

Domination and Democratic Legislation

Sean Ingham* and Frank Lovett†

December 16, 2021

Abstract

Republicans hold that people are unfree if they are dominated, that is, if others have an insufficiently constrained ability to frustrate their choices. Since legislation can frustrate individuals' choices, republicans believe that the design of legislative institutions has consequences for individual freedom. Some have argued that if legislative institutions are democratic, then they need not be sources of domination at all. We argue this view is incorrect: the introduction of legislative authority, even if democratically organized, always creates a new site of domination. However, republicans can defend democratic procedures as the best means of minimizing the degree to which citizens are dominated, subject to the constraint of equalizing everyone's freedom. We formulate and prove this claim within a simple model of legislative authority and domination.

Keywords: republicanism, freedom, domination, majority rule, democratic legislation.

*Department of Political Science, University of California San Diego. Email: singham@ucsd.edu.

†Department of Political Science, Washington University in St. Louis. Email: flovett@wustle.edu.

Republicans hold that people are not free to the extent that they are subject to domination. People are dominated, on the republican view, if others have an unconstrained ability to frustrate their choices.¹ One way to frustrate people's choices is to change the rules governing those choices. Call the ability to change the rules that govern some relationship or association *legislative authority*. Two important questions arise for republicans. Considering any group whose members generally respect the rules governing their mutual affairs:

1. Is it possible for any subset of the group, including the subset that contains all the group's members, to have legislative authority over a person's choice without dominating that person's choice?
2. If not, how might the domination introduced by legislative authority be minimized, and, when minimized, do the benefits of legislative authority justify the costs to freedom as non-domination?

In this paper we argue that, contrary to the view of many republicans and others, the answer to the first question is no. Legislative authority always entails some degree of domination. In response to the second question, however, we argue that some institutional arrangements produce less domination than others, and the value of freedom as non-domination, together with equal concern for everyone's freedom, may give republicans reasons to prefer democratic institutions over alternatives. To support these claims, we employ a formal model to creatively reconstruct a republican argument for democracy as that institutional arrangement which minimizes domination, subject to the constraint of equalizing freedom from domination.

There are, of course, many reasons to value democracy. Even if the republican argument fails, those other reasons might tell decisively in favor of democratic government over any

¹Among other republican writings maintaining these claims, see especially Pettit (1997, 2012), Skinner (1998), Lovett (2010), and many of the papers collected in Laborde and Maynor (2009) and Elazar and Rousselière (2019).

feasible alternative. In this paper, however, we restrict attention to the specifically republican claim that a commitment to freedom from domination entails that legislative authority must be organized democratically.

Preliminaries

All republicans maintain that people are not free to the extent that their choices are dominated by other persons or groups. On the standard republican view, we may define domination as follows:

Domination: A 's choice whether to ϕ is dominated to the extent that some B has an unconstrained ability to frustrate that choice.

Here both A and B are supposed to be agents—either individual persons, or suitably organized groups of persons.² Intuitive examples of frustrating a choice might include physically restraining a person from exercising an option, or issuing a credible coercive threat to punish her in the event that she does. Whether any one agent has the genuine ability to frustrate the choices of another in these or other ways will depend only to a limited extent on the native abilities and resources of those agents. Rather, the various abilities agents have vis-a-vis one another will to a large extent depend on the background structural environment of their relationship. For example, a group of people with legislative authority will have an ability to frustrate the choices of those who are subject to the rules they enact only if other actors—bureaucrats, judges, and so on—help to enforce those rules. For present purposes, however, we mostly abstract from such issues.

²Following List and Pettit (2011) we hold that groups can count as agents when there are institutional mechanisms in place to ensure that the intentional states of the group's members are coherently aggregated on the one hand, and that the group's members are sufficiently motivated to perform a profile of actions suitably answering to those aggregated intentional states on the other.

Supposing that one agent does indeed have the ability to frustrate a given choice of another, under what conditions might that ability be suitably constrained, such that it does not constitute domination? There has been some debate among republicans concerning how to define these conditions precisely.³ The details need not detain us here, however. Instead, we will use a formal characterization of those conditions that is reasonably neutral across those debates, and which can be explained as follows.⁴

Central to the republican view is the idea that an agent can dominate our choices even when that agent is benevolent or indifferent: simply having the ability to frustrate our choices, when that ability is not suitably constrained, is sufficient to undermine our freedom. Accordingly, let us imagine the range of all possible types of potential agents—some benevolent, some malevolent, some indifferent.⁵ The range of all possible types is exceedingly large; in ordinary practice, however, only some are relevant. For example, if two people live together, one may have the ability to physically restrain the other and frustrate their choices about whether to leave their house, and there is a conceivable type who is so maniacally bent on doing so that he would exercise this ability even knowing he will suffer severe legal sanctions as a result. But for all practical purposes the possibility of such a type can be ignored when the rule of law is effective and criminal sanctions for such behavior are sufficiently reliable. For republicans, the possibility of that type is not on par with the possibility that the heretofore benevolent slave owner is the type who (facing no legal constraints) would choose to exercise his ability to restrain the movement of his slave.

The argument we make here is compatible with a range of views about what makes a type “ignorable” or not, but for the sake of concreteness we will proceed with the following provisional definition. Let us say that a given type of an agent is **ignorable** if its becoming

³Among others, see Costa (2007); Lovett (2012); Arnold and Harris (2017).

⁴This formalization follows Ingham and Lovett (2019).

⁵In the case of group agents, these types will correspond to profiles of the types of their constituent members, aggregated according to the internal procedures governing that group: see note 2 above.

common knowledge that the agent does not belong to that type would have no significant practical consequences. Put another way, since everyone already acts in their day-to-day encounters with others as if the agent is not one of the ignorable types, the common revelation that the agent is indeed not such a type will make no practical difference. Now we can formally characterize suitable constraint as follows:

Suitable constraint: The ability of one agent to frustrate the choice of another is suitably constrained if all the types of the former for whom it would be rational to frustrate that choice are ignorable.

Because this is not the most familiar way of describing republican views, a few quick examples should help clarify.

1. *Benevolent slave master:* Suppose that A is a slave and B is her master. Under familiar conditions of the slave relationship, masters will have numerous instruments by which they might frustrate the choices of their slaves. Since the cost of exercising these options is generally low, non-ignorable types of masters will rationally avail themselves of those instruments whenever they care enough about a particular choice. This will be true even if B happens to be benevolent, and even if A knows that B is benevolent. Thus A is dominated.
2. *Law and order:* Suppose that A is an ordinary citizen, and B a potential robber. In the absence of effective law, A will not be free to travel across town because non-ignorable types of B will intervene to rob her: A and others would act differently if it became common knowledge that B is not a type who would choose to do this. Supposing effective laws against robbery, however, A 's choice whether to travel across town will not be dominated. Only ignorable types of B will risk capture and punishment by attempting to rob her; in places where the state provides for law and order, its becoming common knowledge that B is *not* the type who is willing to suffer legal sanctions in

order to rob A will not change anyone's behavior. Thus effective law can provide a suitable constraint on B 's ability.

3. *Freedom of speech*: Suppose that A wants to publish an editorial critical of the government, and that B is a public official with the power to arrest her. If there are laws protecting freedom of speech, and democratic institutions that enable the people to punish public officials who violate those laws, then B 's ability to frustrate A may be suitably constrained: only ignorable types of the public official will risk punishment by arresting A for publishing her article.

Using this framework, let us now describe the central problem to which this paper is addressed.

In the law and order and free speech examples, the abilities of some agents to frustrate the choices of others are constrained by rules of law. We suppose in these examples a context where the system of rules is generally effective, such that only ignorable types would deviate from what the public rules require. That context might, for example, be a healthy democratic polity with a robust rule of law.

The difficulty is that no system of rules can remain static indefinitely. The social and political environment is constantly evolving. Perhaps rules that served in the past no longer accomplish what they were designed to accomplish. Or perhaps entirely new and unanticipated problems emerge that require new rules. Not only will there be good pragmatic reasons for wanting to be able to change the rules, there will also be specifically republican reasons. New social formations can give rise to new forms of domination that cannot be addressed without new rules. For example, protecting the rights of workers to organize unions might have reduced domination in traditional workplaces, but with the emergence of new corporate structures those policies may no longer help.

Let us grant, then, that we might want someone to have the ability to change the rules.

Call this ability **legislative authority**. The challenge for republicans is that legislative authority might itself be an instrument some agents can use to frustrate the choices of others. Consider again the example of free speech. Suppose that everyone follows the rules, whatever they are. Is *A* free to publish her editorial? As things stand in the example, she is not dominated by the public official *B*. Nevertheless she might be dominated by some other agent *C*, if *C* has the ability to change the rule so that *A*'s freedom of speech is no longer protected.

It is important to stress, however, that we cannot know in advance whether it is *changing* the rules or *failing to change* them that presents the relevant danger. If all cases were like the freedom of speech case, in which the danger lies in a possible change to the rules, then we could easily minimize domination by making legislation extremely difficult or, at the limit, impossible.⁶ But many cases will not be like this. The ability to block revisions to the rules can be an ability to cause individuals' choices to be frustrated. The causal effect might be indirect, and yet still intentional—think of segregationist legislators who blocked civil rights legislation and thereby indirectly, but deliberately, sustained patterns of behavior that frustrated African-Americans' choices. And this situation can arise, not simply from past injustices that led to the adoption of unjust rules (as in the example of the American South), but also from unanticipated technological developments or other social changes that alter the distribution of abilities, creating new opportunities for some actors to frustrate others' choices unless the rules are revised. However perfect the initial system of rules, it will need revision at some point. If all cases were like these latter ones, then we could minimize domination by making legislation as easy as possible. But of course both kinds of cases will arise in practice, so neither of these simplistic solutions will be satisfactory.

Accordingly, the central problem is best posed as follows: given that both the enactment

⁶Hayek (1973), for instance, seems to assume it is legal change that usually poses the danger, hence his antipathy towards legislation.

and the blocking of new legislation can cause individuals' choices to be frustrated, is there a way to design legislative authority so that the abilities of different groups to influence decisions on legislation will not themselves be sources of domination?

Some have argued yes. Richard Bellamy, for example, says that when legislative authority “emanates from . . . conditions of equality, it will not dominate” (2007, 160). Similarly, Clarissa Hayward defines non-domination as “that state of power relations in which all are socially enabled to participate effectively in making and remaking the institutions, the policies, the laws, and the other social norms that define their terms” (2011, 478).⁷ Intuitively, the suggestion is that people are dominated only when others have a greater ability to frustrate their choices than vice versa: there can be no domination among equals. And when legislative authority is shared equally, no one has a greater say what the rules will be than anyone else.

In the next section, we show why this is not correct with the help of a model. Then we consider whether Pettit's account of legislative authority avoids the problem, and conclude it does not, or at least not given the usual understanding of that account. Finally, in the last two sections, we use the model to characterize the decision-making rules that would minimize domination subject to a requirement that they equalize the degree of domination everyone expects to suffer.

Legislative authority and domination

We present a largely informal statement of our model in this section; formal definitions can be found in the appendix. Suppose we have a group of individuals who are to share legislative authority. At issue is the design of a procedure through which the rules governing the group's affairs can be changed. To simplify, assume that we are evaluating different procedures for

⁷Similar views are expressed in Bohman (2008) and Kirby (2016), among others.

making binary decisions on proposed changes to the rules governing the group.

Given this assumption, there are two possible decisions in a given case, rejecting or approving the proposed change. Each member of the group can be one of two possible types: they might prefer for the proposal to pass or they might prefer for it to fail. A *profile* of types is simply a list of the types of each member of the group. For example, if a group had just three members, then an example of a type profile would be the list in which the first individual prefers that the proposal pass while the other two prefer that it fail.

A **decision procedure** is a function that maps profiles of types to a collective decision. For example, simple majority rule can be represented by the function that maps a type profile to the decision “approve” if and only if the list has a majority of the individuals preferring for the proposal to pass. Another example would be unanimity rule, represented by the function that maps all type profiles to the collective decision “reject” except for the profile in which everyone’s preference is for the proposal to pass. Note that we use “decision procedure” as a term of art that picks out a concept more abstract than the ordinary connotations of the phrase might suggest. For example, the “decision procedures” in actual democracies are very complex, involving overlapping processes of public deliberation, campaigning, agenda-setting, voting, legal appeal, contestation, and protest. Again for the sake of simplicity, we abstract from these various features of democratic decision-making so as to focus on the relative contribution each individual can make, through some combination of these or other processes, to the collective decision. Individuals choose how to contribute to these processes as a function of their dispositions to support or oppose the proposed decision (their “types”). The functions we call “decision procedures” are descriptions of how individuals’ attitudes, by determining their choices of how to contribute, determine the collective decision taken.

As noted earlier, both the ability to change the rules, and the ability to block such changes, can be sources of domination. A proposed change might either create new social

patterns of behavior that frustrate a person's choices—the adoption of Jim Crow legislation in the American South, for example—or it might eliminate or prevent patterns that have this effect—the adoption of civil rights legislation. A republican analysis of legislative authority should not presuppose that the group faces one type of proposal rather than the other. Decision procedures are put in place, and often entrenched with constitutional rules, in order to handle an indefinite stream of unknown future proposals. Accordingly, we adopt the perspective of an institutional designer who wishes to minimize *ex ante* expected domination, where the expectation is a weighted average of the two possibilities.⁸

We assume that republicans will want any decision procedure to satisfy at least two minimal conditions. The first is that the entire group actually have some ability to revise its rules: it cannot be the case that the decision on a proposed revision is determined independently of the preferences of the group's members with respect to that proposal. We formalize this as:

Weak Pareto condition: if everyone supports (or opposes) a proposal, then that proposal passes (fails).⁹

The condition might fail if, for instance, constitutional provisions render a particular rule change impossible. There may be good reasons for protecting certain basic liberties in this way, but in such cases the issue of legislative authority simply does not arise. Our present concern is whether people are dominated specifically in those choices over which the legislative authority has competence to legislate.

A second condition republicans will wish to impose is that no one in the group be able to unilaterally decide whether a given proposal will be adopted (or rejected) irrespective of

⁸Some proposals might have neither effect, but the power of groups to implement such neutral proposals would not affect the degree to which anyone is dominated, so we can simplify the analysis by stipulating that neutral proposals arise with probability zero.

⁹The weak Pareto condition is stronger than we need: all the arguments would go through if we instead assumed merely that for each of the two possible decisions, there is some type profile that produces that decision under the decision procedure. But the weak Pareto condition, which entails this stronger property, is familiar and has a more natural interpretation and normative justification.

the preferences of all the group's other members. Decision-making power must be dispersed.

Weak Dispersion: for every individual in the group, there are type profiles at which the proposal is approved independently of the individual's preference, and type profiles at which it is rejected independently of the individual's preference.

Decision procedures that violate this condition include dictatorial rules, which empower a single individual to decide all issues, but some non-dictatorial rules as well. For example, requiring unanimity to adopt changes to the group's rules violates the condition. Under this arrangement, each individual would control the ability of the rest of the group to to adopt new rules, and so *that* ability would not dominate anyone. But when the existing rules already permit or sustain patterns of social behavior that frustrate individuals' choices, such as Jim Crow segregation laws, each individual would have an uncontrolled ability to perpetuate such patterns by unilaterally vetoing any proposal to eliminate them, and that ability *would* constitute domination.¹⁰ Our condition excludes both possibilities. We refer to it as "weak" because it does not rule out the concentration of decision-making power in the hands of small groups. Obviously republicans may wish to exclude concentrations of power in small groups, but for our argument the logically weaker condition suffices.

Republicans have reasons to accept each requirement, and yet, if a decision procedure satisfies the Pareto and dispersion conditions, then it will be hard, perhaps impossible, to escape the conclusion that it produces domination. To see why, suppose the group is considering a legislative proposal that would frustrate some of an individual's choices if adopted, and suppose that power is dispersed so that this individual cannot single-handedly determine whether the proposal passes or not. Then there exists at least one type profile at which the proposal will pass and the individual's choices will be frustrated even though the individual prefers for the proposal to fail. Strictly speaking, that fact by itself does not imply that the group has *an ability to frustrate* his choices: it is possible that the decision

¹⁰And accordingly, Pettit (2012, 168) rejects unanimity requirements.

procedure is one under which the proposal passes even if everyone is opposed to it (as in a world where some other agent can enact revisions to the rules, irrespective of what the members of the group think). But if the decision procedure is weakly Paretian, then the frustration results *because* at least some other members of the group wish to see the proposal pass. If instead they all opposed the proposal, then the decision procedure's being weakly Paretian implies that the proposal would be defeated. Thus, it must be that the group has an ability to frustrate the individual's choices by enacting the proposal under consideration.

And to reiterate, the individual's choices will be frustrated in this scenario *irrespective* of whether he supports or opposes the proposal, so he cannot in any relevant sense be said to *control* the ability of the other members of group to frustrate his choices. Thus he is dominated, unless that ability is adequately constrained in some other way. On the view we have proposed, it would be adequately constrained if the profiles of types of the other members of group who would exercise this ability—that is, the ability to frustrate his choices by enacting changes in the rules governing the group—are ignorable. In the case of certain legislative proposals, the relevant type profiles might be ignorable: for example, we can perhaps ignore the possibility of the group deciding to repeal laws against murder or assault. But over an extensive range of open political questions this will not be the case.

A more concrete illustration may help. Assume a group with three members, *A*, *B*, and *C*. *A* is contemplating a costly investment that will pay off if and only if she is permitted to use some controversial technology in the future. For example, the technology might be a ride-sharing app like Uber or Lyft, and her potential investment might be the purchase of a new car. Suppose the current rules governing the group permit her to use the technology, if only by the absence of any explicit mention of the topic. Is her choice whether to make the costly investment dominated or not?

Obviously it would be if either *B* or *C* were a dictator, but we have ruled out this scenario. Consider instead the case where *B* and *C* jointly share legislative authority. Specifically, if

either one proposes a ban on the technology, and both ultimately approve, then the rules will be revised accordingly. Under natural assumptions about the political process, the legislators B and C will have the joint ability to frustrate A 's choice of whether to invest in the technology: if each cares enough about banning the technology, then B and C will each take actions (proposing the ban and voting for it, for instance) that result in A 's choice being frustrated. Nor is this ability suitably constrained. The pairs of possible types for whom it is rational to undertake this intervention are not ignorable: if it became common knowledge that either B or C prefer not to ban the technology, or at least do not care enough to propose and vote in favor of a ban, then people would behave differently. In particular, A would invest in the technology. Thus, giving legislative authority to B and C results in their domination of A .

Here it is important to remember that, on the republican view, whether an individual is dominated does not depend on whether any of her choices are actually frustrated, but rather on whether anyone has the *ability* to frustrate those choices. As Pettit says, “to the extent that I have the uncontrolled power of interfering” in your choices, you are dominated “even if I do not actively impose my will” (2012, 59). If there exists a possible type profile for everyone else in the group such that, if they had those types, they would choose to frustrate the individual's choices through a revision to the rules, then the individual is dominated.¹¹ Or at any rate, she will be dominated provided those possible type profiles are not strictly *ignorable*, as in this example.

Significantly, there is also a sense in which B might unilaterally dominate A and C (and likewise, C might dominate A and B). Suppose the status quo includes a prohibition on the technology. Since a repeal of the ban requires approval from both B and C , by refusing

¹¹Simpson (2017) worries that it might follow we are dominated by an unlimited number of logically possible groups in society. Here we follow Ingham and Lovett (2019) in holding that possible groups can only dominate provided they face no unsolved coordination problems. Standing legislative procedures in effect remove coordination problems for the empowered legislators.

to approve the repeal, B frustrates A . And if we further suppose that C wants to lend A money to finance her investment, then C 's choices are frustrated by B as well. In other words, when some members of a group are excluded from a share of legislative authority, not only are those excluded members exposed to domination, so too are the members of those potential winning coalitions that might have included those excluded members.¹²

These cases are not very controversial, however, because the patterns of domination introduced by legislative authority are asymmetrical: C dominates A and B , or B and C jointly dominate A , and so on. The asymmetry results from an unequal participation in legislative authority. But what of the case where everyone shares *equally* in legislative authority? Consider, in particular, a simple majority rule under which, if anyone proposes a ban on the technology, and at least two individuals ultimately support that ban, the rules will be changed accordingly. This is the most promising case for republicans. As previously observed, many have tried to argue that when legislative authority is shared equally by all members of the group, there is a plausible sense in which no one in the group is dominated by the joint exercise of that ability. On our view, this conclusion cannot be drawn: the existence of legislative authority necessarily introduces at least some domination. Even when legislative authority is equally shared in our example, every potential subgroup of two dominates the remaining individual.

Consider A 's choice whether to invest in the new technology. Acting together, B and C can frustrate A 's choice by proposing and supporting a ban on that technology. So long as the issue of whether to ban the new technology is an open topic of political debate in the group, the possibility that B and C are types who might bring about such a ban will not be ignorable, and thus A 's choice is dominated. Exactly the same thing can be said about B and C with respect to the corresponding majority coalitions (A, C) and (A, B) . Thus, under simple majority rule *everyone* is dominated and *everyone* participates in the domination of

¹²For related points about popular control, see Ingham (2019), §7.3.

everyone else. Observe that the conclusion here follows even if we grant that domination implies inequality, for while there is equality between any two individuals in the model, there remains an inequality between the individual who is dominated and the group of individuals that dominate her.¹³ Democratizing legislative authority does not eliminate domination, it merely distributes that domination equally.

And this conclusion holds more generally. Suppose all we know about a decision procedure is that it satisfies the weak Pareto and dispersion conditions. Since decision-making power is dispersed, A cannot unilaterally block legislative revisions to the rules that would ban the new technology, so there is some possible type profile such that the proposal passes even though A prefers for it to fail. But since the decision procedure satisfies the Pareto condition, the reason the proposal passes in this scenario must be because either B or C or both wish for it to pass: if everyone preferred for it to fail, then weak Pareto would imply that it fails at this type profile. And the explanation for why the proposal passes, despite A 's opposition, when B or C or both wish it to pass, but fails otherwise, must be that B and C have an ability to enact the proposal. That is, they must have actions available to them that they will choose if one or both prefer for the proposal to pass, and that result in its passage despite A 's opposition.

If the pairs of types who would exercise this ability are ignorable, then their ability does not imply domination. But in this particular example—preferences over prohibitions on a new, socially disruptive technology—and many others, there is no basis for assuming the relevant type profiles are ignorable. So A is dominated, and by interchanging the roles of A , B , and C we can conclude that everyone is dominated. The conclusion follows simply from the assumption that the decision procedure satisfies the weak dispersion and Pareto conditions.

Finally, observe that it would not help if the members of the group—in the spirit of

¹³However, see Schmidt (2018) for reasons to doubt that domination implies inequality.

Rousseau’s social contract—were willing to embrace the outcome of the decision process as expressing their own “will” even when they had previously opposed the decision of the majority.¹⁴ On the republican view, we are dominated whenever others have an unconstrained ability to frustrate our choices, whether or not we subjectively identify with their decisions to frustrate our choices. Consider an analogy: the slave is still dominated by his master even if he reliably sees his own will reflected in his master’s decisions, for his master has the ability to interfere with the slave’s choices irrespective of whether the slave identifies his will with that of his master. As Pettit insists, security against domination must be “invariant across relevant alterations *both in your own will* and in the will of others as to what you should do” (Pettit 2012: 67, emphasis added).

Pettit and the Tough Luck Test

In the previous section we used a stylized model to show that the existence of legislative authority necessarily introduces at least some domination.

However, there may be a sense in which one decision procedure can produce more or less domination than another. In this section we consider Philip Pettit’s republican theory of legislative authority. As an argument attempting to demonstrate that democratic legislation does not dominate, the argument fails. But, we show, creatively reconstructed as an argument that democratic legislation *minimizes* domination as far as possible, given other normative constraints, the argument might succeed.

Pettit argues roughly as follows.¹⁵ Consider a large contemporary polity, characterized by reasonable pluralism but also enjoying an effective rule of law. Legislative authority in this polity will not introduce domination, on his view, if the decision process is:

1. *Individualized*: Members of the polity all have an equal opportunity to influence out-

¹⁴See De Dijn (2018) for a recent interpretation of Rousseau along these lines.

¹⁵The central argument described here can be found in Pettit (2012: 160–179).

comes of the legislative process through an open process of public deliberation, voting, contestation, and so forth.

2. *Unmediated*: Outcomes of the legislative process are determined *only* by the equally-weighted inputs of citizen deliberation, voting, and contestation, and thus are not mediated by any exogenous distorting influences (such as multinational corporations or foreign governments).
3. *Efficacious*: The members of the polity have lived together and participated in shared democratic governance long enough to develop a shared sense of the bounds of reasonable political disagreement, which bounds sufficiently constrain the range of legislative discretion that no individual need regard political outcomes as the intrusion of an alien will.

Pettit claims that when these conditions are met, no one in the polity will be exposed to domination in virtue of the citizens' shared exercise of legislative authority over themselves: "if the citizenry control state discretion in a suitable manner," he says, "then the imposition of a social order on those citizens will not take away from their freedom" (2012: 160).

Suppose, for example, that the issue of whether to permit ride-sharing apps lies within the bounds of reasonable political disagreement. Provided everyone has a fair opportunity to participate in the debate and influence the outcome through voting and other means of contestation, and that the process is not pre-empted by sinister corporate influences, the fact that a majority could decide to implement a ban on the technology does not itself dominate anyone, he claims:

The assumption is that while such process constraints will leave the final determination of policy in any area to the luck of the draw—including, for example, the draw that gives one or another side a majority in a legislature or referendum—that mode of determination can be consistent . . . with popular control. If it allows

the wills of some to rule on certain matters, . . . it does so only because such restricted empowerment is consistent with all that popular direction and control requires; like the power that I enjoy when you give me the key to the alcohol cupboard, the power given to the majority party can only be exercised within accepted bounds and need not count as dominating. (2012: 175–6)

But it is not immediately obvious why it need not count as dominating. His individualization and unmediation conditions are simply stronger versions of our weak dispersion and weak Pareto conditions, respectively. In our earlier illustration even his stronger versions of those conditions were met, and we nevertheless found that everyone is dominated, and that everyone participates in the domination of everyone else.

The burden of the argument thus falls on Pettit’s third condition, which effectively requires (in our terms) that most type profiles are ignorable. But even if we suppose the bounds of reasonable disagreement are quite narrow, why should we say that people are *not dominated* with respect to those specific choices falling inside the range of active political disagreement, where the possibility that majorities might form to frustrate those choices cannot be ignored?¹⁶ According to Pettit,

The point of legitimacy is to ensure that you and your fellow citizens are not subject to an alien, controlling will, despite the fact that there may be a good deal of discretion exercised by those in power. Such legitimacy will be adequately ensured . . . to the extent that you and your fellows have good grounds to think that any unwelcome results of public decision-making are just *tough luck* . . . [rather than] the sign of a malign will at work against you (Pettit 2012: 177, emphasis added).

In other words, provided the bounds of reasonable disagreement are sufficiently narrow and

¹⁶Several authors have raised this issue previously: see for example, Oliver (2013), Sharon (2016), or Kolodny (2019).

are settled on in the right way, people can regard specific outcomes within that range as effectively random—as merely “the luck of the draw” in his words—and people are not dominated by random processes.

There is an important difficulty with this, however. On the republican view, people are dominated whenever others have an unconstrained *ability* to frustrate their choices, whether they exercise such abilities or not. Suppose that a master leaves it to chance—a flip of the coin, say—whether he will frustrate his slave’s choices any given day. Whether the slave finds his choices frustrated will thus be random, but surely we would nevertheless insist that he is dominated. This is because the master’s ability is constrained by nothing other than his own whim to persist with the coin-flipping, which might change at any time. Likewise, while it might be the case that democratic political outcomes often exhibit an element of randomness, nothing prevents a determined majority from emerging and using legislative procedures to deliberately frustrate the choices of some minority. Or at any rate, nothing does within the scope of political questions that remain open.

Importantly, however, we should not conflate the question whether institutions give majorities unconstrained abilities to frustrate individuals’ choices, with the question whether such institutions are justified all things considered. When we ask whether individuals are dominated under democratic legislative institutions, we are asking the first question, not the second.

An alternative approach is for republicans to argue that while our choices are dominated by the existence of legislative authority, to the extent that such authority meets the right conditions, the amount of domination involved is sufficiently small that it is a price worth paying for the benefits of having legislative authority in the first place. The “tough luck test” republicans would need to formulate, then, is not one indicating *whether* legislative authority introduces domination, but rather one indicating whether it is so designed that the domination it introduces falls below this acceptability threshold.

Our elaboration of this suggestion will proceed in two steps. First, we formalize a measure of domination that will allow us to make interpersonal comparisons of domination under a given procedure, as well as intrapersonal comparisons of domination across different procedures. Second, we show how democratic decision procedures minimize domination according to that measure, among those rules that equalize everyone’s degree of domination.

Measuring Domination

Intuitively speaking, the greater the range of possible types who would choose to exercise their ability to frustrate my choices, the greater my domination. To illustrate the intuition, compare two hypothetical versions of an academic department: in the first, the twenty tenured members of the department must all vote in favor of B ’s case when she goes up for tenure if she is to receive it; if just one votes against, she is denied tenure. In the second, only a simple majority need vote in favor in order for her case to be approved.

Let us suppose that the rules and culture of the university are such that the faculty’s collective ability to end B ’s employment is unconstrained—i.e., the type profiles who would exercise this ability are not ignorable. Republicans should therefore conclude that B is dominated by the tenured faculty, under either voting rule. But they may plausibly judge the degree of domination to be less under simple majority rule as compared with the unanimity rule. Comparing the two voting rules, the set of possible type profiles for the tenured faculty members that trigger tenure denial under simple majority rule is a strict subset of the set of possible type profiles that trigger denial under the unanimity rule. Because one set is a strict subset of the other, there is straightforward sense in which one of these sets is “larger” than the other; in this sense, we may say that B ’s choices are frustrated in a “greater range” of relevant counterfactual worlds when the unanimity rule is used.

The republican explanation for why this fact should matter, and why we might express

it by saying that B is dominated to a lesser degree under the simple majority rule, cannot be that frustration occurs with higher probability under the unanimity rule. The probability of frustration will not be lower and will almost surely be higher under the unanimity rule, where a single “no” vote triggers a denial of tenure. But a central tenet of republican political theory is that the benevolent master dominates his slave even if the probability that he will frustrate his slave’s choices is low, indeed even if it is effectively zero: for even in that case, he retains the *ability* to do so.

Judgments with a formal structure similar to probabilities may nonetheless have a role to play in the analysis. Here is one possible explanation for why domination is higher under unanimity rule, which illustrates this point. Under normal conditions, the possibility that the tenured faculty might end B ’s employment will have *greater public salience* if the faculty use unanimity rule, compared with simple majority rule. When a single “no” vote is decisive, a cloud will hang over all of B ’s encounters with senior colleagues. Even with colleagues who she is certain would not vote against her case, it will be common knowledge between them that if the senior colleague *did* wish to terminate her employment, he could do so with impunity. That commonly known fact will color their relationship, whatever the objective probability of the senior colleague’s voting against, and whatever subjective probability B attaches to that event. Under simple majority rule, by contrast, the fact that the tenured faculty have a collective ability to end her employment will not have the same public salience. One way to formalize these judgments about salience is to think of them as the probabilities a hypothetical observer would attribute to the group’s exercising its ability, if the observer knew only those facts that are common knowledge among the relevant actors. Thus, the observer would know what the voting rule is, but not all the interior beliefs and dispositions of the tenured faculty.

Other explanations may be available to republicans, and the argument we wish to make does not require accepting exactly the interpretation proposed here. Our argument relies only

on a few more abstract assumptions about the manner in which republicans might conceptualize degrees of domination. We will state those assumptions in their abstract generality, and illustrate their meaning with the more concrete interpretation we have sketched.

We assume the degree to which one is dominated is a *measure* of the set of possible types (for the remaining members of society) who would rationally choose to exercise their uncontrolled, unconstrained abilities to frustrate one's choices. The measure assigns non-negative numbers to every such set of possible types; it assigns a measure of zero to the empty set; and for two mutually exclusive sets of possibilities, the numbers it assigns to them add up to the number it assigns to the set consisting of their union.¹⁷ Under these assumptions, we are dealing with a measure in the technical sense of measure theory; other examples of measures include the integral as a measure of spatial area, or the probability measures that describe long-run frequencies or subjective beliefs. Sticking with our suggested interpretation, the measure of the set of possible types could be interpreted as the probabilities that a hypothetical observer would assign to these events if the observer knew only those facts that are common knowledge among the members of society.

To capture the commitments of republican theory, we make two additional assumptions about the domination measure:

- (A1) The measure assigns strictly positive numbers to every non-empty set of possible types.
- (A2) The identities of the legislative proposal's supporters and opponents at a type profile do not matter when one takes the measure of a set containing just that single profile of types; only the number of individuals at the profile who support or oppose the legislative proposal can matter.

The first assumption is needed to capture the republican idea that even perfectly benevolent masters dominate. In the present context, where the type profiles in question describe

¹⁷See the appendix for a formal definition.

legislators' attitudes towards legislation, every logically possible profile of types is relevant when considering the degree to which the rest of the group dominates an individual through their possession of legislative authority. Taking recourse to our running interpretation, the idea is that a hypothetical observer, knowing only commonly known facts, would not assign probability zero to any possible profile of types.

To illustrate the rationale for the second assumption, consider a slave, B , who has two joint owners, A_1 and A_2 , either of whom can unilaterally frustrate the slave's choices. The degree to which the slave is dominated reflects some measure of the set of possible worlds in which one or the other owner cares enough to frustrate the slave's choices. Intuitively, the possibility that A_1 chooses to intervene but not A_2 contributes just as much to the slave's unfreedom—by way of its contribution to the measure of this set—as the possibility that A_2 chooses to intervene but not A_1 . In the absence of any extra information about their abilities or the constraints on their abilities, the relevant feature of a possible scenario in which the slave's choices are frustrated is not the identity of the intervening agent, but just whether an agent is intervening and (perhaps) how many agents are intervening or supporting the intervention.¹⁸

This second assumption may not always be justified. In the case of our running interpretation, it would be satisfied if the possibility of one legislative coalition forming and enacting or blocking new legislation has no greater public salience than the possibility of any other coalition, of the same size, doing the same. In some contexts, that assumption might not hold. For example, in a society where a racial majority has historically oppressed a racial minority, the possibility of racially homogeneous coalitions might have greater public

¹⁸In this example, the assumption leaves open the possibility, *but does not require*, that the scenario in which both agents intervene to frustrate B 's choices contributes more than, say, the scenario in which A_1 intervenes but not A_2 . More generally, the assumption is *not* that the measure of a type profile (technically speaking, the measure of the singleton set containing the type profile) be equal to the number of individuals' supporting the proposal at that profile, nor does it require that the measure increase as that number increases. The assumption is merely that the measure remains constant across different profiles if that number remains constant across those profiles.

salience than the possibility of racially diverse coalitions of the same size. This is obviously an important limitation to the results we present in the following section, which are derived on the basis of an assumption that excludes such cases.

Minimizing Domination

Having formalized a measure of degrees of domination, we can now compare different institutional arrangements according to the amount of domination they introduce.

As we explained above, new legislation does not always have the effect of frustrating individuals' choices. For every piece of legislation that has this effect, there is another legislative proposal to repeal it that has the opposite effect. And of course some legislation has the effect of preventing private citizens from frustrating other citizens' choices. Thus the goal of republican institutional design cannot simply be to make it as hard as possible to enact legislation. In one kind of situation, domination stems from an unconstrained, uncontrolled ability to enact a proposed piece of legislation; in the other, it stems from an unconstrained, uncontrolled ability to block a proposed piece of legislation.

The appropriate objective for the republican institutional designer (if we momentarily set aside values besides freedom from domination) should therefore be to minimize the *expected* degree of domination that a decision procedure produces, averaging over the degrees of domination it will produce in each of these two kinds of situations. The **expected degree of domination** is defined as the probability of the first kind of situation, in which legislation would have the effect of frustrating an individual's choices, multiplied by the measure of the set of types for the remaining members of the group who, whether the individual likes it or not, would exercise their ability to enact the legislation; plus the probability of the second kind of situation, in which blocking the legislation has the effect of frustrating the individuals' choices, multiplied by the measure of the set of types who, whether the individual likes it or

With probability Φ , passage causes frustration		With probability $1 - \Phi$, defeat causes frustration	
Type profile	Majority rule decision	Type profile	Majority rule decision
(y, y, y)	passes	(y, y, y)	passes
(y, y, n)	passes	(y, y, n)	passes
(y, n, y)	passes	(y, n, y)	passes
(n, y, y)	passes	(n, y, y)	passes
(y, n, n)	defeated	(y, n, n)	defeated
(n, y, n)	defeated	(n, y, n)	defeated
(n, n, y)	defeated	(n, n, y)	defeated
(n, n, n)	defeated	(n, n, n)	defeated

Table 1: Each table displays all the possible type profiles. Individual B 's type is listed second in the profiles. The type profiles whose possibility contributes to individual B 's domination under simple majority rule are in red. Whether the possibility of a type profile contributes to the measure of domination depends on whether the legislative proposal would cause B 's choices to be frustrated (left-hand table) or instead prevent B 's choices from being frustrated (right-hand side).

not, would exercise their ability to defeat the legislation.¹⁹

The following example illustrates the model. Assume a society of three individuals, A , B , and C , each a member of the legislative assembly which follows simple majority rule. At some point in the future, new legislation will be proposed (we bracket the admittedly important questions about the agenda-setting process and treat it as exogenous to the model), which might have either the effect of frustrating B 's choice, or the effect of securing B 's choice from frustration. Let Φ refer to the probability of the first kind of situation, in which legislation frustrates B 's choice, and $1 - \Phi$ refer to the probability of the other kind. Table 1 displays the first scenario on the left and the second on the right, and indicates (in red) which of the eight possible profiles of types contributes to the measure of B 's domination (B 's type is listed second in the type triplets).

In the first scenario, the profile (y, y, y) contributes to B 's domination because it describes a situation in which the group comprising A and C will exercise its ability to frustrate B by passing the legislation. In the second scenario, on the right, this type profile does not

¹⁹See the appendix for a formal definition.

With probability Φ , passage causes frustration		With probability $1 - \Phi$, defeat causes frustration	
Type profile	Unanimity rule	Type profile	Unanimity rule
(y, y, y)	passes	(y, y, y)	passes
(y, y, n)	defeated	(y, y, n)	defeated
(y, n, y)	defeated	(y, n, y)	defeated
(n, y, y)	defeated	(n, y, y)	defeated
(y, n, n)	defeated	(y, n, n)	defeated
(n, y, n)	defeated	(n, y, n)	defeated
(n, n, y)	defeated	(n, n, y)	defeated
(n, n, n)	defeated	(n, n, n)	defeated

Table 2: Each table displays all the possible type profiles. Individual B 's type is listed second in the profiles. The type profiles whose possibility contributes to individual B 's domination under unanimity rule are in red. Whether the possibility of a type profile contributes to the measure of domination depends on whether the legislative proposal would cause B 's choices to be frustrated (left-hand table) or instead prevent B 's choices from being frustrated (right-hand side).

contribute to the degree to which B is dominated because, in this second scenario, passing the legislation protects B from frustration. In the second scenario, where blocking the legislation results in B 's choice being frustrated, the possibility of the type profile (n, n, n) contributes to B 's domination. Note that this profile contributes to the measure of B 's domination even though it is a scenario in which the vote goes the way B desires. It is a scenario in which A and C exercise their ability to intervene, even though B happens to approve.

Now let us compare simple majority rule with a decision procedure under which unanimity is necessary and sufficient for passage, displayed in Table 2. Type profiles contributing to B 's domination are again indicated in red.

To assess whether B is dominated to a greater degree under the unanimity rule, we need to know the relative likelihood of the two scenarios, and how much weight to give to the different possible type profiles. As remarked earlier, from the point of the view of the institutional designer, we should assume either scenario is equally likely to arise over the long run (so assume $\Phi = 1/2$). One might also argue that, to capture the commitments of republican theory, the measure ought to give equal weight to each of the possible type

With probability Φ , passage causes frustration		With probability $1 - \Phi$, defeat causes frustration	
Type profile	Oligarchic rule.	Type profile	Oligarchic rule
(y, y, y)	passes	(y, y, y)	passes
(y, y, n)	defeated	(y, y, n)	defeated
(y, n, y)	passes	(y, n, y)	passes
(n, y, y)	defeated	(n, y, y)	defeated
(y, n, n)	defeated	(y, n, n)	defeated
(n, y, n)	defeated	(n, y, n)	defeated
(n, n, y)	defeated	(n, n, y)	defeated
(n, n, n)	defeated	(n, n, n)	defeated

Table 3: Each table displays all the possible type profiles. Individual B 's type is listed second in the profiles. The type profiles whose possibility contributes to individual B 's domination under the oligarchic rule are in red. Whether the possibility of a type profile contributes to the measure of domination depends on whether the legislative proposal would cause B 's choices to be frustrated (left-hand table) or instead prevent B 's choices from being frustrated (right-hand side).

profiles. Making both these assumptions, then, the degree of B 's domination under the unanimity rule is $1/2 \cdot 6/8 = 3/8$, while B 's degree of domination under simple majority rule is $1/2 \cdot 2/8 + 1/2 \cdot 2/8 = 1/4$, so simple majority rule is superior.

Finally consider an “oligarchic” rule under which B is disempowered and the proposal passes if and only if A and C both vote for it. Table 3 displays this decision procedure, with type profiles contributing to B 's domination again indicated in red.

It is easy to see that B is dominated to a greater degree under the oligarchic rule compared with simple majority rule. Every scenario in which B would suffer frustration under simple majority rule is also a scenario in which she would suffer frustration under the oligarchic rule; but there are additional scenarios where she would suffer frustration if the oligarchic rule is used but not simple majority rule. Given the assumption (A1) that the measure assigns positive weight to every possible scenario, it follows that the degree of domination is greater under the oligarchic rule—and indeed, is greater *regardless of the relative likelihood of the two scenarios*, and regardless of the weight the measure assigns to the different type profiles. This is a significant result for republicans, insofar as it formally demonstrates what should be

intuitively obvious, namely that inclusion in a decision procedure can only reduce a person's domination. And this likewise holds comparing the oligarchic rule with the unanimity rule: provided there is any possibility at all of the first scenario arising, B 's domination will be less under the latter than under the former, simply because B has not been excluded from the legislative process.

We can also make interpersonal comparisons of domination (although the color-coding in the tables only depicts the scenarios that contribute to the measure of B 's domination). For example, under the oligarchic rule, there are no scenarios that contribute to the measure of A 's domination when the legislative proposal would cause her choices to be frustrated, for in that case A can unilaterally prevent that frustration by voting against the proposal. As with the previous conclusions, these are independent of the assumption that future legislative proposals are just as likely to create as to eliminate sources of frustration.

What else can be said about the decision procedures that are best if one is concerned to minimize domination? To simplify the analysis, we now restrict attention to **voting rules**, by which we mean decision procedures under which proposals are approved if and only if the weighted sum of supporters exceeds the weighted sum of opponents, with possibly different weights attached to different individuals.²⁰ One can derive two conclusions about voting rules (theorems 1 and 2) from the assumptions (A1) and (A2), together with the generic properties of measures, and a third (theorem 3) if one makes a further assumption about the measure of domination.²¹

²⁰Recall, from above, that a "decision procedure," as we use the term, is just an abstract description of how decisions depend on the distribution of support for a proposal among the group. The functional dependence might be mediated by various actions, from casting votes to deliberating to protesting and so on. The same goes for "voting rules," as defined: these are functions describing how decisions depend on the distribution of support for a proposal, assuming they depend just on the weighted sum of support. The mechanism by which an individual's support for a proposal contributes a certain amount of "weight" towards its approval might be the literal casting of a weighted ballot, but it could also be more complicated and include non-voting behavior such as deliberating, protesting, and so on. To call these functions "voting rules" is thus to invoke a metaphor for the purpose of giving them an evocative label, not to give a literally descriptive name to them.

²¹From a formal point of view, theorems 2 and 3 are similar to results in the literature on measures of voting power (see Felsenthal and Machover 1998). Their substantive interpretation is quite different,

Theorem 1 is intuitive: everyone's preference must be given equal weight if the voting rule is to be consistent with everyone's equal freedom. Say that a decision procedure *preserves equal freedom* if, for any two individuals, their expected degrees of domination under the decision procedure are the same, and say that it is a *counting rule* if the outcome depends only on the number of votes cast in favor of the proposal, but not on the identities of those voting in favor (which implies that the weights attached to votes are, for all practical purposes, equal).

Theorem 1. A voting rule f preserves equal freedom if and only if it is a counting rule.

Simple majority rule as well as supermajority rules are counting rules.

Which decision procedure is best from the standpoint of minimizing domination? Recall again that we adopt the perspective of an institutional designer who does not know *ex ante* whether future legislative proposals will cause individuals' choices to be frustrated or instead prevent them from being frustrated, so the goal is not to make legislation as hard as possible. The next proposition characterizes the nature of the optimal decision procedure, among those procedures that preserve equal freedom, under the assumption that the two possibilities are equally probable.

Theorem 2. Fix any member of the legislature, and assume $\Phi = 1/2$. If a voting rule minimizes the individual's expected degree of domination among those voting rules that preserve equal freedom, then it maximizes the measure of the event in which the individual's vote is pivotal.

To understand the claim, note first that, having fixed an individual and comparing two decision procedures, the events in which the individual is pivotal are different for the two

however. Note, in particular, that the assumptions about the measure of domination are defended here by appeal to the substantive commitments of republican political theory, not by appeal to intuitions about how to measure power.

different procedures, and the measure of these events—the weights that the measure assigns to the events—may be different as well. In the three-person example from above, the event in which B is pivotal under simple majority rule is the event in which A and C feel differently about the proposal, with one supporting and one opposing, while the event in which B is pivotal under unanimity rule is the event in which A and C each support the proposal. Thus, the theorem says that if one wishes to minimize B 's expected degree of domination, simple majority rule is better if and only if the first event has greater measure than the second.

Keep in mind our assumption that $\Phi = 1/2$, i.e., legislative proposals are as likely to create as to eliminate frustrations in the individual's choices: that assumption is necessary for this result. As a stark example to illustrate the point, the optimal counting rule would be unanimity rule if $\Phi = 1$ and legislation were certain to cause rather than prevent the frustration of her choices, because this rule eliminates domination under this assumption. It need not maximize pivotality, however.

A second observation to bear in mind is that, having fixed a counting rule with a threshold k (that is, a procedure where the proposal passes if and only if at least k individuals support it), the events in which two different individuals are pivotal are each events in which exactly $k - 1$ of the remaining n voters support the proposal and $n - 1 - (k - 1) = n - k$ oppose it. Even though the two events will involve coalitions with different members, the two events must receive the same measure given the assumption (A2). Thus, one can fix any individual member of the group and, comparing voting rules from the perspective of that individual's expected degree of domination, the optimal rule will be the same, irrespective of the individual one picked.

We can say more about which voting rule is optimal if we make further assumptions about the weight the measure assigns to different possible scenarios. One natural assumption to consider is that the measure ought to assign the same weight to every possible scenario. For example the possibility that a minimum decisive coalition will exercise its power to frustrate

a person's choices by passing legislation contributes just as much to the overall measure of domination as the possibility that a coalition of everyone will exercise its power to do so, and so on.

Theorem 3. Fix any member of the legislature, and assume $\Phi = 1/2$. If the measure assigns the same weight to any two type profiles, then simple majority rule minimizes the individual's expected degree of domination among those voting rules that preserve equal freedom.

The result is a corollary of the previous theorem: under the assumption that the measure assigns the same weight to every type profile, the pivotality event associated with simple majority rule has greater measure than the pivotality events associated with any other voting rule.

Conclusion

A group with legislative authority, no matter how democratically constituted, necessarily dominates those choices over which it is formally and practically competent to legislate. That is not a reason for republicans to reject the legitimacy of legislation, however. Given that social and economic conditions are constantly evolving, new forms of domination will always emerge, and legislation will be required to address them. The relevant question for republican theory is thus how to minimize all forms of domination over time. Even if the democratic constitution of legislative authority does not eliminate its dominating character, it may minimize the degree of domination, compared with the relevant alternatives.

We have proposed a model to articulate a sense in which domination can come in degrees, and used it to establish three results that speak to this question. Republicans who are also egalitarians will want to equalize the degree to which people are dominated by the

existence of legislative authority, and our first theorem confirms the intuitive link between equal freedom and equal voting power. The second theorem says that, assuming that legislative proposals are as likely to create as to eliminate interference, the domination-minimizing decision procedure is the one that maximizes the measure of events in which the individual legislator is pivotal. And the third result, a corollary of the second, is that this optimal decision procedure is simple majority rule if our measure assigns equal weight to every possible configuration of types. In other words, unless there is some special reason to attach more weight to certain exercises of legislative interference than others in one's measure of an individual's overall degree of domination, or some reason to believe that future legislation will more often create rather than remove obstacles that frustrate an individual's choices, we cannot do better from a republican point of view than simple majority rule.

Acknowledgements

An earlier version of this paper was presented at the Manchester Centre for Political Theory Workshop in September 2021. The authors would like to thank the workshop's participants as well as this journal's editors and anonymous referees for helpful comments on previous drafts.

Author biographies

Sean Ingham is Associate Professor of Political Science at the University of California San Diego. Frank Lovett is Professor of Political Science at Washington University in St. Louis.

Appendix

We first give formal definitions of the elements of our model and then state and prove the theorems from the main text. Let $G = \{1, \dots, n\}$ be the set of individual legislators, with n odd (so that we can bracket the question of ties). Let $t_i \in \{0, 1\}$ represent individual i 's type, with $t_i = 1$ indicating that i prefers for the proposal to pass and $t_i = 0$ indicating i prefers for it to fail. A *profile* of types is simply a list of the types of each of the members of G , formally denoted $t = (t_1, \dots, t_n) \in \{0, 1\}^n$. Given lists $t, t' \in \{0, 1\}^n$, we will let t_i and t'_i denote the i th components of the list with matching superscript.

Let Σ be the set of all subsets of $\{0, 1\}^n$. A *measure* is a function $m : \Sigma \rightarrow \mathbb{R}$ that satisfies (i) $m(E) \geq 0$ for all $E \in \Sigma$, (ii) $m(E) = 0$ if $E = \emptyset$, and (iii) $m(E \cup F) = m(E) + m(F)$ for any $E, F \in \Sigma$ such that $E \cap F = \emptyset$. As explained in the text, we make the further assumptions that (A1) $m(E) > 0$ for all nonempty $E \in \Sigma$, and (A2) $m(\{t\}) = m(\{t'\})$ if $\sum_{i \in G} t_i = \sum_{i \in G} t'_i$.

A decision procedure is a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, mapping profiles of types to a collective decision (zero for “rejection,” one for “approval”). By a *voting rule*, we mean a procedure in which decisions are made by adding up (possibly weighted) preferences: f is a voting rule if there are numbers $w_i \geq 0$ for each $i \in G$, a number q such that $q > \sum_{i \in G} w_i/2$, and, for all $t \in \{0, 1\}^n$, $f(t) = 1$ if and only if $\sum_{i \in G} w_i t_i \geq q$. While the focus on voting rules is a significant restriction from a mathematical point of view, we make it to simplify the analysis and because the restriction seems less significant from a substantive point of view. A *counting rule* is a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ such that for some number $k \in \{1, \dots, n\}$, $f(t) = 1$ if and only if $\sum_{i \in G} t_i \geq k$.

Having fixed a voting rule, let $\mathcal{D}_1 := \{D \subseteq G : \sum_{i \in D} w_i \geq q\}$ refer to the set of coalitions that can unilaterally secure passage of the proposal, and let $\mathcal{D}_0 := \{D \subseteq G : \sum_{i \in G \setminus D} w_i < q\}$ refer to the set of coalitions that can unilaterally prevent passage. For a set $D \subseteq G$, let E_1^D

(and E_0^D) refer to the event in which all and only those members of D support (oppose) a proposal. For example, if $n = 3$ and the decision procedure is unanimity rule (e.g., $w_1 = w_2 = w_3 > 0$ and $q = w_1 + w_2 + w_3$), then $\{1, 2\} \in \mathcal{D}_0$ but $\{1, 2\} \notin \mathcal{D}_1$, and $E_1^{\{1,2\}} = \{(1, 1, 0)\}$ and $E_0^{\{1,2\}} = \{(0, 0, 1)\}$.

For an individual $h \in G$, the *expected degree of domination* under a decision procedure, as defined on p. 23, is equal to

$$\Phi \sum_{D: D \setminus \{h\} \in \mathcal{D}_1} m(E_1^D) + (1 - \Phi) \sum_{D: D \setminus \{h\} \in \mathcal{D}_0} m(E_0^D) \quad (1)$$

where, recall, Φ is the probability that the future legislation would have the effect of frustrating the individual's choice, and $1 - \Phi$ the probability that blocking the legislation would have this effect. For example, with unanimity rule and three individuals, the expected degree of domination for an individual $h = 1$ is

$$(1 - \Phi) \left[m(\{(1, 0, 1)\}) + m(\{(1, 1, 0)\}) + m(\{(0, 0, 1)\}) + m(\{(0, 1, 0)\}) + m(\{(1, 0, 0)\}) + m(\{(0, 0, 0)\}) \right],$$

because $\mathcal{D}_1 = \{G\}$ and $G \setminus \{1\} \notin \mathcal{D}_1$, so the first term in (1) just drops out; and, among the nonempty subsets of G , all are elements of \mathcal{D}_0 and only the set $\{1\}$ fails to satisfy the condition that $D \setminus \{1\} \in \mathcal{D}_0$.

An observation used below is that if the expected degrees of domination for two individuals h and j are the same under some voting rule then

$$\left(\frac{\Phi}{1 - \Phi} \right) \left(\sum_{D: D \setminus \{h\} \in \mathcal{D}_1} m(E_1^D) - \sum_{D: D \setminus \{j\} \in \mathcal{D}_1} m(E_1^D) \right) = \sum_{D: D \setminus \{j\} \in \mathcal{D}_0} m(E_0^D) - \sum_{D: D \setminus \{h\} \in \mathcal{D}_0} m(E_0^D) \quad (2)$$

Finally, say that a decision procedure f is *pairwise anonymous* if, for any two individuals

$h, j \in G$, $f(t) = f(t')$ whenever $t_h = t'_j$, $t_j = t'_h$, and $t_k = t'_k$ for all $k \in G \setminus \{h, j\}$.

Observation: if a voting rule f is pairwise anonymous, then it is a counting rule. As this observation is intuitive and easy enough to prove, we omit the proof here.

Theorem 1. *A voting rule f preserves equal freedom if and only if it is a counting rule.*

Proof of Theorem 1. To prove sufficiency, suppose the decision procedure is a counting rule with threshold k . Let m_x refer to the measure assigned to any singleton set $\{t\}$ such that $\sum_{i \in G} t_i = x$ (this quantity is well-defined because of the assumption (A2)). If $k < n$, then, for every individual $h \in G$, the expected degree of domination is just

$$\Phi \sum_{j=k}^{n-1} \binom{n-1}{j} (m_j + m_{j+1}) + (1 - \Phi) \sum_{j=0}^{k-2} \binom{n-1}{j} (m_j + m_{j+1}) \quad (3)$$

and if $k = n$, the expected degree of domination is $(1 - \Phi) \sum_{j=0}^{n-2} \binom{n-1}{j} (m_j + m_{j+1})$. (The derivation from (1) uses property (iii) of measures together with assumption (A2), which is implicit in the definition of m_x .) As these quantities are independent of h , we conclude the decision procedure preserves equal freedom. To prove necessity, suppose a voting rule f preserves everyone's equal freedom but (to derive a contradiction) does not satisfy pairwise anonymity. Then there are two individuals $h, j \in G$ and type profiles $t, t' \in \{0, 1\}^n$ such that $t_h = t'_j = 0$, $t_j = t'_h = 1$, $t_k = t'_k$ for all $k \in G \setminus \{h, j\}$, but $f(t) = 1$ and $f(t') = 0$. It follows that $w_j > w_h$. Construct the bijection $g : 2^G \rightarrow 2^G$ as follows (2^G refers to the set of all subsets of G): for each $D \in 2^G$, if $h \in D$ and $j \in D$, or $h \notin D$ and $j \notin D$, then $g(D) = D$; if $h \in D$ but $j \notin D$, $g(D) = \{j\} \cup D \setminus \{h\}$, and if $j \in D$ but $h \notin D$, $g(D) = \{h\} \cup D \setminus \{j\}$. For all $D \subseteq G$,

$$D \setminus \{j\} \in \mathcal{D}_x \implies g(D) \setminus \{h\} \in \mathcal{D}_x, \forall x \in \{0, 1\}$$

because $w_j > w_h$. Assumption (A2) ensures that the events E_x^D and $E_x^{g(D)}$, for $x \in \{0, 1\}$, have the same measure under m because $|D| = |g(D)|$. It follows that the left-hand side of

(2) must be at least as great as zero, and the right-hand side must be no greater than zero (informally speaking, for every set D such that the event E_x^D contributes to the measure of j 's domination, there is a unique set $g(D)$ that makes the same contribution to h 's domination). But the left-hand side must in fact be strictly greater than zero because $\{i \in G : t_i = 1\} \setminus \{h\} \in \mathcal{D}_1$, but $g(\{i \in G : t_i = 1\} \setminus \{j\}) \notin \mathcal{D}_1$. As this contradicts our assumption that f preserves equal freedom, we can reject the supposition that f does not satisfy pairwise anonymity. Since it satisfies pairwise anonymity, it is a counting rule, by the Observation above. ■

Theorem 2. *Fix any $h \in G$ and assume $\Phi = 1/2$. If f minimizes the expected degree of h 's domination among those voting rules that preserve equal freedom, then it maximizes (among these voting rules) the measure of the event in which h is pivotal.*

Proof of Theorem 2. Fix an individual $h \in G$, and assume $\Phi = 1/2$. Because f preserves equal freedom, it is a counting rule, by theorem 1. Setting $\Phi = 1/2$ and rearranging terms in (3), from the proof of theorem 1, the expected degree of domination for h is

$$1/2 \sum_{j=0}^{n-1} \binom{n-1}{j} (m_j + m_{j+1}) - 1/2 \binom{n-1}{k-1} (m_{k-1} + m_k)$$

Choosing k (which parameterizes the decision procedure) to minimize this expression is equivalent to choosing k to maximize

$$\frac{(n-1)!}{(k-1)!(n-k)!} (m_{k-1} + m_k)$$

which is just the measure assigned to the scenario in which the $n-1$ other legislators are split, with $k-1$ supporting and $n-1-(k-1) = n-k$ opposing the measure. Since k is the minimum number of yes-votes needed to pass the measure under a k -rule, this is just the scenario in which h is pivotal. ■

In theorem 3, by *simple majority rule*, we mean the function $f_M : \{0, 1\}^n \rightarrow \{0, 1\}$ defined by $f(t) = 1$ if and only if $\sum_{i \in G} t_i \geq \frac{n+1}{2}$, for all $t \in \{0, 1\}^n$.

Theorem 3. *Fix any $h \in G$ and assume $\Phi = 1/2$. If $m(\{t\}) = m(\{t'\})$ for all $t, t' \in \{0, 1\}^n$, then simple majority rule minimizes the expected degree of domination for individual h among all those voting rules that preserve equal freedom.*

Proof of Theorem 3. By theorem 2, the voting rule that minimizes the individual's expected degree of domination, among all voting rules that preserve equal freedom, is the counting rule corresponding to the number k that maximizes the expression

$$\frac{(n-1)!}{(k-1)!(n-k)!} (m_{k-1} + m_k)$$

By hypothesis, there is some number $m > 0$ such that $m_k = m$ for any integer $k \in \{1, \dots, n\}$. (This number must be positive because the measure assigns positive weight to every type profile, by assumption (A1).) Since the choice of k does not affect the second term in the maximand, the optimal k maximizes $\frac{(n-1)!}{(k-1)!(n-k)!}$. As is well-known, this expression, viewed as a function of k , attains its maximum when $k = \frac{n+1}{2}$, which corresponds to simple majority rule. ■

References

Arnold S and Harris JR (2017) What is Arbitrary Power? *Journal of Political Power* 10(1): 55–70.

Bellamy R (2007) *Political Constitutionalism: A Republican Defence of the Constitutionality of Democracy*. Cambridge: Cambridge University Press.

Costa VM (2007) Freedom as Non-domination, Normativity, and Indeterminacy. *Journal of*

Value Inquiry 41(2): 291–307.

De Dijn A (2018) Rousseau and Republicanism. *Political Theory* 46(1): 59–80.

Elazar Y and Rousselière G (2019) *Republicanism and the Future of Democracy*. Cambridge: Cambridge University Press.

Felsenthal DS and Machover M (1998) *The Measurement of Voting Power: Theory and Practice, Problems and Paradoxes*. Cheltenham: Edward Elgar.

Hayward C (2011) What Can Political Freedom Mean in a Multicultural Democracy?: On Deliberation, Difference, and Democratic Governance. *Political Theory* 39(4): 468–497.

Ingham S (2019) *Rule by Multiple Majorities: A New Theory of Popular Control*. Cambridge: Cambridge University Press.

Ingham S and Lovett F (2019) Republican Freedom, Collective Action, and Popular Control. *American Journal of Political Science* 63(4): 774–787.

Laborde C and Maynor J (2009) *Republicanism and Political Theory*. New York: John Wiley & Sons.

Kolodny N (2019) Being under the Power of Others. In *Republicanism and the Future of Democracy*, ed. Yiftah Elizar and Geneviève Rousselière. Cambridge: Cambridge University Press. Pp. 96–114.

List C and Pettit P (2011) *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.

Lovett F (2010) *A General Theory of Domination and Justice*. New York: Oxford University

Press.

Lovett F (2012) What Counts as Arbitrary Power? *Journal of Political Power* 5(1): 137–152.

Pettit P (1997) *Republicanism: A Theory of Freedom and Government*. New York: Oxford University Press.

Pettit P (2012) *On the People's Terms: A Republican Theory and Model of Democracy*. Cambridge: Cambridge University Press.

Schmidt AT (2018) Domination without inequality? Mutual domination, republicanism, and gun control. *Philosophy & Public Affairs* 46(2) 175–206.

Sharon A (2016) Domination and the Rule of Law. In *Oxford Studies in Political Philosophy*, vol. 2, ed. David Sobel, Peter Vallentyne, and Steven Wall. Oxford: Oxford University Press. Pp. 128–155.

Simpson TW (2017) The Impossibility of Republican Freedom. *Philosophy & Public Affairs* 45(1) 27–53.

Skinner Q (1998) *Liberty before Liberalism*. Cambridge: Cambridge University Press.