note that the probability that 1 is pivotal approaches π_1 as $v_2 \rightarrow \infty$, while

the probability that 2 is pivotal approaches π_2 as $v_1 \rightarrow \infty$. And, for each $i = 1, 2$, $\frac{w}{c} f_i(b_i/v_i) \rightarrow \frac{w}{c} f_i(0)$ as $v_i \rightarrow \infty$. Thus, one can choose \bar{v} large enough so that in the equilibria to the counterfactual games in which $v_1, v_2 \geq \bar{v}$, the incumbent will choose e_i so that it is close to $\frac{\pi_i w}{c} f_i(0)$, for $i = 1, 2$, implying that e' is not chosen.

Cases (i) and (ii) exhaust the ways in which 1 and 2 could each strictly prefer e to e' . Thus, when they share a preference over two vectors of effort levels, the incumbent will not choose the less preferred of the two, provided $|v_1|$ and $|v_2|$ are large enough. Obviously, nothing depended on the choice of voters 1 and 2 in particular, so we conclude that every majority has control over the incumbent's choice of effort levels.

Chapter 7

On a normative theory of democracy

One desideratum for a concept of popular control is that it be useful for a normative theory of democracy, by which I mean a theory of the reasons that democracy, or certain kinds of democracy, could be better or worse than the alternatives. To be useful for this purpose, a concept of popular control should pick out features of democracy that are relevant to these comparisons.

In this chapter, I lay some groundwork for the theory that democracy's tendency to give all majorities' an adequate and comparable degree of control is one of the primary reasons it is preferable to alternative types of political regime. The justification for that theory would involve familiar kinds of arguments, appealing to the idea that citizens' legitimate interests are more secure when they have some control over their leaders, and the idea that a system of democratic, popular control contributes to relations of social equality among citizens. The full elaboration and defense of such a theory is far too ambitious a project for this chapter. My more modest goal is to soften the reader's resistance to the idea that the concept of popular control, as I have analyzed it, could play this role in a normative theory of

democracy. I will examine two sources of resistance, in particular, and then point out some implications of giving all majorities' control that may help to make the theory more attractive.

One likely source of resistance is the intuition that in a democracy, decisions are not supposed to reflect the intensity of citizens' preferences. Indeed, democratic theorists often reject the idea that democratic decisions should reflect even citizens' "mere" *preferences*, as opposed to their deliberative *judgments*. Yet, on my account, it may seem we have to credit preferences and their intensity with normative significance that they do not actually have. I explain in the next section why these are not good reasons to resist the normative theory I am proposing.

A second likely source of resistance is the thought that majorities' control represents too minimal an achievement to be a plausible reason for preferring democratic regimes to the alternatives. After all, even if all majorities have "adequate" control (§5.4), it is possible that some majorities have merely adequate control while others have much greater control. In this chapter, I articulate a sense in which two groups can be said to have equal degrees of control and use it to formulate a requirement that fully democratic systems of popular control must satisfy.

Having addressed these two sources of resistance to the theory, I show that the concept of rule by multiple majorities also changes the way one thinks about the status of minorities in democracy. For example, I explain why disenfranchising a minority is incompatible with giving all majorities control, and, building on an argument by McGann (2006), explain why giving majorities control can be a means of protecting minorities' rights and

interests. These conclusions differ markedly from the conclusions to which standard intuitions about democracy and majority rule lead one. They must be born in mind when one considers whether a regime's tendency to establish democratic, popular control is one of the reasons to value it.

§7.1 Preferences and their intensity

On the view I have defended, whether a group has control depends on whether and how outcomes are sensitive to the strength of the group members' preferences, and I have suggested that a democracy's tendency to produce popular control could be reason it is superior to alternatives. A skeptic might object that it is a mistake to treat facts about preference intensity as normatively significant. As Shapiro (2016) says, "there are good reasons why democracies count how many people want something rather than how much they want it" (p. 48).

The amount of passion with which Southern whites wanted to preserve Jim Crow, or how much utility they derived from its existence, is beside the point; it was an oppressive social order.
(Shapiro, 2016, p. 57)

I expect many readers will agree with Shapiro's observation. Since the observation could easily be mistaken for a reason to reject the account of popular control I have defended, let me explain why it is in fact perfectly compatible with that account.

Shapiro's example is imperfect rhetorically because it seems unlikely that Jim Crow policies would have survived if Southern states had counted how

intensely Southern blacks preferred the abolition of those policies. But even supposing that “counting how much people want something” would have led to the adoption of such policies, the example works against Shapiro’s purposes. *How many* Southern whites wanted to preserve Jim Crow seems to be just as much “beside the point” as *how intensely* they wanted to preserve it. Both kinds of facts are clearly beside the point in one sense: when it comes to matters of basic justice, neither the number of people preferring a policy nor the intensity of their preferences is relevant to the determination of whether the policy is just (Rawls, 1971, p. 230). One might argue further that political institutions should be designed, as far as possible, so that the protection of certain basic rights does not depend on *either* the mere number of people who wish to violate these rights *or* the intensity with which they hold their attitudes. Neither claim stands in necessary tension with the account of popular control I have given. In particular, the latter claim just implies that decisions on certain policies should not be under popular control or, indeed, *anyone’s* control.

As a cleaner test case, let us take a situation where the policy in question arguably does not concern matters of basic justice or rights. Take the decision whether to require background checks on all guns sales. As Shapiro notes, large majorities of Americans consistently support universal background checks.¹ Let us stipulate, for the sake of the argument, that

¹Tom Kertscher, “Do 90% of Americans support background checks for all gun sales?” *Politifact*, October 3, 2017, <http://http://www.politifact.com/wisconsin/statements/2017/oct/03/chris-abele/do-90-americans-support-background-checks-all-gun-/>.

the explanation for why Congress nevertheless fails to pass such measures is that the small minority of citizens who oppose the measures hold their preferences much more intensely than the large majority who support them. And let us also grant that under this description, the case of gun control legislation is a democratic failure, a deviation from how the democratic process ought to work.² Let us grant, in particular, that it is a case in which, intuitively speaking, policy fails to be under popular, democratic control. The question is whether the account of popular control I have defended is compatible with those intuitive judgments, even though, on that account, popular control is partly a matter of whether and how outcomes respond to the intensity of preferences.

On the account I have given, popular control requires that majorities have control. As I noted in chapter 5, that concept of popular control differs from the claim, associated with theories of pluralism, that the distinctive feature of democracies (or “polyarchies”) is that they establish something like “government by *minorities*” (Dahl, 1956, pp. 133, 134). Popular control, as I have characterized it, does not require that minorities have control, so it does not require that the intense preferences of a minority prevail over those of a majority. Thus, it does not require that the minority of gun enthusiasts be allowed to prevail over the supermajority favoring gun control.

It is true that on my account, we might still want to say of the gun control case that the supermajority retains *some* control over gun control

²For an alternative view, defending a voting scheme in which the intense preferences of minorities are allowed to outweigh the weaker preferences of majorities, see Posner and Weyl (2018), chapter 2.

legislation, even though they do not get their way. The supermajority could have some control if it is the case that the gun control measures would pass, provided the members of the supermajority cared enough about the issue. But that judgment should not strike one as counterintuitive or objectionable so long as one keeps in mind that control comes in degrees. Our intuitions about the gun control case may be naturally captured by the judgment that the supermajority ought to have more control over gun control legislation than it actually does. Our intuition may be that in a better functioning democracy, the members of the supermajority would not need to hold their preferences any more intensely than they actually do, in order for their preferences to prevail over those of the minority.

Indeed, we may think that in a well-functioning democracy, the majority's preferences over the gun control measure ought to prevail no matter how intensely they are held. On the view I have defended, that would be equivalent to the judgment that the majority ought to have maximal control. In other words, in the case of a binary decision, a procedure like simple majority rule, under which a majority's preference is decisive irrespective of how strongly it is held, maximizes majorities' control. (The reader may wish to revisit the explanation, from chapter 5, for why quadratic voting, which *would* allow a minority with sufficiently intense preferences to defeat a majority, does not maximize majorities' control.)

If that is a reasonable judgment about the gun control case, why not say that popular control always requires giving majorities maximal control, such that their preferences always prevail irrespective of how intensely they are held? The answer is that popular control would then be impossible in

situations where no Condorcet winner exists. That view of popular control is vulnerable to the impossibility objection, from chapter 3—it entails the critical premise about the concept of popular control that figures in the skeptic’s argument for its impossibility.

Apart from the misplaced worry about the role of intense minorities, a critic might contend that my account of popular control errs in giving a fundamental role to preferences, as opposed to judgments about the common good or justice. (The additional role given to preference intensity only compounds the error, the critic is likely to think.) In epistemic theories of democracy (Cohen, 1986; Estlund, 2008; Landemore, 2012), for example, democratic citizens (at least in an ideal case) deliberate about what is just or advances the common good, and when they vote, their “voting expresses beliefs about what the correct policies are according to the independent standard [e.g., of justice or the common good], not personal preferences for policies” (Cohen, 1986, p. 34). One might worry that the focus on preferences in my account of popular control means that the concept could not play any useful role in an epistemic theory of democracy.

Let me reiterate the operative definitions of preference and preference intensity. The terms are not to be construed according to their connotations in ordinary language; they are technical terms of art. A person is said to prefer one option to another if she would choose the first over the second, given a choice between the two (§2.2). The intensity of this preference is measured by what she would be willing to give up, with respect to other variables she cares about, in order to secure the choice of the first option over the second (§4.2). If someone deliberates with her fellow citizens about

the requirements of justice and, in light of that deliberation, would choose a higher minimum wage over the existing one given a choice, then she “prefers” a higher minimum wage. If, in one scenario, she would be willing to march in the streets in order to secure the passage of the higher wage, while, in another, she would not be willing to give up anything in order to secure the passage of the higher minimum wage, then (all else being equal between the two scenarios) she holds her preference “more intensely” in the first.

Given how the terms are to be understood, it should be clear that a regime in which citizens’ preferences and their intensity influence decisions could still be a regime in which citizens deliberate and vote in the manner that epistemic democrats envision. Consistent with epistemic democrats’ assumptions about such deliberation and voting behavior, it may be the case that decisions are influenced by facts about which of two decisions a citizen would choose given the choice, and facts about what a citizen would be prepared to sacrifice in order to secure those choices. The account of popular control may raise new questions about whether a democratic regime of popular control has the epistemic properties that epistemic democrats have described in the context of their simpler models of democracy, which restrict attention to binary decisions. But there is no obvious reason why it could not play a role in an epistemic theory of democracy.

§7.2 Equal control

The second source of concern with my account of popular control is that it requires too little, because the conditions for a majority’s control are too

minimal. Even if every majority has “adequate” control, as defined in §5.4, there may be vast inequalities in the control that different majorities, or majorities and minorities, enjoy. So, without a further requirement on what counts as a democratic system of popular control, one may wonder whether the concept of popular control can play the role in a normative theory of democracy that I have suggested it can. To address this concern, I propose in this section an additional requirement for a “fully democratic” system of popular control. The requirement concerns equality in the distribution of control. In addition to making the ideal of democratic popular control more demanding, it represents an original interpretation of the kind of political equality that democracy requires.

A society and its institutions are not fully democratic unless there is the right kind of political equality. What kind of political equality is necessary?

In empirical studies of representation and the effects of public opinion on policymaking, it is commonly assumed that “a key characteristic of a democracy is the continued responsiveness of the government to the preferences of its citizens, considered as political equals” (Dahl, 1971, p. 1). Dahl’s claim is the point of departure for a recent wave of scholarship on unequal representation (Jacobs and Page, 2005; Gilens, 2012; Gilens and Page, 2014; Bartels, 2016). Political theorists sometimes defend principles in the same general vicinity as Dahl’s, although they more often emphasize equality of *opportunities* for political influence rather than equality of influence per se (Rawls, 1971; Christiano, 1996; Knight and Johnson, 2011; Pettit, 2012; Kolodny, 2014; Saunders-Hastings, 2018).

Alternatively, one might reasonably take the view that the political

equality that democracy requires is not equality of influence or opportunity for influence per se, but rather equality in the distribution of *control*.³ One reason to shift our focus from influence, or responsiveness, to control is that influence over political decisions is far too broad a category, comprising phenomena that there is little reason to care about and that have little connection with democracy. As Pettit (2008) points out, “people would have equal, non-zero influence on government, if government was conducted, perversely, with a view to frustrating anything they equally supported” (p. 47). Influence could be haphazard. Since mere influence could take the form of purposeless or even perverse and counterproductive influence, as in Pettit’s example, it is not clear why there should be any connection between equality of generic influence, or equal opportunities for generic influence, on the one hand, and moral goods like democratic legitimacy, justice, or social equality, on the other. In light of these observations, we might try capturing our intuitions about political equality with principles that regulate the distribution of control, rather than mere influence or opportunities for influence.⁴

³Some writers have taken this approach. Viehoff (2014), for example, argues that individuals relate to each other as equals only when they have “equal control over [their] relationship,” and that egalitarian procedures for deciding on the rights and regulations structuring citizens’ relationship are therefore necessary for social equality. He does not give a detailed account of what equal control requires, however.

⁴This is Pettit’s (2012) thesis in the case of political legitimacy and non-domination: legitimacy requires non-domination, which requires a system of popular influence resulting in equal control over the state. On Pettit’s

In Pettit’s hypothetical example, citizens have equal influence but they do not have control over what the government does.

In what sense might democracy require an equal distribution of control? There are two questions to consider here. First, what does it mean for a distribution of control over something to be equal? Second, over *what* should the distribution of control be equal in a democracy—political decisions, selection of leadership, policy outcomes, the “aims” of elected representatives? I give an answer to the first question, but bracket the second. For the sake of concreteness, I will refer to control over political decisions, but it may be that the principle of equal control I go on to formulate would work better as a constraint on the distribution of control over other variables besides political decisions.

One might first try the idea that democracy requires all citizens to have equal control. As a rule, actual democracies give most citizens no control over political decisions, however.⁵ Nor would it be desirable to give each individual some rather than no control over political decisions, because then the extremely intense preferences of a single person could prevail over the preferences of an overwhelming majority. Even in the case of an idealized direct democracy, the only reason we would expect everyone to have equal control is that we would expect everyone, equally, to have no control over view, a kind of equality in influence is still necessary, but it is equality in a particular kind of influence.

⁵Brennan (2016, ch. 4) and Dworkin (2000, p. 191) make essentially the same point, albeit couched as a claim about “power” rather than control as I have defined it.

political decisions.

It may be that democracy still requires that all individuals have equal control, even if the only democratically acceptable way of equalizing control is for everyone to have no control. Perhaps the only alternative is inequalities in control that are unacceptable in a democracy. However, even if equality across individuals is necessary, it is not interesting or important in its own right. It does not illuminate a connection between democracy and any of the moral goods associated with it. Prior to women's suffrage, the typical individual man and the typical individual woman equally had no control over political decisions, even though one and not the other enjoyed the right to vote. More generally, the point is that even in hierarchical societies almost everyone, equally, lacks control over political decisions. The exceptions are political elites close to the levers of power, but their counterparts in democracies—even idealized representative democracies—will also generally have more control over political decisions than anyone else.

A more promising thought is that a principle of democratic equality requires equality in the distribution of control across *groups*, not just individuals. We already saw one minimal version of this idea in chapter 5. There, in the reasoning leading to the multiple majorities view of popular control, it was assumed that the distribution of control across different groups must be minimally egalitarian, meaning that if one group has control over a variable, then any other group that is the same size or larger should also have control over it, all else being equal. (Remember that here and throughout the book, the default meaning of *groups* is just sets of individuals, not particular sets defined by social categories like ethnicity, class, gender, etc.) If groups whose

members all lack the right to vote would lack control in a democracy, then this minimal principle would appropriately register their disenfranchisement as a breach of democratic equality.

The minimal version of the principle would still fail to pick up other deviations from political equality, however. Onerous voting identification requirements might raise the costs to certain individuals of exercising their voting rights. Groups of these affected citizens would still have *some* control over political decisions if they were prepared to exercise their voting rights, despite the cost, whenever they attached sufficient importance to political decisions, and if the effect of being so prepared was that their preferences constrained decisions whenever they attached sufficient importance to them. But intuitively speaking, their control would not be comparable in its significance to the control of similarly sized groups whose members were not affected by the identification requirements. The minimal version of the principle would let slide even worse offenses. For even in the case of outright, *de jure* disenfranchisement, the disenfranchised could still have some control over decisions. They might still have means of influencing policymakers—taking to the streets, going on strikes, and so on—and if they would choose to use these means when they cared strongly enough, then their preferences might constrain policymakers when held with sufficient intensity. If so, they would have control, just as the enfranchised have control. But intuitively speaking, there is a difference in the degree or significance of the two groups' control, owing to the inequality in voting rights. A principle of equal control should treat that difference in the practical significance of the two groups' control as a problem.

Here is a stronger principle that captures these intuitive judgments: in a democracy, if two groups are of the same size, then they ought to have *equally significant* degrees of control over a political decision. Groups with equally significant degrees of control will be said to have *equal control*, for brevity's sake. A principle of democratic equality would then still require all individuals to have equal control—each individual corresponds to a group of one—but it would also require more. Similarly sized groups of men and women did not have equal degrees of control over political decisions prior to women's suffrage, violating the principle. In fact, the principle may seem *too* strong, as it requires that, say, the country's chief executive and an ordinary citizen must each have the same degree of control over political decisions. I consider amendments below. There is a prior question to answer about what it means for two equally sized groups to have equally significant degrees of control. Let me start off with an example to prime some intuitions.

Imagine an elected legislator who is conflicted about a piece of upcoming legislation and announces to his constituents that he will hold a town hall meeting to hear their concerns and opinions. To keep things simple, suppose he has just three constituents, *A*, *B*, and *C*. The legislator plans to vote against the legislation unless at least two constituents show up at the meeting and express their support for the legislation. *A* lives one mile from the location of the town hall meeting, while *B* and *C* each live 200 miles away.

Assume that in actual fact, all three constituents are indifferent as to whether the legislation passes, and none of them enjoys participating in a

town hall meeting for its own sake, so they all choose to stay home.⁶

In this scenario, does the majority comprising B and C have control over the legislator's action? To answer the question, we consider what would happen in counterfactual scenarios where B and C felt differently.⁷ Suppose first that B and C each preferred for the legislation to fail. Then, in the unique equilibrium to that counterfactual game, neither would attend the meeting and the legislation would fail. Now suppose that B and C each preferred for the legislation to pass, and they would each be willing to incur the costs involved in traveling 400 miles, round trip, and participating in the meeting, in order to secure passage of the legislation.⁸ In one equilibrium to that

⁶Recall two points from our analysis in chapter 4. First, a group can have control over a decision even if, in actual fact, its members are indifferent (§4.5). Second, and relatedly, a judgment about an agent's control is always a judgment about their control "in equilibrium," that is, their control in a particular situation, with all of the relevant agents acting in the way one would expect, given their actual preferences, beliefs, and so on (§4.2).

⁷That is, we consider what happens in the equilibria to counterfactual games corresponding to alternative specifications of B and C 's preferences, holding A 's preferences fixed. For readers who like their models expressed formally, this game effectively has three players (the legislator is like an automaton, game-theoretically), A , B , and C , and each player $k \in \{A, B, C\}$ simultaneously chooses whether to attend, $a_k \in \{0, 1\}$. Player k gets a payoff of $v_k - c_k a_k$ if the legislation passes, and $-c_k a_k$ otherwise, where $c_A = 1$ and $c_B = c_C = 200$. In the actual game, $v_k = 0$ for each $k \in \{A, B, C\}$.

⁸That is, consider the counterfactual game in which $v_B, v_C > 200$, while

counterfactual game, they would each choose to attend and the legislation would pass. Assume that is how they would act. There is also an equilibrium in which neither attends, but let us suppose, for now, that B and C will successfully coordinate on the Pareto optimal equilibrium, perhaps because they have means, not modeled here, of solving the coordination problem, or perhaps because the Pareto optimal equilibrium is a “focal point.” Then the majority comprising B and C has control over the legislation.

What about the majority comprising A and B ? They also have control. Suppose that A and B each preferred for the legislation to fail. Then neither would attend the meeting and it would fail. Suppose that they each preferred for it to pass, and, in order to secure the passage of the legislation, A would be willing to incur the costs of driving two miles, round trip, plus the time spent in the meeting, and B would be willing to incur the costs of driving 400 miles, round trip, plus the time spent in the meeting.⁹ Then, in one equilibrium to that game, they would each choose to attend and it would pass. (As before, there is a Pareto suboptimal equilibrium in which neither attends.) Assuming that is how they would act in that counterfactual scenario, then they also have control over the legislation.

Each majority has control over the legislation, but there is an intuitive sense in which the majority comprising B and C , who must each travel 200 miles in order to influence the legislator, has a less significant degree of control than the majority comprising A and B , one of whom need only $v_A = 0$ as before.

⁹That is, consider the counterfactual game in which $v_A > 1$ and $v_B > 200$, while $v_C = 0$ as in the original, “actual” version of the game.

travel one mile to influence the legislator. *B* and *C*'s shared preferences constrain the outcome only if they are *each* willing to give up the time and resources involved in traveling 400 miles (plus attending the meeting), while *A* and *B*'s shared preferences constrain the outcome if they are each willing to make that sacrifice, but they will also constrain the decision even if *A* is unwilling to make *that* sacrifice, so long as *B* is willing to do so and *A* is willing to make the smaller sacrifice involved in traveling two miles (plus attending the meeting). It is more costly, in the aggregate, for *B* and *C* to exercise their control than for *A* and *B* to do so. One majority need not hold its preferences over the legislation as "intensely" as the other, in order for its shared preferences to constrain the decision. What the members of one majority must be willing to give up, in the aggregate, if their preferences are to have any effect on the decision, is more practically significant than what the members of the other majority must be willing to give up, in the aggregate. We might record this fact by saying that although the two overlapping majorities both have control over the legislation, the majority comprising *B* and *C* has a less significant degree of control than the majority comprising *A* and *B*.

Interpersonal comparisons of "preference intensity" may make some readers uneasy, so it is worth emphasizing that we are not pretending here to make interpersonal comparisons of unobservable mental states or quantities of satisfaction. We are making comparisons of the practical significance of what different people must be willing to give up if their preferences are to constrain a decision. Having to forfeit the time involved in 400 miles of travel is a more significant cost than having to forfeit the time involved in

two miles of travel. It is more significant in the sense that it is more likely to be a deterrent to the exercise of control.

To make sure the reader has the relevant intuition here, let me suggest an even simpler example. An item is being auctioned off at a starting price of \$50,000, and there are two potential bidders, *A* and *B*. In actual fact, neither is willing to pay the asking price. But it is true of each of them that if he preferred to have the item and were willing to pay the asking price, then he would get it, given that the other person is in fact unwilling to pay the asking price. Thus, each has some control over the outcome of the auction. But suppose *A* is a billionaire, while *B* would have to exhaust all of his savings and take out loans in order to pay the asking price. Who has a more significant degree of control over the outcome? Surely *A*. What fact are we trying to express when we make that judgment? The price is likely to be a greater deterrent to the exercise of *B*'s control than *A*'s control. The range of "psychological types" who, in *A*'s position, would acquire the item is larger than the range of types who, in *B*'s position, would do so.

This is admittedly still a bit vague, but I will continue on the assumption that the idea is clear enough for present purposes. My goal is merely to persuade the reader that the concept of popular control can be made suitably more demanding if we add the further requirement that the distribution of control across similarly sized groups be approximately equal in its practical significance. The reader can form a reasonable opinion about that question even if there remain outstanding questions about how exactly to flesh out the notion of equal practical significance.

Now, the principle of equal control as stated is no doubt too strong a

principle. For one thing, a group will have more or less significant control, depending on contextual factors that vary from decision to decision. Relevant factors include the preferences that other people, not in the group, happen to have—recall that the preference condition for control is assessed “in equilibrium,” holding fixed the strategies and preferences of other agents. Depending on the mechanisms by which decisions are made, it may be that the more intense the “opposition,” the less control a group has. For another, the principle would rule out the inequality in control between an elected representative and one of his constituents, as each individual makes up a similarly sized group of one. In response to those observations, we might make two further amendments so that the principle runs: in a democracy, there are no *persistent, durable* inequalities in the practical significance of equally sized groups’ control over political decisions, *apart* from those inequalities arising from democratically authorized delegations of decision-making. When there are no persistent inequalities in the control of similarly sized groups, and all majorities have control, I will say that majorities have *comparable* control.

The rider is of course an ad hoc fix to the apparent counterexample represented by political inequality between public officials and ordinary citizens. That inequality is sometimes treated as a reason to reject the view that democracy requires equality of power (Dworkin, 2000; Landa and Pevnick, N.d.). I will not try to meet this objection here. Ideally, one would specify the category of exceptions to equal control in a less ad hoc fashion, appealing to the underlying democratic values that explain why there are such exceptions. That is, one would specify the rider by appeal to the values that explain why it makes perfect sense to say, on the one hand, that per-

sistent and durable inequalities between a career politician and an ordinary citizen are compatible with a democratic system of popular control, but, on the other, persistent inequalities in control over decisions between men and women are not. I will proceed on the assumption that this can be done. I am mainly concerned here to convince the reader that a suitably demanding ideal of democratic popular control can be gotten from adding the requirement of equal control to the requirement that all majorities have control. Meeting the objection that the resulting ideal is *too* demanding is a future task.

To that end, consider the judgments that our amended principle of equal control supports about standard cases of political inequality. Under plausible assumptions about the mechanisms by which groups could come to have control over political decisions in a representative democracy, this principle will be violated when some citizens are denied the right to vote; or when some citizens' votes count for more than other citizens', as in Mill's plural voting scheme; or, indeed, a system like the U.S. electoral college, in which majorities comprising voters from small states, which are "overrepresented" in the electoral college, can be expected to have more control than similarly sized majorities comprising voters from large states, which are underrepresented. If some citizens could cast more votes than others, or their votes are more consequential because of a system like the electoral college, then elected officials and political parties will give more weight to groups comprising those privileged voters than similarly sized groups without privileged voters, all else being equal. Majorities comprising the privileged voters will not need to hold their preferences over policy decisions with as much in-

tensity, in order for those preferences to constrain the choices of the elected officials and political parties.

Formal inequalities in political rights are not the only source of persistent, durable inequalities in control. In our toy example from above, the control that a majority had over the legislator's vote depended on a majority's ability to coordinate on a Pareto optimal equilibrium. When *A* and *B* both prefer for the legislation to pass, and each would be willing to incur the costs of attending the meeting in order to secure passage of the legislation, there are two equilibria. There is one in which they both attend the legislation passes, and one in which neither attends and the legislation fails. The assumption that they had control depended on the assumption that in such scenarios, they would successfully coordinate their actions and attend the meeting. Generalizing from the toy model, one source of durable, persistent inequalities in groups' control could be durable, persistent differences in their abilities to overcome collective action problems. Such inequalities are of course compatible with formal equality in political rights.

With this potential source of inequality in mind, consider the judgments that the principle of equal control would support in the case of the U.S., in light of the evidence of unequal responsiveness to public opinion. Gilens (2012) reports evidence that in the U.S., policy changes are responsive to the preferences of the affluent, but not low- or middle-income citizens, when the preferences of the different groups diverge.¹⁰ Is that pattern of unequal responsiveness evidence of unequal control? It depends on the mechanism

¹⁰See also Jacobs and Page (2005), Bartels (2016), Gilens and Page (2014).

that produces it. Policy could be unequally responsive to the preferences of two groups that have the same degree of control if the intensity of one group's preferences exceed the necessary threshold for influence, while the intensity of the other group's preferences fall short of a comparable threshold. To his credit, Gilens considers this explanation of the observed pattern, and notes that in surveys that ask respondents "How important is this [policy] issue to you personally?", "low-, middle-, and high-income respondents expressed nearly identical levels of importance," averaging across policy issues (Gilens, 2012, pp. 90, 91). The survey evidence is suggestive, but a skeptic might doubt the validity of respondents' self-assessed preference intensity. And even if the self-assessments track facts about what respondents would be willing to give up in order to get their preferred policies, the survey evidence that Gilens describes leaves open the possibility that on those issues where preferences of the rich and poor diverge—and where the evidence of unequal responsiveness emerges—the preferences of the rich are held more strongly than the preferences of the poor, even if no differences in intensity exist on average. I am not suggesting that this is a plausible hypothesis, but merely pointing out the limits of the survey evidence on its own. We need to supplement that evidence with a theory of the mechanism that explains unequal responsiveness.

One plausible hypothesis is that unequal responsiveness results from unequal political participation. Wealthier citizens vote at higher rates and give more money on average than poorer citizens (Verba, Schlozman and Brady, 1995; Gilens, 2012, figure 8.1). Let us assume this explanation is correct,

for the sake of argument, although there are of course other hypotheses.¹¹ My main purpose here is just to illustrate what the principle of equal control requires, not to arrive at a firm conclusion about whether American democracy satisfies it. For that purpose, it is instructive to examine the implications of the principle under the assumption that unequal responsiveness results from unequal political participation.

A naive explanation for differences in participation would be that wealthier citizens care more about the outcome of the democratic process than poorer citizens. Differences in the importance that the two groups tend to attach to the outcomes of the political process would explain why one group participates at higher rates than the other, and why one group's shared preferences tend to prevail over the other group's shared preferences. That explanation would be compatible with its being the case that the intensity with which groups of wealthy citizens must hold their preferences, in order for their preferences to constrain what policymakers do, is comparable to the intensity with which poor citizens must hold their preferences in order

¹¹Another hypothesis is that imperfections in electoral competition and the principal-agent relationship between voters and elected officials leave the latter with slack to act on their personal policy preferences. Elected representatives have personal policy preferences that reflect their more affluent class backgrounds (Carnes, 2013). See also Butler (2014), presenting evidence, from field experiments, that elected representatives are unequally responsive to constituents from different socioeconomic backgrounds, even when constituents participate in the same way and at the same level.

to achieve comparable effects. Thus, on the naive explanation of unequal participation, we would not be justified in inferring inequalities in control from the inequalities in responsiveness.

However, that explanation is not convincing. Each individual has a vanishingly small probability of affecting the outcome through her participation, and the importance she attaches to the outcome is, *on its own*, insufficient to explain why a minimally rational agent would choose to participate. It may play *a* role in conjunction with other factors. Perhaps mechanisms of social pressure—which have been shown in field experiments to affect voter turnout (Gerber, Green and Larimer, 2008)—are more effective when they target someone who is publicly known to have intense preferences over the outcome of the political process. The perception of being a free-rider is perhaps more socially costly for those people. I only toss out such possibilities to illustrate an idea. The importance someone attaches to the outcome of the political process may play *some* role in explaining their participation, but only in conjunction with other factors (such as social networks that pressure one into compliance with the certain norms), and these other factors will likely be unequally present among rich and poor. Thus, one cannot naively infer differences in the importance that rich and poor attach to the outcome of the political process from differences in participation rates. Moreover, if these other participation-inducing factors—different social networks that apply different norms, or differences in how much someone cares about violating a social norm relative to the costs of turning out to vote—are distributed unequally between rich and poor, then participation rates will be unequal even when rich and poor care about the outcome with equal

intensity.

Given the probable connection between political participation and control, the implication is likely to be that the shared preferences of an affluent majority (or even just an affluent minority) will constrain political decisions, but the shared preferences of a similarly sized group of less affluent citizens will not, even if everyone holds their preferences with comparable intensity. The distribution of control across groups will be unequal in virtue of inequalities in those factors, besides strength of preferences over the outcome of the political process, that explain participation. Or, to put it another way, unequal control will result from differences in groups' abilities to solve collective action problems.

This conclusion illustrates, I think, one way in which the concept of equal control may lead us to think differently about the implications of a commitment to political equality, compared with the usual approach that interprets that commitment as a constraint on the distribution of *opportunities* for political influence. As the previous line of speculation about the empirical research on unequal responsiveness shows, the principle of equal control might plausibly justify a measure like compulsory voting. That would erase the collective action problem that groups of voters face on election day, when they choose whether to turn out to vote, and it would remove one likely source of inequality in the control of rich and poor. I am not claiming that such a measure is justified all things considered; I am merely pointing out that whatever the other reasons for or against it, it would plausibly remove one source of inequality in control across groups. That conclusion is different from the view that one would presumably come to on the basis of a

commitment to equal opportunities for political influence. Forcing someone to vote deprives someone of an ability to refrain from exercising the opportunity for political influence they already had. It does not give them a new *opportunity* for influence.

§7.3 The status of minorities

What difference does it make whether we adopt the concept of popular control as rule by multiple majorities or, instead, a variation on the traditional view that identifies it with the rule of “the” majority? While I criticized the traditional view for its lack of realism in chapter 2, I now show that its defect is also consequential for normative democratic theory. It prevents one from thinking clearly about the connection between popular control and the place of minorities in a democracy.

The traditional view pictures democracy as a game with two players. Democratic regimes, at least those of a majoritarian cast, are said to empower “the” majority to rule and to dominate “the” minority. “What is a majority, in its collective capacity,” de Tocqueville (2000, p. 251) asks, “if not an individual with opinions, and usually with interests, contrary to those of another individual, called the minority?”

Sometimes the best response to a rhetorical question like Tocqueville’s is to give the uninvited answer. What is a majority? A set of citizens who outnumber those citizens not included in the set. Each member of a majority also belongs to multiple distinct majorities. That is not an empirically testable statement about cross-cutting social cleavages and the multiple di-

mensions of any person's identity and interests. It is a trivial fact about the concept *majority of the group*, as defined here: for any member of a group—a group of citizens, a group of books, of anything—the group member belongs to multiple, distinct subsets of the group that each constitute numerical majorities of the whole.

If there were just two collective actors on the democratic stage, then empowering one would be compatible with disempowering the other. One could empower “the majority” to rule even if one disempowered “the minority.” In particular, denying members of the minority the right to vote would not be expected to diminish the majority's control over collective decisions if one subscribed to the standard view, illustrated by Tocqueville's rhetorical question.

However, democracy is not a contest between two collective actors, the majority and the minority. There are multiple, overlapping subsets of the population that constitute numerical majorities of the whole. If they are *all* to have adequate control over a political decision (or elected officials or some other outcome), then *no* minority can be disenfranchised. Not, anyway, if the right of citizens to vote is the means by which groups of citizens acquire adequate control over political decisions (or elected officials and so on on). When a minority is deprived of its right to vote, the effect is that various majorities of the *citizen* population are no longer majorities of the *enfranchised* population. For example, if ten percent of the citizen population is disenfranchised, then there is a group comprising 51 percent of the citizen population but whose enfranchised members make up only $\frac{41}{90}$, or about 46 percent, of the electorate. This majority of the citizen population

will presumably lack adequate control over decisions—and it will certainly have less control than other, similarly sized majorities—because it is no longer a majority of the electorate. In other words, if a majority enjoys an adequate degree of control only in virtue of the right to vote, then the effect of disenfranchising a minority is that various majorities will have inadequate control.

This observation points us toward a reason for supporting constraints on majorities that is inaccessible to anyone who is in the grip of the standard view. Suppose that one does not believe that democratic rights have any intrinsic value, but one believes that it would be desirable for instrumental reasons if certain decisions are under democratic, popular control, specifically the equal and adequate control of all majorities, and one is wondering how to institutionalize a system of rights and decision-making procedures that will reliably provide all majorities with equal and adequate degree of control over these decisions. If a majority can only hope to have adequate control if its members enjoy the full suite of democratic rights of participation, then one should support constraints on majorities that prevent them from infringing on minorities' democratic rights. Protecting the democratic rights of minorities is a means of ensuring all majorities have the desired control. If members of minorities are not able to participate fully, then the various bare majorities to which they belong will not have adequate control. The point is a simple one, but obscured from view if one misdescribes democracy in terms of Tocqueville's two group agents. In that imaginary world, disempowering "the minority" would not undercut the power of "the majority."

Notice that this argument does not depend on any assumption about the distribution of preferences in the population, so it is sound even if majority preference cycles never occur. Compared with the traditional alternatives, the multiple majorities concept of popular control is a superior tool for thinking about democracy even when we bracket the challenges from social choice theory.

Disabusing oneself of the singular majority fallacy brings other insights into the status of minorities in a regime of popular democratic control. It puts one in a position to think more carefully about the connection between majorities' control and the security of minorities' rights and interests more broadly (not just their democratic rights narrowly construed). The argument I have in mind builds on McGann's (2006) argument that "the presence of multiple, cycling majorities [...] provides the possibility of a check on majorities without artificially empowering a minority" (McGann, 2006, p. 109). He gives two reasons:

First, as Miller (1983) argues, [cycling] can ensure that there are no permanent losers. A group out of power can always expect to be able to defeat the incumbents in the future and thus have an incentive to keep playing the game, which enhances the stability of the system. Second, [because] any winning [majority] coalition can be split, the minority always has a means of retaliation. [...]

If the winning coalition is too harsh to the other players, then these players may approach part of the winning coalition and sell their support for a very low price. (McGann, 2006, p. 109)

As formulated, the claims about the consequences of cycles are too strong. Recall our distinction, in §3.2, between majority *voting* cycles and majority *preference* cycles. A majority voting cycle refers to a dynamic process in which first one majority adopts one policy, then another defeats it and replaces it with a new policy, and so on, in a cycle. A majority preference cycle refers to a property of certain configurations of individuals' preferences. That a majority preference cycle occurs does not, by itself, imply anything about how anyone will behave. In particular, it may be that majority preference cycles occur all the time, yet the outcomes of majoritarian decision-making processes are stable—so that some people are “permanent losers” in the sense of suffering badly under the stable outcome. Shepsle's (1979) model of “structure-induced equilibrium,” Baron and Ferejohn's (1989) model of majority rule bargaining in legislatures, models of electoral competition with probabilistic voting—they are all stories about majoritarian decision-making processes that produce “stable” outcomes in equilibrium, despite the presence of majority preference cycles. The outcomes are stable in the sense that if the underlying preferences, beliefs, and other elements of the strategic environment remain unchanged, so too does the outcome. Moreover, minorities can fare quite poorly in some of these stories. In Baron and Ferejohn's model of majority rule bargaining, a majority divides the available resources amongst its members and leaves nothing for the minority. In the hypothetical worlds described by the models, minorities do not have “means of retaliation,” despite the presence of majority preference cycles.

Whatever the limitations of those models, they demonstrate that the

claims about minorities do not follow from the assumption of majority preference cycles. The claims may be true. But just because a majority preference cycle occurs and a group of actors use “majority rule” to make a decision, we cannot infer that there will be no permanent losers or that minorities will have means of breaking up majority coalitions. To quote a familiar slogan from the literature in the 1980s and ’90s that emerged from the engagement with social choice theory, “institutions matter.” Whether the claims are true will depend not just on the presence of majority preference cycles, but also on the particular institutional embodiment of “majority rule.” One lesson of social choice theory and the research it inspired is that “there are, in fact, *many majority rules*” (Shepsle, 1986, p. 54).

Some institutional embodiments may not give all majorities control. If institutions concentrate agenda-setting power, then majorities not comprising the agenda-setters may not have control, even if proposed legislation is decided by majority voting. More generally, if a minority is faring badly under the status quo, but the majorities that prefer alternatives to the status quo lack control, then the intuitive basis for McGann’s predictions is missing, even if the institutions in place instantiate a version of “majority rule.” If those majorities lack control, then no matter how strongly they care about the outcome, their shared preferences will not constrain the choice between the status quo and those alternatives. (That is what it means to lack control.) The mere fact that those majorities prefer alternatives to the status quo would then be no reason to expect any relief for the minority that is faring badly under the status quo. If McGann’s claims are plausible, we should therefore investigate them the context of those institutional embodiments of

“majority rule” that give all majorities control.

Even with the assumption that all majorities have control, we still cannot infer protection of minorities’ interests from the presence of a majority preference cycle. That a majority comprising a vulnerable minority has control means that if *every* member of that majority felt sufficiently intensely, then their shared preference for an alternative to the status quo would bring about a change. But no matter how badly the minority is faring under the status quo, and no matter how intensely the minority prefers the alternative, it may get no relief if the other members of the “latent” majority coalition do not also feel strongly enough. This is just to repeat the point made above in the discussion of Shapiro’s worry about intense minorities: every majority’s having control does not imply that the preferences of intense minorities will be decisive.

We can make a more tentative and qualified claim, however. Suppose that a group is considering several policy proposals, including a perpetuation of the status quo. Suppose that among the options, x is the worst alternative for each member of some vulnerable minority of the group. If no Condorcet winner exists among the alternatives, then there is some alternative, y , that a majority prefers to x . Since x is the worst alternative for the vulnerable minority, the majority that prefers y to x must comprise the vulnerable minority. If that majority has control, then their shared preference would prevent x from being chosen if they all felt strongly enough about the issue. Thus, if the vulnerable minority feels strongly about the issue, then one element of a safeguard for the minority is in place; all that is still needed, to protect the minority’s interest in avoiding the worst outcome, is for enough of

the remaining members of the latent majority coalition—another minority—to care enough about the issue. It would be too strong to say that the minority’s interests are secure simply in virtue of there being a majority preference cycle. But the fact that a majority, comprising the threatened minority, can be counted upon to prefer an alternative to the option that the minority considers worst,¹² together with the fact that the majority has control, means that the minority’s interest in avoiding the worst outcome is *conditionally secure*. It is secure conditional on another minority—the other latent coalition partners—feeling sufficiently strongly about the issue, too.

To form a reasonable belief about whether and when such conditional security will translate into actual security, one would need a model of the democratic process and the particular mechanism underlying majorities’ control. Abstracting from the details of any particular embodiment of “majority rule” and of that mechanism, the most one can say is that if the members of a vulnerable minority all agree on which outcome is worst, then their interest in avoiding that outcome is conditionally secure when no Condorcet winner exists and all majorities have control. While that conclusion is weaker than the one McGann advances, it is not insignificant, practically

¹²Note that the existence of such an alternative is guaranteed by the assumption that no Condorcet winner exists, but not by the assumption that a majority preference cycle occurs. So the conjectural argument I am sketching is not an argument about the consequences majority preference cycles per se. See example 3.2 in §3.2 for the explanation of why majority preference cycles are compatible with the existence of a Condorcet winner.

speaking.

Coming back to our overarching theme, the discussion illustrates again how normative democratic theory stands to benefit from abandoning the traditional understanding of popular rule and adopting in its place the concept of rule by multiple majorities. One is not in a position to appreciate the conditional security that minorities can enjoy in a regime of popular control if one operates with the “monolithic majority” model. If one imagines a democracy as a regime that empowers “the” majority at the expense of “the” minority, then the interests of the disempowered minority will be entirely contingent on the goodwill of the majority in control. That is of course precisely the feature of this imagined democracy that causes consternation for traditional liberal and republican writers, who warn of the “tyranny of the majority.” Their error is not the thought that things could go terribly wrong for minorities in democracy, but rather a simplistic assumption about the mechanism thought to be responsible for the outcomes of the democratic process. When minorities suffer ill treatment in democracies, the best explanation is not necessarily that a tyrannical group agent is abusing its monopoly on power. The best explanation may be rather that some other minority—the potential coalition partners of the ill-treated minority—does not care enough about those decisions. That, anyway, is the best explanation if all majorities’ have control in a democracy. The most effective remedy may be, not to restrict the institutional powers of majorities, as anti-majoritarian liberals and republicans typically suppose, but rather to reform institutions so that the potential coalition partners of vulnerable minorities do not need to care so much in order for the preferences that they

share with the vulnerable to constrain decisions.

§7.4 Conclusion

I have suggested that one dimension along which political regimes can be compared, as better or worse, is the degree to which they give all majorities control. The considerations that lend that claim *prima facie* plausibility are familiar arguments about the importance of social equality and the instrumental benefits of a group's control for the protection of its legitimate interests. The reasons some political theorists might initially greet the claim with skepticism are that democratic theorists are generally dismissive of the view that preference intensity has any legitimate role in the democratic process, and, prior to this chapter, the concept of popular democratic control was missing an important egalitarian requirement.

Neither basis for skepticism is sound, however. Among democrats, the potentially controversial claim about preference intensity is that the intense preferences of a minority should be allowed to prevail over the preferences of a majority, but the thesis that popular control is desirable does not entail that claim. It entails only that it is desirable that if a majority holds its preferences with sufficient intensity, then they should prevail. And as for equality, we have seen how one might formulate a principle of equal control, which characterizes the kind of equality that a fully democratic system of popular control requires. The conclusion that the principle might require compulsory voting, or other measures to eliminate the collective action problems facing particular groups, should suffice to convince the reader that the

concept of control can be used to formulate a suitably demanding normative ideal of democratic, popular control.

Finally, we saw in this chapter how the concept of rule by multiple majorities provides a fresh perspective on questions about minorities' status in democracy. Even if one thought that rights to political participation were only instrumentally valuable as means of giving majorities' control, one would have reason to deprive majorities of the power to disenfranchise minorities. Majorities' control rests on a foundation of universal suffrage and universal protection of democratic rights more generally. And building on McGann's (2006) argument, we identified reasons why giving all majorities control could be a means of extending at least conditional protection to minorities' interests more broadly, a good whose provision in a democracy would be hard to appreciate if one wrongly interpreted democracy as rule by "the" majority, rather than rule by multiple majorities.

Chapter 8

Conclusion

Robert Dahl once characterized democratic theory as concerned, “at a minimum, . . . with processes by which ordinary citizens exert a relatively high degree of control over leaders” (Dahl, 1956, p. 3). To study the *processes* by which citizens do this, one needs a well-defined *concept* of what they would allegedly be doing. What does it even mean, or what could it mean, to say that citizens have control over their leaders? No empirical or normative inquiry into popular control can get far without an answer to the conceptual question.

The main contribution of the book is a novel concept of popular control that satisfies two desiderata that pull in different directions. On the one hand, a concept of popular control should respect common intuitions to some degree; on the other, it should be useful to democratic theory and empirical social science. These two sets of criteria are in tension. On the one hand, we have the intuition that if citizens have control over their government, then the government should do what they want. On the other, a concept of popular control would be more useful to democratic theory and social science if it helped organize and clarify our thinking about democracies even in those

conditions where citizens' preferences do not aggregate into anything like a popular will, and where it makes little sense to refer to what "they" want.

The concept of rule by multiple majorities does well by each desideratum. Popular control over government does not imply that the government acts as "the citizens" want, or even as "the majority" wants. But it implies that every majority has control over the government, in the sense that the majority's shared preferences would constrain the government if the majority held its preferences with enough intensity. As I argued in chapter 4, that characterization of a majority's control does sufficient justice to our intuitions about control that it is not misleading to describe the result as a concept *of* popular control rather than something else. And, as I explained in chapter 5, every majority can have control—and an adequate degree of control—at the same time, even if citizens' preferences do not aggregate into a popular will, even if majority preference cycles occur, even if no Condorcet winner exists, and even if the government is therefore bound to frustrate the preferences of a majority. The concept of rule by multiple majorities expands the range of conditions in which democratic theorists and social scientists can examine the core intuitions about democracy that people are trying to express with the language of popular control.

A well-defined concept of popular control puts scholars in a position to investigate empirical and normative questions about democracy without the simplifying assumptions of a "representative voter" or a singular, homogeneous majority. Under what conditions can democratic elections establish popular control over decision-makers? If democracies have a tendency to bring collective decisions or public officials under popular control, is this

fact a reason to prefer them to non-democratic alternatives? Inquiry on both questions has so far operated with an informal understanding of popular control that does not square well with insights from social choice theory. Theoretical and empirical work on the mechanisms of electoral accountability usually rests, implicitly or explicitly, on the assumption that “the voters” all have the same preferences over the elected official’s behavior, or the assumption that their preferences aggregate into consistent majority preferences. Similarly, normative democratic theory typically relies, more or less explicitly, on a simple model of democracy as a regime that empowers “the” majority to rule.

These simplifying assumptions stand in tension with what social choice theory has taught us. Citizens’ preferences need not aggregate into anything resembling a collective will. In particular, it is not in general meaningful to refer to “the will of the majority.” No majority is homogeneous; for each majority subset of the population, its members will share preferences over some pairs of options but not others. And there is in general no reason to assume that the pairwise preferences of different majorities can all be satisfied at the same time. Thus, simplifying assumptions are just that: simplifications of a more complicated reality. But using the book’s concept of popular control, meaningful inquiry into democracy’s connection with popular control and its normative significance can proceed without those assumptions. It will still sometimes be expedient to approach a question under the assumption of a binary decision or single-peaked preferences. But now that we have a concept of popular control, the coherence of which does not *depend* on those assumptions, we can ask, in any particular context, whether they are merely

innocent simplifications of a more complicated reality or in fact misleading distortions of it.

Within political theory, the conventional term for the intellectual exercise that led to this result would be *conceptual analysis*, but that is not the most descriptively accurate label. It suggests that we start with a given concept and then arrive at the finished product by analyzing what we were given into its constituent parts. No one could say of the concept of rule by multiple majorities that it is the same as some concept of popular control that we started with, only more clearly defined and with its constituent parts and their relationships traced out. A better term for the exercise is *conceptual exploration*. We are searching logical space for the set of conditions we will use to define a new concept. The search does not take place in a historical and social vacuum. Although we want a set of conditions that are useful for the explanatory and evaluative aims of social science and political theory, we also want a set of conditions that mark out a phenomenon not too radically different from what other people have in mind when they use terms like *popular control*. Otherwise whatever we do with the resulting concept will not contribute to the existing fields of inquiry we aim to contribute to.

The method of search borrowed tools and strategies of argument from game theory and social choice theory. That makes the book somewhat unusual within political theory, which usually does not draw on those approaches. However, the infrequent use of game theory and social choice theory in the discipline cannot be explained by their lack of useful applications. I hope the reader agrees that the book is proof. There are in fact clear continuities between the kinds of arguments that political theo-

rists conventionally make and the kinds of arguments that “formal political theorists” make with the help of tools from game theory and social choice theory (Johnson, 2014; Ingham, 2015; Patty and Penn, 2015). Each type of theory is concerned with figuring out “what follows from what”—the structure of arguments, what different combinations of assumptions imply, what assumptions are needed to reach a conclusion. And in that endeavor, each type of theory helps itself to “models” that describe hypothetical scenarios, constructs thought experiments, draws conceptual distinctions. The distinction between the two kinds of political theory is useful if one is doing an ethnography of political science, but not if one is trying to understand the kinds of arguments that people are making. I hope the book contributes to a broader understanding of what political theory is, or can be, and the kinds of methods it can usefully employ.

The goal of our conceptual exploration was an account of popular control that is useful both to social science and normative democratic theory. I tried to show, in particular, that it is useful both for the study of existing democracies, in which retrospective voting has been thought to be a plausible mechanism of popular control over elected officials (chapter 6), and for a normative theory of democracy, which aims to identify the features of democratic regimes that could justify preferring them to non-democratic alternatives (chapter 7). The arguments of those two chapters may seem to stand in some tension, and I would like to conclude with thoughts on why the tension is only apparent.

On the one hand, the reader might think that the explanation for why retrospective voting could be a plausible mechanism of giving majorities control works only because the concept of control is insufficiently demanding. On the other, a democracy's tendency to establish popular control will be a reason to value it over non-democratic alternatives only if popular control is a non-trivial accomplishment. However, the tension between those two thoughts is only apparent. Control comes in degrees. That majorities of ordinary citizens have more significant degrees of control over political leaders in democratic regimes than in non-democratic alternatives is plausibly one reason to prefer the former to the latter, all else being equal. That is true even if one thinks that a majority's merely having *adequate* control, as defined in chapter 5, is not very significant. The fact (if it is a fact) that the actually existing regimes we designate as democracies tend to give majorities some control over elected officials, through retrospective voting or some other mechanism of accountability, is important even if we think the degree of control is insufficient from a normative point of view. It is important because it affects whether we think of these regimes as at least imperfect instantiations of our idea of democracy. And it is important because it is easier to imagine how popular control might be strengthened in actually existing democracies if we can at least identify a mechanism by which they already establish some degree of popular control.

The arguments about retrospective voting were meant to bolster the plausibility of seeing it as a mechanism of popular control, but they were not meant to show that all is well in existing democracies, as far as the ideal of popular democratic control goes. Those arguments leave open the pos-

sibility that actual democracies stand in need of institutional reforms that would strengthen and expand popular control, perhaps through the creation of sites of direct popular participation in decision-making (McCormick, 2011; Landemore, 2017; Rahman, 2017). As we saw in the last chapter, there are also good reasons to believe that existing democracies distribute control unequally to different majorities. Majorities comprising society's most socially advantaged citizens probably have control over public officials, and as much control as any other groups of comparable size. It is the majorities comprising society's least advantaged that arguably have too little control over public officials. Are there institutional reforms that would address this problem?

McCormick (2006, 2011) proposes a "Tribunate Assembly," modeled on the popular assemblies of the Roman Republic and Machiavelli's interpretation of their role as mechanisms of popular control over elites. Its 51 members, selected at random from the population but excluding households in the top ten percent of the distribution of wealth, would be like the tribunes of the ancient Roman Republic. As applied to the United States, the proposal is that during one-year terms, they would have the power to veto a piece of legislation, a Supreme Court decision, and an executive order; the power to call a national referendum; and the power to initiate impeachment proceedings against a federal official (McCormick, 2006, p. 160).¹ The Tribunate Assembly would be an example of a "deliberative minipublic": as-

¹The suggestion is that that the assembly be able to exercise each of these powers once per term.

semblies of ordinary citizens, selected at random from the population, who gather information and deliberate together about a policy problem before either making a binding decision themselves or conveying their judgments to the public at large or other decision-makers. Examples include James Fishkin's Deliberative Polls, the British Columbia Citizens' Assembly, and the Oregon Citizens' Initiative Review (Warren and Pearse, 2008; Fishkin, 2009; Gastil, Knobloch and Richards, 2015). McCormick's proposal would differ from other minipublics in the institutional powers assigned to the assembly and in its demographic composition. While other proponents of minipublics aim for assemblies that reproduce the entire citizen population in miniature, the Tribune Assembly would exclude the wealthy.

In what sense could the creation of such an assembly strengthen popular control in democracies? It might bring the policy agenda or political and economic elites under the control of a *popular* body, that is, a body of "ordinary" citizens. But if one identifies popular control with a system in which *majorities* of the citizen population have control, as I claim one should, then empowering a tiny minority of the population does not in fact establish popular control even if this group has a popular character.

The assembly could strengthen popular control by a different mechanism, however. Its power to propose national referendums would put certain policy decisions under the control of majorities of the electorate. Less obviously, the exercise of its other powers could strengthen the control of majorities over elected officials, in virtue of what the assembly's actions would *signal* to the general electorate. If the Tribune Assembly initiated impeachment proceedings against the President, or vetoed legislation passed by a majority

party, voters might reasonably infer that these elected officials are probably not pursuing the objectives that the members of the Tribune Assembly would like them to pursue. And to the extent that voters can trust the Assembly's members to share *their* views about the objectives that policymakers should be pursuing, voters might reasonably infer that the elected officials are not pursuing the objectives that *they* would like them to pursue. Anticipating such inferences, elected officials would have stronger incentives to target the policy outcomes that citizens would prefer for them to target. This is similar to the retrospective voting mechanism described in chapter 6. Citizens' preferences over policymakers' choice of "targets," or "aims" (Christiano, 1996), might then have constraining effects on what elected officials do, and have such effects under a wider range of circumstances than in the absence of the Tribune Assembly, as citizens would have access to a new, informative "measure" of what elected officials are up to, namely the Assembly's actions.

The assembly's signaling function could be made more explicit and direct. In addition to the functions McCormick describes, it might conduct periodic reviews of elected officials and release statements to voters prior to elections, similar to a proposal made by Gastil (2000). Or it might conduct audits of various measures of social prosperity—measures of economic growth, inequality, health and mortality, environmental pollution, or any other outcomes which elected officials can potentially influence through their choice of policy and which we might wish to bring under popular control.²

²The assumption behind these conjectures is that voters would be prepared to revise their beliefs about elected officials on the basis of the As-

Through such functions, the Tribune Assembly would be playing a role similar to that of a free and independent media or an independent auditor, institutions which have been shown to strengthen electoral accountability in actual democracies (Besley and Burgess, 2002; Ferraz and Finan, 2008; Snyder and Strömberg, 2010).

I have referred here to “voters” as though they were a single mass, but of course the point of departure for the book is the recognition that we are not generally justified in referring to “voters’ preferences.” Nor should we commit the same mistake and treat “the non-wealthy,” or “ordinary citizens,” as a single mass with homogeneous preferences. From the perspective adopted in the book, then, the way to think about McCormick’s proposal for a Tribune Assembly is that it would recruit representatives of one particular majority, namely the majority comprising all those whose wealth falls below the threshold for eligibility. Through its power to call national referendums and to signal information to the general electorate, it might strengthen the control of this one particular majority. The justification for institutional reforms with this anticipated effect might appeal to the principle of equal control, from chapter 7: given the evidence that democracies do not currently give majorities comprising non-wealthy citizens the same sembly’s actions and reports. While the empirical literature on the question is not extensive, there is some evidence that citizens would revise their beliefs about policies upon learning a deliberative minipublic’s conclusions, although the effects appear to depend on the policy in question (Cutler et al., 2008; Boulianne, 2018; Gastil et al., 2018; Ingham and Levin, 2018*b,a*).

degree of control as majorities comprising the wealthy, equalizing control across groups requires strengthening the control of the former, or reducing the control of the latter.

While there are surely ways to strengthen the control of particular majorities, it would be a mistake to view popular democratic control over policymakers as a utopian ideal that actual democracies fail to realize. Control comes in degrees, and there are good reasons to believe that majorities have at least some control over the behavior of elected officials. They may not exercise control over the details of the policies that their representatives choose, but it is reasonable to believe that majorities have some control over what their leaders are trying to achieve with their policy choices—lowered unemployment, increased economic growth, peace, and so on.

Claims about popular control over elected officials can be made to look ridiculous if they are assessed with unreasonable interpretations of what popular control requires. That is a favorite rhetorical move of skeptics of popular rule. “To suggest that we can ever hope to have the power to make [our ostensible rulers] act just as we would wish them to suggests that it is really we, not they, who are ruling. This is an illusion, and probably a somewhat malign illusion” (Dunn, 1999, p. 20). The illusion, I take it, is the belief that we could ever make our ostensible rulers act just as we would wish, not the conceptual claim that if we could do so, then it would be really we who are ruling. The factual claim does indeed seem far-fetched. But notice how much work is being done by that little word *just*. The claim is perhaps

not so outlandish if we substitute *roughly as we would wish*. And it is not outlandish at all if, in place of *just as we would wish*, we substitute *consistent with widely shared preferences, if there are any widely shared preferences, and if they are held with sufficient strength*. Now the factual claim is modest: majorities will sometimes share preferences, and they will sometimes feel strongly about their preferences, and it is not unreasonable to expect that with democratic institutions in place, their shared preferences will constrain the behavior of public officials. When that happens, we are justified in describing majorities as controlling their public officials. And since rulers everywhere exercise their power by controlling the agents who act on their behalf, perhaps we are also justified in saying that it is really they, majorities of ordinary citizens, who are ruling.

Bibliography

- Abizadeh, Arash. 2012. "On the Demos and Its Kin: Nationalism, Democracy, and the Boundary Problem." *American Political Science Review* 106(4):867–882.
- Achen, Christopher H. and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton, NJ: Princeton University Press.
- Alt, James, Ethan Bueno de Mesquita and Shanna Rose. 2011. "Disentangling Accountability and Competence in Elections: Evidence from US Term Limits." *The Journal of Politics* 73(01):171–186.
- Arneson, Richard. 2009. The Supposed Right to a Democratic Say. In *Contemporary Debates in Political Philosophy*, ed. Thomas Christiano and John Christman. Oxford: Oxford: Wiley-Blackwell pp. 197–212.
- Arrow, Kenneth. 1951. *Social Choice and Individual Values*. New York: John Wiley & Sons, Inc.
- Ashworth, Scott. 2012. "Electoral Accountability: Recent Theoretical and Empirical Work." *Annual Review of Political Science* 15:183–201.

- Ashworth, Scott and Ethan Bueno de Mesquita. 2006. "Delivering the Goods: Legislative Particularism in Different Electoral and Institutional Settings." *The Journal of Politics* 68(1):168–179.
- Ashworth, Scott, Ethan Bueno de Mesquita and Amanda Friedenberg. 2018. "Learning about voter rationality." *American Journal of Political Science* 62(1):37–54.
- Austin-Smith, David and Jeffrey Banks. 1999. *Positive Political Theory I: Collective Preference*. Ann Arbor, MI: University of Michigan Press.
- Baron, David P and John A Ferejohn. 1989. "Bargaining in Legislatures." *American political science review* 83(4):1181–1206.
- Barro, Robert J. 1973. "The Control of Politicians: An Economic Model." *Public Choice* 14(1):19–42.
- Barry, Brian. 2002. "Capitalists Rule OK? Some Puzzles about Power." *Politics, Philosophy & Economics* 1(2):155–184.
- Bartels, Larry M. 2016. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton, NJ: Princeton University Press.
- Bechtel, Michael M. and Jens Hainmueller. 2011. "How Lasting Is Voter Gratitude? An Analysis of the Short-and Long-Term Electoral Returns to Beneficial Policy." *American Journal of Political Science* 55(4):852–868.
- Beerbohm, Eric. 2012. *In Our Name: The Ethics of Democracy*. Princeton, NJ: Princeton University Press.

- Besley, Timothy. 2006. *Principled Agents? The Political Economy of Good Government*. Oxford: Oxford University Press.
- Besley, Timothy and Robin Burgess. 2002. "The Political Economy of Government Responsiveness: Theory and Evidence from India." *The Quarterly Journal of Economics* 117(4):1415–1451.
- Black, Duncan. 1958. *The Theory of Committees and Elections*. London: Cambridge University Press.
- Boulianne, Shelley. 2018. "Mini-publics and Public Opinion: Two Survey-Based Experiments." *Political Studies* 66(1):119–136.
- Brennan, Jason. 2016. *Against Democracy*. Princeton, NJ: Princeton University Press.
- Broome, John. 1999. *Ethics out of Economics*. Cambridge University Press.
- Busby, Ethan C, James N Druckman and Alexandria Fredendall. 2017. "The Political Relevance of Irrelevant Events." *The Journal of Politics* 79(1):346–350.
- Butler, Daniel M. 2014. *Representing the Advantaged: How Politicians Reinforce Inequality*. Cambridge: Cambridge University Press.
- Carnes, Nicholas. 2013. *White-Collar Government: The Hidden Role of Class in Economic Policy Making*. Chicago: University of Chicago Press.
- Chen, Jidong, Jennifer Pan and Yiqing Xu. 2016. "Sources of Authoritarian Responsiveness: A Field Experiment in China." *American Journal of Political Science* 60(2):383–400.

- Christiano, Thomas. 1993. Social Choice and Democracy. In *The Idea of Democracy*, ed. David Copp, Jean Hampton and John E. Roemer. Cambridge: Cambridge University Press.
- Christiano, Thomas. 1996. *The Rule of the Many: Fundamental Issues in Democratic Theory*. New York: Routledge.
- Cohen, Joshua. 1986. "An Epistemic Conception of Democracy." *Ethics* 97(1):26–38.
- Cohen, Joshua. 1997. Deliberation and Democratic Legitimacy. In *Deliberative Democracy: Essays on Reason and Politics*, ed. James Bohman and William Rehg. Cambridge, MA: MIT Press.
- Crespin, Michael H. and David W. Rohde. 2010. "Dimensions, Issues, and Bills: Appropriations Voting on the House Floor." *The Journal of Politics* 72(4):976–989.
- Cutler, Fred, Richard Johnston, Kenneth R. Carty, André Blais and Patrick Fournier. 2008. Deliberation, Information, and Trust: the British Columbia Citizens' Assembly as Agenda Setter. In *Designing Deliberative Democracy*.
- Dahl, Robert. 1971. *Polyarchy: Participation and Opposition*. New Haven: Yale University Press.
- Dahl, Robert A. 1956. *A Preface to Democratic Theory: Expanded Edition*. Chicago, IL: The University of Chicago Press.

- Dahl, Robert A. 1957. "The Concept of Power." *Behavioral Science* 2(3):201–215.
- Dahl, Robert A. 2005. *Who governs?: Democracy and power in an American city*. second ed. Yale University Press.
- Dasgupta, Partha and Eric Maskin. 2008. "On the Robustness of Majority Rule." *Journal of the European Economic Association* 6(5):949–973.
- de Tocqueville, Alexis. 2000. *Democracy in America*. New York: Perennial.
- Dixit, Avinash and John Londregan. 1996. "The Determinants of Success of Special Interests in Redistributive Politics." *The Journal of Politics* 58(4):1132–1155.
- Dougherty, Keith, Michael Lynch and Anthony Madonna. 2014. "Partisan Agenda Control and the Dimensionality of Congress." *American Politics Research* 42(4):600–627.
- Dowding, Keith. 2003. "Resources, Power and Systematic Luck: A Response to Barry." *Politics, Philosophy & Economics* 2(3):305–322.
- Dowding, Keith. 2006. "Can populism be defended? William Riker, Gerry Mackie and the interpretation of democracy." *Government and Opposition* 41(3):327–346.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. Boston, MA: Addison-Wesley Publishing Company, Inc.
- Dryzek, John. 2000. *Deliberative Democracy and Beyond*. New York: Oxford University Press.

- Dryzek, John S. and Christian List. 2003. "Social Choice Theory and Deliberative Democracy: A Reconciliation." *British Journal of Political Science* 33:1–28.
- Duch, Raymond M. and Randolph T. Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results*. Cambridge: Cambridge University Press.
- Dunn, John. 1999. Situating Political Accountability. *in* Przeworski, Manin and Stokes (1999).
- Dunn, John. 2005. *Democracy: a history*. New York: Atlantic Monthly Press.
- Dworkin, Ronald. 2000. *Sovereign Virtue: The Theory and Practice of Equality*. Cambridge, MA: Harvard University Press.
- Egan, Patrick J. 2014. "Double-Peaked Policy Preferences and 'Do Something' Politics." *Journal of Politics* 76(2):333–49.
- Elster, Jon. 2013. *Securities Against Misrule: Juries, Assemblies, Elections*. Cambridge: Cambridge University Press.
- Estlund, David M. 2008. *Democratic Authority*. Princeton, NJ: Princeton University Press.
- Farrar, Cynthia, James S. Fishkin, Donald P. Green, Christian List, Robert Luskin and Elizabeth Levy Paluck. 2010. "Disaggregating Deliberation's Effects: An Experiment within a Deliberative Poll." *British Journal of Political Science* 40:333–347.

- Fearon, James. 1999. Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. *in* Przeworski, Manin and Stokes (1999).
- Ferejohn, John. 1986. "Incumbent Performance and Electoral Control." *Public Choice* 50:5–26.
- Ferraz, Claudio and Frederico Finan. 2008. "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes." *The Quarterly Journal of Economics* 123(2):703–745.
- Fiorina, Morris P. 1981. *Retrospective Voting in American National Elections*. New Haven, CT: Yale University Press.
- Fishkin, James S. 2009. *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford: Oxford University Press.
- Fowler, Anthony and Andrew B. Hall. 2016. "The Elusive Quest for Convergence." *Quarterly Journal of Political Science* 11(1):131–149.
- Fowler, Anthony and Andrew B. Hall. 2018. "Do shark attacks influence presidential elections? Reassessing a prominent finding on voter competence." *The Journal of Politics* .
URL: <https://doi.org/10.1086/699244>
- Fowler, Anthony and Pablo B. Montagnes. 2015a. "College Football, Elections, and False-Positive Results in Observational Research." *Proceedings of the National Academy of Sciences* 112(45):13800–13804.

- Fowler, Anthony and Pablo B. Montagnes. 2015*b*. “Reply to Healy et al.: Value of Ex Ante Predictions and Independent Tests for Assessing False-Positive Results.” *Proceedings of the National Academy of Sciences* 112(48):E6592–E6592.
- Gasper, John T. and Andrew Reeves. 2011. “Make It Rain? Retrospection and the Attentive Electorate in the Context of Natural Disasters.” *American Journal of Political Science* 55(2):340–355.
- Gastil, John. 2000. *By Popular Demand: Revitalizing Representative Democracy Through Deliberative Elections*. Berkeley: University of California Press.
- Gastil, John, Katherine Knobloch and Robert Richards. 2015. Empowering Voters through Better Information: Analysis of the Citizens’ Initiative Review, 2010–2014. Technical report the Democracy Fund.
- Gastil, John, Katherine R Knobloch, Justin Reedy, Mark Henkels and Katherine Cramer. 2018. “Assessing the Electoral Impact of the 2010 Oregon Citizens’ Initiative Review.” *American Politics Research* 46(3):534–563.
- Gehlbach, Scott. 2013. *Formal Models of Domestic Politics*. Cambridge: Cambridge University Press.
- Gerber, Alan S, Donald P Green and Christopher W Larimer. 2008. “Social pressure and voter turnout: Evidence from a large-scale field experiment.” *American Political Science Review* 102(1):33–48.

- Gilens, Martin. 2012. *Affluence and Influence: Economic Inequality and Political Power in America*. Princeton, NJ: Princeton University Press.
- Gilens, Martin and Benjamin I. Page. 2014. "Testing theories of American politics: Elites, interest groups, and average citizens." *Perspectives on Politics* 12(03):564–581.
- Goodman, Al. 2003. "Polls: 90 percent of Spaniards against war.". CNN.com; accessed August 6, 2018.
URL: <http://www.cnn.com/2003/WORLD/europe/03/29/sprj.irq.spain/>
- Gourevitch, Alex. 2015. *From Slavery to the Cooperative Commonwealth: Labor and Republican Liberty in the Nineteenth Century*. Cambridge: Cambridge University Press.
- Gourevitch, Alex. 2016. "The limits of a basic income: Means and ends of workplace democracy." *Basic Income Studies* 11(1):17–28.
- Green, Jeffrey E. 2010. *The eyes of the people: democracy in an age of spectatorship*. Oxford: Oxford University Press.
- Hamilton, Alexander, James Madison and John Jay. 2003. *The Federalist with Letters of "Brutus"*. Cambridge: Cambridge University Press.
- Hausman, Daniel M. 1992. *Essays on philosophy and economic methodology*. Cambridge: Cambridge University Press.
- Healy, Andrew J. Neil Malhotra and Cecilia Hyunjung Mo. 2010. "Irrelevant Events Affect Voters' Evaluations of Government Performance." *Proceed-*

ings of the National Academy of Sciences of the United States of America 107(29):12804–12809.

Healy, Andrew and Neil Malhotra. 2009. “Myopic Voters and Natural Disaster Policy.” *American Political Science Review* 100(3):387–406.

Healy, Andrew and Neil Malhotra. 2010. “Random Events, Economic Losses, and Retrospective Voting: Implications for Democratic Competence.” *Quarterly Journal of Political Science* 5(2):193–208.

Healy, Andrew and Neil Malhotra. 2013. “Retrospective Voting Reconsidered.” *Annual Review of Political Science* 16:285–306.

Healy, Andrew, Neil Malhotra and Cecilia Hyunjung Mo. 2015. “Determining False-Positives Requires Considering the Totality of Evidence.” *Proceedings of the National Academy of Sciences* 112(48):E6591–E6591.

Hibbs, Douglas A. 2006. *The Oxford Handbook of Political Economy*. Oxford: Oxford University Press chapter Voting and the Macroeconomy.

Hobbes, Thomas. 1998. *On the Citizen*. Cambridge University Press.

Hotelling, Harold. 1929. “Stability in Competition.” *The Economic Journal* 39(153):41–57.

Hsieh, Nien-hê. 2005. “Rawlsian justice and workplace republicanism.” *Social Theory and Practice* 31(1):115–142.

Huber, Gregory A., Seth J. Hill and Gabriel S. Lenz. 2012. “Sources of Bias in Retrospective Decision Making: Experimental Evidence on Vot-

- ers' Limitations in Controlling Incumbents." *American Political Science Review* 106(04):720–741.
- Ingham, Sean. 2015. "Theorems and Models in Political Theory: An Application to Pettit on Popular Control." *The Good Society* 24(1):98–117.
- Ingham, Sean. 2016a. "Popular Rule in Schumpeter's Democracy." *Political Studies* 64(4):1071–1087.
- Ingham, Sean. 2016b. "Social Choice and Popular Control." *Journal of Theoretical Politics* 28(2):331–349.
- Ingham, Sean. 2018. "Why Arrow's theorem matters for political theory even if preference cycles never occur." *Public Choice* .
URL: <https://doi.org/10.1007/s11127-018-0521-9>
- Ingham, Sean and Ines Levin. 2018a. "Can Deliberative Minipublics Influence Public Opinion? Theory and Experimental Evidence." *Political Research Quarterly* 71(3):654–667.
- Ingham, Sean and Ines Levin. 2018b. "Effects of Deliberative Minipublics on Public Opinion: Experimental Evidence from a Survey on Social Security Reform." *International Journal of Public Opinion Research* 30(1):51–78.
- Jacobs, Lawrence R. and Benjamin I. Page. 2005. "Who influences U.S. foreign policy?" *American Political Science Review* 99(1):107–123.
- Johnson, James. 2014. "Models among the Political Theorists." *American Journal of Political Science* 58(3):547–560.

- Karol, David and Edward Miguel. 2007. "The Electoral Cost of War: Iraq Casualties and the 2004 U.S. Presidential Election." *The Journal of Politics* 69(3):633–648.
- Key, V.O. 1966. *The Responsible Electorate: Rationality in Presidential Voting, 1936–1960*. Cambridge, MA: The Belknap Press of Harvard University Press.
- Kingdon, John W. 1989. *Congressmen's voting decisions*. Ann Arbor, MI: University of Michigan Press.
- Knight, Jack and James Johnson. 1994. "Aggregation and Deliberation: On the Possibility of Democratic Legitimacy." *Political Theory* 22(2):277–296.
- Knight, Jack and James Johnson. 2011. *The Priority of Democracy: Political Consequences of Pragmatism*. Princeton, NJ: Princeton University Press.
- Kolodny, Niko. 2014. "Rule over none II: social equality and the justification of democracy." *Philosophy & Public Affairs* 42(4):287–336.
- Kriner, Douglas L. and Francis X. Shen. 2007. "Iraq Casualties and the 2006 Senate Elections." *Legislative Studies Quarterly* 32(4):507–530.
- Kurrild-Klitgaard, Peter. 2017. "Trump, Condorcet and Borda: Voting paradoxes in the 2016 Republican presidential primaries." *European Journal of Political Economy* 50:xxx–xxx.
- Landa, Dimitri and Ryan Pevnick. N.d. "What Could Justify Representative Democracy." Unpublished manuscript.

- Landemore, Hélène. 2012. *Democratic reason: Politics, collective intelligence, and the rule of the many*. Princeton, NJ: Princeton University Press.
- Landemore, Hélène. 2017. “Deliberative Democracy as Open, Not (Just) Representative Democracy.” *Daedalus* 146(3):51–63.
- Lane, Melissa. 2016. Popular sovereignty as control of office-holders: Aristotle on Greek democracy. In *Popular Sovereignty in Historical Perspective*, ed. Richard Bourke and Quentin Skinner. Cambridge University Press.
- Lax, Jeffrey R and Justin H Phillips. 2009. “Gay Rights in the States: Public Opinion and Policy Responsiveness.” *American Political Science Review* 103(03):367–386.
- Le Bihan, Patrick. 2016. “Popular referendum and electoral accountability.” *Political Science Research and Methods* pp. 1–17.
- Le Breton, Michel. 1987. “On the Core of Voting Games.” *Social Choice and Welfare* 4:295–305.
- Lee, David S., Enrico Moretti and Matthew J. Butler. 2004. “Do Voters Affect or Elect Policies? Evidence from the U. S. House.” *The Quarterly Journal of Economics* 119(3):807–859.
- Lewis-Beck, Michael S. and Mary Stegmaier. 2009. *Oxford Handbook of Political Behavior*. Oxford chapter Economic Models of Voting.
- Lindbeck, Assar and Jörgen W. Weibull. 1987. “Balanced-budget redistri-

bution as the outcome of political competition.” *Public Choice* 52(3):273–297.

List, Christian. 2002. “Two Concepts of Agreement.” *The Good Society* 11(1):72–79.

List, Christian and Laura Valentini. 2016. “The methodology of political theory.” *The Oxford Handbook of Philosophical Methodology* pp. 525–550.

List, Christian and Philip Pettit. 2011. *Group agency: The possibility, design, and status of corporate agents*. Oxford University Press.

List, Christian, Robert C Luskin, James S Fishkin and Iain McLean. 2012. “Deliberation, single-peakedness, and the possibility of meaningful democracy: evidence from deliberative polls.” *The Journal of Politics* 75(1):80–95.

Locke, John. 1988. *Two Treatises of Government*. Cambridge: Cambridge University Press.

Lovett, Frank. 2010. *A General Theory of Domination and Justice*. Oxford: Oxford University Press.

Lupia, Arthur and John G. Matsusaka. 2004. “Direct Democracy: New Approaches to Old Questions.” *Annual Review of Political Science* 7:463–482.

Mackie, Gerry. 2003. *Democracy Defended*. Cambridge: Cambridge University Press.

- Manin, Bernard. 1997. *The Principles of Representative Government*. Cambridge: Cambridge University Press.
- Mansbridge, Jane. 2009. "A "Selection Model" of Political Representation." *Journal of Political Philosophy* 17(4):369–398.
- Marx, Karl. 1977. *Karl Marx: Selected Writings*. Oxford: Oxford University Press.
- Maskin, Eric and Jean Tirole. 2004. "The Politician and the Judge: Accountability in Government." *The American Economic Review* 94(4):1034–1054.
- Matsusaka, John G. 2010. "Popular Control of Public Policy: A Quantitative Approach." *Quarterly Journal of Political Science* 5(2):133–167.
- May, Kenneth O. 1952. "A Set of Independent Necessary and Sufficient Conditions For Simple Majority Decision." *Econometrica* 20:680–684.
- McCormick, John. 2006. "Contain the Wealthy and Patrol the Magistrates: Restoring Elite Accountability to Popular Government." *American Political Science Review* 100(2):147–163.
- McCormick, John P. 2011. *Machiavellian Democracy*. Cambridge: Cambridge University Press.
- McGann, Anthony. 2006. *The Logic of Democracy*. Ann Arbor, MI: The University of Michigan Press.
- McKelvey, Richard. 1979. "General Conditions for Global Intransitivities in Formal Voting Models." *Econometrica* 47:1085–1112.

- McKelvey, Richard and Norman Schofield. 1987. "Generalized Symmetry Conditions at a Core Point." *Econometrica* 55:923–933.
- Medearis, John. 2001. *Joseph Schumpeter's Two Theories of Democracy*. Cambridge: Harvard University Press.
- Miller, David. 1992. "Deliberative Democracy and Social Choice." *Political Studies* XL(special issue):54–67.
- Müller, Jan-Werner. 2016. *What Is Populism?* Philadelphia, PA: University of Pennsylvania Press.
- Pattanaik, Prasanta K. 2014. *The Arrow Impossibility Theorem*. New York: Columbia University Press chapter Introduction.
- Patty, John W. and Elizabeth Maggie Penn. 2014. *Social Choice and Legitimacy: The Possibilities of Impossibility*. Cambridge: Cambridge University Press.
- Patty, John W. and Elizabeth Maggie Penn. 2015. "Aggregation, evaluation, and social choice theory." *The Good Society* 24(1):49–72.
- Penn, Elizabeth Maggie, John W. Patty and Sean Gailmard. 2011. "Manipulation and Single-Peakedness: A General Result." *American Journal of Political Science* 55(2):436–449.
- Persson, Torsten and Guido Enrico Tabellini. 2000. *Political economics: explaining economic policy*. Cambridge, MA: MIT Press.

- Pettit, Philip. 1991. Decision Theory and Folk Psychology. In *Foundations of Decision Theory*, ed. Michael Bacharach and Susan Hurley. Basil Blackwell.
- Pettit, Philip. 2008. "Three Conceptions of Democratic Control." *Constellations* 15(1):46–55.
- Pettit, Philip. 2012. *On the People's Terms: A Republican Theory and Model of Democracy*. Cambridge: Cambridge University Press.
- Pildes, Richard H. and Elizabeth S. Anderson. 1990. "Slinging Arrows at Democracy: Social Choice Theory, Value Pluralism, and Democratic Politics." *Columbia Law Review* 90(8):2121–2214.
- Plott, Charles R. 1967. "A Notion of Equilibrium and Its Possibility Under Majority Rule." *American Economic Review* 57:787–806.
- Poole, Keith and Howard Rosenthal. 2007. *Ideology and Congress*. 2nd edition ed. New Brunswick: Transaction Publishers.
- Posner, Eric A and E Glen Weyl. 2015. "Voting squared: Quadratic voting in democratic politics." *Vanderbilt Law Review* 68:441.
- Posner, Eric A. and E. Glen Weyl. 2018. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Princeton, NJ: Princeton University Press.
- Przeworski, Adam, Bernard Manin and Susan Stokes, eds. 1999. *Democracy, Accountability, and Representation*. Cambridge: Cambridge University Press.

- Rahman, K. Sabeel. 2017. *Democracy Against Domination*. Oxford: Oxford University Press.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Regenwetter, Michel, Bernard Grofman, A. A. J. Marley and Ilia Tsetlin. 2006. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge: Cambridge University Press.
- Riker, William. 1982. *Liberalism Against Populism*. New York: Waveland Press.
- Risse, Mathias. 2004. "Arguing for majority rule." *Journal of Political Philosophy* 12(1):41–64.
- Saari, Donald. 1997. "The Generic Existence of a Core for q-Rules." *Economic Theory* 9(2):219–260.
- Saari, Donald G. 2003. "Capturing the "Will of the People"." *Ethics* 113(2):333–349.
- Saunders-Hastings, Emma. 2018. "Plutocratic Philanthropy." *The Journal of Politics* 80(1):149–161.
- Schumpeter, Joseph. 1942. *Capitalism, Socialism, and Democracy*. New York: Harper & Brothers.
- Scwhartzberg, Melissa. 2014. *Counting the Many: The Origins and Limits of Supermajority Rule*. Cambridge: Cambridge University Press.

- Sen, Amartya. 1977a. "On weights and measures: informational constraints in social welfare analysis." *Econometrica* 45(7):1539–1572.
- Sen, Amartya. 1977b. "Social Choice Theory: A Re-Examination." *Econometrica* 45(1):53–88.
- Shapiro, Ian. 2003. *The State of Democratic Theory*. Princeton, NJ: Princeton University Press.
- Shapiro, Ian. 2016. *Politics Against Domination*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.
- Shepsle, Kenneth A. 1979. "Institutional arrangements and equilibrium in multidimensional voting models." *American Journal of Political Science* 23(1):27–59.
- Shepsle, Kenneth A. 1986. Institutional Equilibrium and Equilibrium Institutions. In *Political Science: The Science of Politics*, ed. Herbert F. Weisberg. Algora Publishing p. 51.
- Shepsle, Kenneth A. and Barry R. Weingast. 1981. "Structure-Induced Equilibrium and Legislative Choice." *Public Choice* 37(3):503–519.
- Sher, Itai. N.d. "Neutral Freedom and Freedom as Control."
- Skinner, Quentin. 1998. *Liberty before Liberalism*. Cambridge: Cambridge University Press.
- Snyder, James M. and David Strömberg. 2010. "Press Coverage and Political Accountability." *Journal of Political Economy* 118(2):355–408.

- Snyder, James M. and Michael M. Ting. 2003. "Roll calls, party labels, and elections." *Political Analysis* 11(4):419–444.
- Thompson, Dennis. 2002. *Just Elections: Creating a Fair Electoral Process in the United States*. Chicago: The University of Chicago Press.
- Tideman, Nicolaus. 2006. *Collective Decision and Voting: The Potential for Public Choice*. Burlington, VT: Ashgate Publishing Company.
- Tuck, Richard. 2016. *The Sleeping Sovereign: The Invention of Modern Democracy*. Cambridge: Cambridge University Press.
- Verba, Sidney, Kay Lehman Schlozman and Henry E Brady. 1995. *Voice and equality: Civic voluntarism in American politics*. Cambridge, MA: Harvard University Press.
- Viehoff, Daniel. 2014. "Democratic equality and political authority." *Philosophy & Public Affairs* 42(4):337–375.
- Waldron, Jeremy. 1999. *The Dignity of Legislation*. Cambridge: Cambridge University Press.
- Warren, Mark and Hilary Pearse, eds. 2008. *Designing Deliberative Democracy*. Cambridge: Cambridge University Press.