

Republican Freedom, Popular Control, and Collective Action

Sean Ingham* and Frank Lovett†

May 4, 2019

Abstract

Republicans hold that people are dominated merely in virtue of others' having unconstrained abilities to frustrate their choices. They argue further that public officials may dominate citizens unless subject to popular control. Critics identify a dilemma. To maintain the possibility of popular control, republicans must attribute to the people an ability to control public officials merely in virtue of the possibility that they might coordinate their actions. But if the possibility of coordination suffices for attributing abilities to groups, then, even in the best case, countless groups will be dominating because it will be possible for their members to coordinate their actions with the aim of frustrating others' choices. We argue the dilemma is apparent only. To make our argument, we present a novel interpretation of the republican concept of domination with the help of a game-theoretic model that clarifies the significance of collective action problems for republican theory.

*University of California San Diego. Email: singham@ucsd.edu

†Washington University in St. Louis. Email: flovett@wustl.edu.

Republicans hold that people are dominated, and therefore unfree, merely in virtue of others' having unconstrained abilities to frustrate their choices, even if the latter are benevolent and therefore unlikely to exercise those abilities.¹ Republicans further argue that public officials may dominate citizens unless subject to popular control.² Some recent critics identify a dilemma.³ To maintain the possibility of popular control, republicans must attribute to the people an ability to control public officials merely in virtue of the possibility that they might coordinate their actions. But if the mere possibility of coordination suffices for attributing abilities to collections of individuals, then, as one critic puts it, "republican freedom is impossible," for "there are always collections of agents who could, by coordinating, invade any individual," and who thus would count as sources of domination (Simpson, 2017, p. 41).

We argue that the dilemma raised by these critics is apparent only. We present a novel interpretation of the republican concept of domination, paying special attention to cases involving collections of agents, and show how it allows for the possibility of popular control, without making domination ubiquitous and therefore unavoidable. For collections of individuals to have abilities in the relevant sense, the mere *possibility* of coordination is insufficient: in a sense we will make precise, it must be the case that they would actually coordinate if each placed sufficient value on achieving their shared aim. Not all collections of individuals meet this condition; most are unable to surmount the barriers to

¹Earlier versions of this paper were presented at the Workshop in Politics, Ethics, and Society at Washington University in St. Louis; at the Philosophy Colloquium at Davidson College; and at the 2019 meeting of the Philosophy, Politics, and Economics Society. The authors would like to thank the participants at each of these events for their very helpful feedback. The authors are also especially grateful for comments and suggestions received from Matthew Babb, Randy Calvert, John Lawless, Daniel Layman, Brian Tamanaha, David Wiens, and the anonymous reviewers at *AJPS*.

²Significant contributions to the contemporary republican literature maintaining these claims include Pettit (1997, 2012, 2014); Skinner (1998, 2002); Viroli (2002); Lovett (2010). See also the papers collected in Laborde and Maynor (2009).

³See most importantly Simpson (2017), but similar arguments can be found in Dowding (2011), Sharon (2016), List and Valentini (2016), Kolodny (2019).

successful collective action. In a well-ordered republic, however, barriers to the collective action sustaining popular control over at least some public officials can be overcome. Republicans can thus explain how citizens exercise popular control, while avoiding the implication that generic collections of individuals have abilities to frustrate individuals' choices.

To make this argument, we present a game-theoretic model that clarifies the significance of collective action problems for the republican view. Though primarily a contribution to republican theory, our paper illustrates the value of bringing insights and tools from social science to bear on conceptual problems in political theory. Specifically, it shows how formal models can be useful, even when they do not generate testable predictions, as tools for conceptual analysis (Johnson, 2014; Ingham, 2015).

We do not here present a positive argument for the republican conception of freedom, except indirectly. By refining and clarifying republican ideas, our analysis aims to preempt important recent criticisms. If one is going to attack republicanism, one ought to attack the best version of that theory. The work in this paper is thus preliminary to either a successful defense or a persuasive critique of the republican conception of freedom.

1 Preliminaries

What are our criteria for success in characterizing republican freedom?

Our goal is to develop a *political* conception of freedom—that is, a conception designed to serve as a political ideal. This differs from analyzing the *meaning* of freedom as a concept. Generally speaking, we take freedom to be something a well-ordered society should honor or promote. Accordingly, we adopt a

- *Pragmatic constraint*: a political conception of freedom should be useful,

meaning roughly that it should have critical force without being hopelessly demanding.

The point of a political conception of freedom is to provide actionable guidance in the assessment of various social arrangements as better or worse. A conception which defined freedom in such a way that it would be impossible to secure in any reasonable degree no matter how hard we tried would presumably fail the pragmatic constraint: it would be hopelessly demanding.⁴

We assume that any attractive political conception of freedom would satisfy the pragmatic constraint. But what would qualify a conception of freedom as republican? Here we adopt a

- *Republican constraint*: to be a republican conception of freedom, the conception must capture several distinctively republican intuitions about freedom, including:
 - the subject of a benevolent master is unfree;
 - freedom is institutionally constituted; and
 - freedom requires popular control over at least some public officials.⁵

Each of these intuitions is well-supported in the classical republican tradition, and has also been embraced and elaborated by contemporary republican authors.⁶ Each also runs counter to the familiar conception of freedom as non-interference, as articulated by Berlin (1969) and others.

On the non-interference view, a person is free provided her choices are not actually frustrated, as they might not be by a benevolent master. Not so, argue

⁴For further discussion, see Pettit (2005).

⁵Popular control over all public officials—including public prosecutors, judges, central bankers, etc.—may be neither necessary for republican freedom nor desirable. Social conventions, legal constraints, or other forms of oversight will often be more appropriate. Republicans differ as to which public officials need to be subject to popular control under which circumstances. Our argument does not depend on which of these views one holds.

⁶Here the “classical republican tradition” refers to the conventionally recognized canon of early modern neo-roman republican authors from Machiavelli to Harrington to Madison. For contemporary republican authors, see footnote 1 above.

republicans. Algernon Sidney writes that the “the weight of chains, number of stripes, hardness of labour, and other effects of a master’s cruelty, may make one servitude more miserable than other,” but nevertheless “he is a slave who serves the best and gentlest man in the world, as well as he who serves the worst” (Sidney, 1996, p. 441). Glossing this sentiment, Pettit writes:

[W]hen a dominating agent is benign . . . , of course, the likelihood of interference will be that much lower. But it is important to see that domination goes with the accessibility of arbitrary interference to another, and improbability of the kind in question here does not make for inaccessibility. Someone can be in a position to interfere with me at their pleasure, even when it is very improbable that they will actually interfere. (Pettit, 1997, p. 64)

The non-interference view also holds that freedom is natural or pre-political: each law, institution, or convention represents a limitation on our freedom, even if that limitation is compensated for in other ways. Contrary to this view, William Blackstone famously insisted that “laws, when prudently framed, are by no means subversive but rather introductive of liberty” (Blackstone, 1979, p. 122). Pettit expands on this claim as follows:

[T]he law need not itself represent a form of domination There will be systems of law available, at least in principle, which are entirely undominating and entirely consistent with freedom: not only will they inhibit potential dominators and reduce unfreedom, they will do so without representing a form of domination in their own right. (Pettit, 1997, p. 66)

Finally, if, as on the non-interference view, we are free provided no one actually frustrates our choices, then it would seem that freedom can be enjoyed in any

form of polity, provided the rulers are suitably inclined to respect that freedom. Against this view, Machiavelli counts as free states only those “accustomed to living by their own laws,” not laws imposed by their rulers (Machiavelli, 1998, p. 20). As Skinner explains:

Free states, like free persons, are thus defined by their capacity for self-government. A free state is a community in which the actions of the body politic are determined by the will of the members as a whole. (Skinner, 1998, p. 26)

This is only a spare outline of some of the main features of republican theory, but it will serve for present purposes.

The particular worry motivating this paper concerns a possible inconsistency between the pragmatic constraint on the one hand, and the republican constraint on the other. The issue, pressed most forcefully by Thomas Simpson but noted also by others, can be summarized as follows.⁷

Suppose, to honor the republican constraint, we grant that freedom requires citizens to be able to control certain high-ranking public officials. To honor the pragmatic constraint, we must then believe such popular control is feasible. Unlike the state or a business corporation, however, the people are not a group agent in the proper sense: there are none of the institutionalized mechanisms for ensuring the sort of ongoing consistency in action and intention needed for group agency (List and Pettit, 2011; Pettit, 2012).⁸ If the people have the ability to control public officials, then they must have this ability in virtue of the fact that they might possibly coordinate their actions in the way needed to sanction and discipline wayward public officials.

But if the conditions for collections of individuals to have such abilities are

⁷See footnote 2 above.

⁸Here we accept List and Pettit’s account of group agency which is not at issue in the dispute that is the subject of this paper.

minimal enough that citizens have the ability to control public officials, then one may worry that we are surrounded by groups with the ability to frustrate our choices — indeed, nearly any arbitrary set of individuals we might imagine. It is possible, for example, that the strangers we encounter on entering a restaurant might spontaneously team up to prevent us from entering. Republicans cannot, so the argument goes, appeal to the improbability of this happening without abandoning their commitment to the view that probabilities do not matter in the analysis of freedom. Nor can they appeal to the fact that such collections of strangers fail to qualify as group agents, because if citizens are to have the ability to control public officials, then collections of individuals need not be group agents in order to be the bearers of abilities. Having excluded those grounds for appeal, it seems we must concede that domination is ubiquitous, and freedom so defined impossible to achieve.

To avoid this dilemma, we need an analysis of freedom as non-domination that captures all the distinctive republican intuitions, while respecting the pragmatic constraint. We make no attempt here to defend these constraints. Perhaps not everyone shares the intuition that the subject of a benevolent master is unfree and so on; neither, perhaps, does everyone share the intuition that the republican conception of freedom would be useless if domination turned out to be ubiquitous. For present purposes we simply grant both constraints and try to show that, within such parameters as they impose, an analysis of freedom consistent with republican commitments is available that does not succumb to the challenge raised by Simpson and others.

We use simple game-theoretic models to clarify the meaning and implications of our interpretation of republican freedom. Formal models are useful tools for conceptual analysis generally (Johnson, 2014; Ingham, 2015), but the nature of the dilemma motivating our analysis provides an additional justification. The

question at the heart of the dilemma is when a collection of individuals, who do not amount to a group agent, has an ability to frustrate another agent's choice. We argue that any compelling answer must recognize the importance of collective action problems, which critics like Simpson (2017) fail to do, as we explain at the end of section 3. Since formal models are especially useful aids for articulating and evaluating claims about collective action, they have a valuable role to play in this argument.

Before presenting a definition of domination by groups, and a model that illustrates how groups can fail to dominate in virtue of collective action problems, we first present a definition that covers the case of a single agent's dominating another, and a baseline model to illustrate the moving parts of our account.

2 A basic model of domination

Consider a simple model of an interaction between two agents, A and B . B chooses whether to perform the action ϕ or not. Before she makes her choice, A chooses whether to **intervene**, which refers to taking an action that affects either the possibility of B 's ϕ -ing or the payoff-relevant consequences of B 's ϕ -ing. An example of intervention would be a master's instructing his foreman to punish his slave if the slave performs the action ϕ . Having so intervened, ϕ -ing and not ϕ -ing now produce different payoffs for the slave than they would have, had the master not intervened.⁹ Another example would be a master's chaining his slave to a post, which affects the possibility of the slave's walking away.

We will wish to consider counterfactual variation in A 's preferences. To model such variation, we imagine "nature," as a metaphorical player, choosing A 's "type," that is, whether A prefers for B to ϕ or not, and how much impor-

⁹Some interventions may change the *expected* payoff to B 's ϕ -ing, but not the actual payoff. For example, if A makes a credible threat he does not intend to carry out, then his threat changes B 's expected payoffs from ϕ -ing, but not her actual payoff from ϕ -ing. Such actions are interventions, as we use the term.

tance A attaches to B 's ϕ -ing relative to the possible costs of intervention. (This terminology should not be understood to imply that A 's preferences are rooted in fixed “natural” facts about his psychological constitution, nor even that they are stable over time—the word *type*, conventional in game theory, is merely intended to refer to A 's different possible preferences, abstracting from their underlying causes.) Player A observes nature's choice before choosing whether to intervene.¹⁰

Figure 1 depicts this scenario. The pairs of numbers at the end of each branch of the game tree represent the players' preferences over the possible outcomes of their interaction: the first number represents A 's preferences and the second number B 's. Nature chooses v , which parameterizes A 's preferences. Specifically, if $v > 0$, then A prefers—all else being equal—for B not to ϕ , while if $v < 0$, he prefers for B to ϕ . He also prefers—again all else being equal—not to intervene, as intervention incurs a cost $c > 0$, but he would be willing to intervene, as a means of preventing B from ϕ -ing if (but only if) $v \geq c$. The payoffs for B indicate that in this scenario she would prefer to ϕ , all else being equal, but is better off not ϕ -ing if A has intervened.

We should emphasize that v is to be interpreted as encoding information about A 's psychology (e.g., his goals and the tradeoffs he is willing to make), not information about A 's external environment (e.g., the monetary costs or benefits of intervening). Thus, when we refer to counterfactual variation in nature's choice of v , we are referring to variation in facts about A 's psychology, not variation in A 's external circumstances. Let $\mathcal{T} = (-\infty, \infty)$ be the set of A 's possible types (the possible values of v).

A *strategy* for player A specifies the action that he takes—intervene, re-

¹⁰The type space is used to model counterfactual variation in A 's preferences and its effects, not B 's uncertainty about A 's type. We do not specify B 's beliefs about A 's type because they are unimportant for the analysis of the game's perfect Bayesian equilibrium.

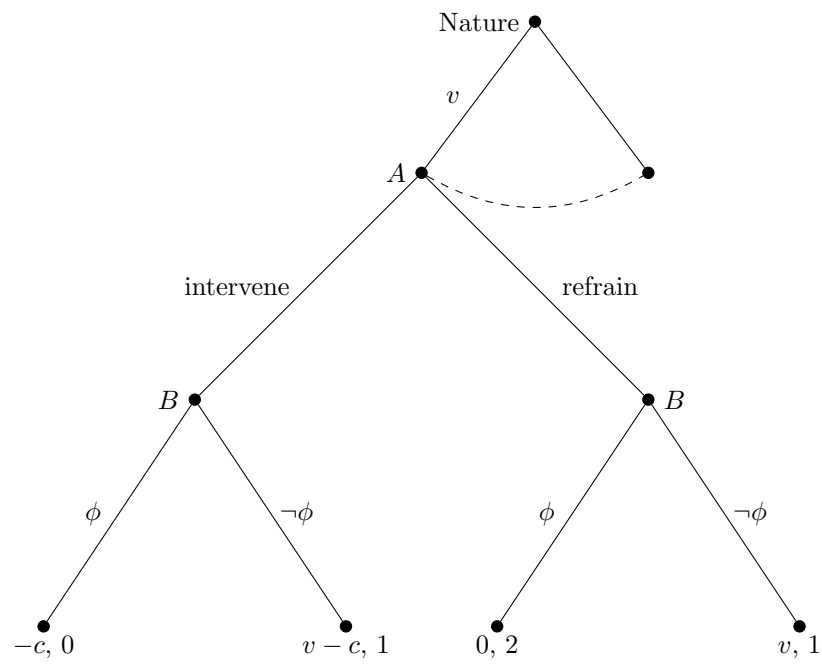


Figure 1: A situation in which A has an ability to frustrate B 's ϕ -ing.

frain — as a function of his type, v , while a strategy for player B specifies the action that she takes — ϕ , not- ϕ — as a function of player A 's action. A plausible assumption is that the players will choose sequentially rational strategies: player B 's strategy will be to choose ϕ unless player A has chosen to intervene, and player A 's strategy will be to intervene if $v > c$ and refrain from intervening if $v < c$. (Either action is consistent with sequential rationality if $v = c$.)

We now state definitions of domination and republican freedom, as they apply to a relationship between just two agents.

D: A dominates B if A has an insufficiently constrained ability to deliberately frustrate B 's choice whether to ϕ .

F: B 's choice whether to ϕ is not free if A has an insufficiently constrained ability to deliberately frustrate her ϕ -ing.

Note that **F** specifies a necessary condition for political freedom. All republicans are committed to this claim, but they disagree whether **F** is also a sufficient condition for freedom.¹¹

When does A have an ability to frustrate B 's choice whether to ϕ ? In the context of our simple model, we can say that

Player A has the **ability to frustrate** B 's choice whether to ϕ if

- i) even when preferring to ϕ , B will choose not to ϕ if A intervenes;
- and
- ii) if A prefers for B not to ϕ and cares enough, then A will choose to intervene.

In our example, condition (i) holds in the postulated equilibrium because B 's strategy was to ϕ if and only if A refrains from intervention. In concrete

¹¹Skinner (1998, p. 83) and Viroli (2002, p. 54), for example, argue that freedom additionally requires non-interference.

terms, this condition is met when the structural context gives A power over B or makes B dependent on A .¹² A has power over B in the sense that he has an action (intervention) by which he can make ϕ -ing less appealing to B , thereby changing what she would otherwise prefer to do. Condition (ii) holds in our example because if $v > c$ (meaning A prefers for B not to ϕ and also cares enough about B 's ϕ -ing, relative to the costs of intervention), then A will rationally choose to intervene. In cases where it is possible for a single agent to intervene in another's choice, and condition (i) is met, it would be unusual for condition (ii) to fail.¹³ The same will not be true, however, when we later generalize the condition to cover groups of agents.

When is such an ability sufficiently constrained? Let $\mathcal{P} \subset \mathcal{T}$ denote the subset of types of player A for whom it would be rational to intervene. In the example, $\mathcal{P} = \{v \in \mathbf{R} | v \geq c\}$. This is the set of the types who prefer for B not to ϕ and care enough about B 's ϕ -ing, relative to the associated costs of intervening, to make intervention worthwhile. The intuitive idea is that A 's ability to frustrate B 's choice whether to ϕ is sufficiently constrained if the set \mathcal{P} comprises — in some sense to be made more precise shortly — only types whose logical possibility is justifiably “ignored.”

As just one example to elicit the relevant intuition, suppose that in our example, *ϕ -ing* refers to B 's leaving her house, *intervening* refers to A 's holding B captive against her will, and the cost of intervention c is a lengthy prison sentence for A . Then \mathcal{P} comprises only those types of A who are so deranged that they prefer to hold B captive even though doing so will land them in prison.

¹²In general, the relevant structural context will be the upshot of the equilibrium behavior of further agents, C , D , and so on. The master's ability to frustrate the slave's choices depends not only on the foreman who executes the master's instructions, but also on the courts and officers of the state, among other collaborators, who enforce the laws upholding property rights in human beings. Contrary to a common misperception, the republican conception of freedom is not dyadic: we model a two-player game for expositional purposes only.

¹³One might imagine exotic cases where A suffers from a debilitating psychological fear of confrontation, such that even if she cares enough for it to be rational for her to intervene, she will nevertheless fail to do so.

Typically, it is reasonable to ignore the logical possibility that A could have such preferences. In this case, we would say that the law is a sufficient constraint on A 's ability to frustrate B 's choice whether to leave the house, since the law ensures that only ignorable types of player A would act to frustrate B 's choice.

The distinctive quality of the republican conception of freedom partly derives from its regarding certain possible worlds—here, certain possible types for player A —as significant even though they are merely counterfactual and, indeed, even though they are improbable. Whether B 's choice is free depends, not on A 's *actual* preferences— not on the actual realization of v — but instead on what A would do in various counterfactual scenarios in which his preferences are possibly different from what they actually are. Put differently, not all counterfactual types of player A are “ignorable.”

Consistent with that general thought, one can still generate weaker or stronger interpretations of the republican position from different interpretations of what makes certain types ignorable. Consider:

Strong republicanism. No logically possible type is ignorable.

This position is too strong. We introduce it here only to motivate a more reasonable interpretation of the republican view. One can see that it is too strong from reflecting on the example above. Suppose A and B are friends who live together. There is a logically possible type of player A who prefers for B not to leave the house and holds this preference so intensely that he is willing to suffer imprisonment in order to prevent B from leaving the house. According to the “strong republican” view, it follows that A has an insufficiently constrained ability to frustrate B 's choice whether to leave the house, and so B 's choice is unfree. On the strong view, freedom will be generally impossible, violating the pragmatic constraint.¹⁴

¹⁴This seems to be an implication of the analysis proposed in List and Valentini: rather

Here is a more plausible interpretation of the republican view:

Moderate republicanism. A subset of an agent’s possible types is ignorable if its becoming common knowledge that the agent’s type does *not* belong to that subset would have no significant practical consequences.

It is *common knowledge* that the agent’s type v does not belong to a subset $\mathcal{P} \subset \mathcal{T}$ if everyone knows that $v \notin \mathcal{P}$, everyone knows that everyone knows this, and so on, ad infinitum. Recall the example above, in which A and B are roommates, and the effective rule of law ensures that A would suffer severe punishment if he tried to hold B captive. Under normal circumstances, neither B nor anyone else would act differently if it somehow became common knowledge that A is *not* prepared to spend time in prison in order to keep B from leaving the house. Everyone already acts as if this is not a possibility, and as if everyone already knows that it is not a possibility, and as if everyone knows that everyone knows that it is not a possibility, and so on. Thus, if it became common knowledge that A is not this type, no one would act any differently. Its becoming common knowledge would have no significant practical consequences.

Contrast this situation with one in which the costs that A incurs from intervention are not severe legal repercussions, but rather the kind of trivial opportunity costs that accompany most action. For example, suppose that to intervene to prevent B from ϕ -ing, A , a slave owner, must rise from his chair, walk across the lawn in the blistering sun, and issue instructions to his foreman. All else being equal, he would rather not have to do this in order to keep B from ϕ -ing. It is costly, albeit trivially so. In anticipation of such situations, B , the owner’s slave, may have reason to ingratiate herself with A , or plead with A , with the

than characterize some choices as free and others as unfree, they prefer to focus on “the ‘size’ and ‘shape’ of the set of possible worlds across which the agent suffers no constraints” (List and Valentini, 2016, p. 1071). The analysis we propose is more tractable, in our view.

hope of shaping A 's preferences over B 's ϕ -ing, specifically with the hope of inducing in A an attitude of indifference or indulgence towards B 's ϕ -ing. She may have such reasons because she is unsure whether A has a preference over B 's ϕ -ing, and specifically whether A not only has a preference but cares enough to rise from his chair, walk across the lawn in the blistering sun, and issue instructions to his foreman. The possible type of A who is willing to incur those costs is not “ignorable,” because B would act differently — she would no longer have the same incentives to ingratiate herself with A , for example — if it became common knowledge that A is not that type. That such non-ignorable types of A would rationally choose to intervene to prevent B from ϕ -ing means that A has an insufficiently constrained ability to frustrate B 's choice whether to ϕ .

Because our analysis relies on a common knowledge condition, the determination of which types are ignorable will to some extent be specific to the cultural-historical context. Our condition is similar in spirit to Pettit's “eyeball test” (Pettit, 2012, p. 84): roughly speaking, one might say that people can look one another in the eye without reason for fear or deference — the eyeball test is met — when, according to local standards, all types worthy of fear or deference are ignorable. Importantly, however, whether the condition is satisfied does *not* depend merely on whether the agent subject to power is uncertain about the preferences of the agent with the ability to intervene.

Think of the position of a woman under traditional family law, who may be confident that her husband is not the type to exercise his legal powers of interference in her choices. Even if the wife is confident, the types of husband who would exercise these powers are not ignorable. Those types are non-ignorable because if it became, not just known to the wife, but also common knowledge that the husband is not one of those types, then the wife and others would act differently. For example, if the husband has the power to prevent his wife from

accepting positions of employment, then prospective employers, unsure about the husband's preferences and willingness to exercise this power, will have diminished incentives to try to recruit the woman. The woman, anticipating their behavior, will have diminished incentives to develop the kinds of productive skills that would make her valuable to prospective employers. If it were common knowledge that the husband is not the type to exercise his power to prevent her from taking employment outside the home, then she and the prospective employers would act differently. On the moderate republican view we have defined, such types of husbands are therefore non-ignorable *even though the wife may assign them zero probability*.

As this example should make clear, it would be false to say that what determines, on our account, whether A dominates B is the probability that A will wish to intervene. That judgment would rest on a misunderstanding of what makes one of A 's possible types ignorable. Nor is non-domination simply a matter of *everyone's* knowing that the husband is very unlikely to be the intervening type. Let everyone know this, and yet the wife may not know that prospective employers know this, and so she may choose not to invest in developing productive skills. Even supposing the woman knows that employers know this, there may be schools whose administrators do not know that she knows that employers know that her husband would not intervene, and thus do not expect her to be receptive to invitations to enroll in their schools, anticipating that she will view enrolling as fruitless. And so on. Since its becoming *common knowledge* that the husband is not the intervening type would be practically consequential, this type is not ignorable, and so his wife is dominated.

Our analysis generates standard republican judgments about canonical examples of domination. Consider:

- a) **The benevolent master.** The slave's choices are unfree even if the

slave's master does not wish to intervene and frustrate the slave's choices. What matters is not the master's actual type — not his *actual* preferences — but rather whether the master would rationally choose to intervene if the master preferred for the slave not to ϕ and cared enough about the slave's ϕ -ing to make the costs of intervention worthwhile (that determines whether the master has an ability to frustrate the slave's choice); and whether its becoming common knowledge that the master does not care enough to intervene would be practically consequential (that determines whether the “interventionist” types of the master are “ignorable” and, accordingly, whether the master's ability is sufficiently constrained). Given background assumptions about the master-slave relationship, both conditions will be met: non-ignorable types of the master will rationally choose to intervene in the slave's choices. That is true whether the master is actually benevolent or not.

- b) **The rule of law.** In the absence of law, people will generally be unfree in their choices. For example, B 's choice whether to travel unmolested will not be a free one, as various potential bandits will have insufficiently constrained abilities to frustrate her choice: non-ignorable types of the people she might encounter in her travels would rationally choose to intervene to prevent her from traveling unmolested, given the payments they would be able to extract from her in the absence of any legal repercussions. Thus, as the classical republicans claim, freedom is institutionally constituted, not natural: only with the effective rule of law do people become free to travel unmolested. Note, however, that even with the introduction of legal sanctions, there will still be some types — hardened criminals, say — who will rationally choose to intervene and frustrate B 's choices. But, where the rule of law accomplishes what republicans expect it to accomplish,

those types will be “ignorable” in ordinary circumstances.

- c) **Domination in the workplace.** Suppose A is an employer and B his employee, and A threatens to fire B unless she dresses in a particular way, which she would otherwise prefer not to do. This can be viewed as a reduction in effective compensation in so far as the employee is loathe to accommodate the request. In a perfectly competitive labor market, B would just quit and work for a different employer willing to pay her marginal revenue product. Anticipating this reaction, A would therefore not choose to intervene, condition (ii) would fail, and so A would not dominate B . But under realistic conditions, it will not be so easy for many employees to quit.¹⁵ Knowing this, non-ignorable types of employers may be disposed to intervene in the employee’s choices. Under such conditions, employers will dominate their employees in various workplace choices unless suitably constrained. Public provisions for basic needs can mitigate this domination, as can policies that strengthen workers’ collective bargaining rights: the former place an absolute limit on how far employers can reduce the value of employment, while the latter enhance the threat of strikes which might deter intervention by all but ignorable types of employers.

3 Domination by plural agents

So far, we have only considered the sense in which a single agent, A , may dominate another agent, B . We now extend the analysis and articulate the sense in which a plurality of agents might dominate B .

As before, it will be useful to have a concrete model to illustrate the concepts

¹⁵This may be because market concentration grants employers monopsony power, or because employees must bear the costs of searching for alternative employment. See Posner and Weyl (2018), ch. 4, for discussion of relevant literature.

we define below. Consider the following variation on our initial example: as before, an agent B chooses whether to perform an action ϕ or not. Before she makes her choice, two agents, A_1 and A_2 , simultaneously choose whether to contribute to a joint effort of intervening in B 's choice. Specifically, if, but only if, both of the agents choose to contribute, the intervention is successful and the consequences of B 's ϕ -ing are modified (just as when, in the original example, A intervenes). For example, the two agents may be the joint owners of a business and B their employee, and “contributing to a joint effort of intervening” may refer to agreeing to work together to draft a new workplace policy according to which employees are fired if they choose to ϕ . Alternatively, A_1 and A_2 may be two strangers who, acting together, could intervene to prevent B from traveling freely through town.

As before, we will wish to consider counterfactual variation in the preferences of A_1 and A_2 over B 's ϕ -ing. To model such variation, we again imagine “nature,” as a player, choosing their types—specifically the payoffs v_1 and v_2 that they receive when player B chooses not to ϕ —at the beginning of the game. Let $\mathcal{T} = (-\infty, \infty)$ be the set of A_i 's possible types (the possible values of v_i). To simplify the model, we will assume that players A_1 and A_2 each observe nature's choices, that is, they each know both their own type and the other player's type.¹⁶

Figure 2 depicts the situation arising after the determination of the players' types. Players A_1 and A_2 simultaneously choose whether to contribute. (Equivalently, first A_1 chooses whether to contribute and then A_2 , not having observed A_1 's choice and unsure whether he is at the decision node on the left or the node on the right, chooses whether to contribute.) B then chooses whether to ϕ . In the vectors of payoffs after each branch, the first number is A_1 's payoff, the sec-

¹⁶Nothing important hinges on this assumption. See the online supplementary materials for an explanation.




Figure 2: A situation in which A_1 and A_2 may, or may not, have an ability to frustrate B 's ϕ -ing, depending on which equilibrium they find themselves in.

and A_2 's, and the third B 's. The numbers represent the assumption that player B would prefer to ϕ , all else being equal, and A_1 and A_2 would each prefer not to contribute to the joint preventative effort, all else being equal (hence the $-c$ appearing in A_i 's payoff after any branch in which A_i contributes). If $v_i > 0$, then A_i prefers for B not to ϕ , and if $v_i \geq c$, then A_i would be willing to incur the costs of contributing to the intervention if A_i expected the intervention to be successful.

A strategy for player A_i , for $i = 1, 2$, specifies the action that he takes — contribute or not — as a function of nature's determination of the types, v_1 and v_2 , while a strategy for player B specifies the action that she takes — ϕ , not- ϕ — as a function of the actions taken by players A_1 and A_2 . The game has multiple perfect Bayesian equilibria; we focus on two that are plausible focal points.¹⁷

¹⁷See the online appendix for further discussion.

In each, player B chooses to ϕ unless players A_1 and A_2 have successfully intervened (that is, each chosen to contribute). In one, players A_1 and A_2 solve their coordination problem, each using the strategy of contributing if and only if $v_1, v_2 \geq c$. In the second equilibrium, they each use the strategy of never contributing. In that equilibrium, they suffer from a collective action problem: even when $v_1, v_2 \geq c$ and they would each benefit from a successful intervention, neither has any reason to contribute given his correct belief that the other will not contribute.

With this game as a point of reference, we now state general definitions of domination and republican freedom.

D: a group of agents A_1, \dots, A_m dominate B if they have an insufficiently constrained (collective) ability to deliberately frustrate B 's choice whether to ϕ .

F: B 's choice whether to ϕ is not free if a group of agents A_1, \dots, A_m has an insufficiently constrained (collective) ability to deliberately frustrate her ϕ -ing.

When $m = 1$, these definitions reduce to the original definitions, which describe the special case in which a group of one agent dominates another agent.

When does a group A_1, \dots, A_m have the collective ability to frustrate B 's choice whether to ϕ ? Let us say that:

A group of agents A_1, \dots, A_m has the **ability to frustrate** B 's choice whether to ϕ if

- i) even when preferring to ϕ , B will choose not to ϕ if each agent A_1, \dots, A_m chooses an action such that the profile of their actions constitutes a joint intervention; and
- ii) if each agent A_1, \dots, A_m prefers for B not to ϕ and cares

enough, then each will choose an action such that the profile of their actions constitutes a joint intervention.¹⁸

In the two equilibria described, condition (i) holds because B 's postulated strategy was to ϕ unless A_1 and A_2 each contribute and successfully intervene. Condition (ii) holds in one, but not the other equilibrium. In the equilibrium where players A_1 and A_2 overcome their potential collective action problem, they each use a strategy of contributing if and only if $v_1, v_2 \geq c$. In other words, if each of them prefers for B not to ϕ (meaning $v_1, v_2 > 0$) and each cares enough about B 's ϕ -ing relative to the costs of intervention (specifically, if $v_1, v_2 \geq c$), then each chooses to contribute, and the profile of their actions is such that they successfully intervene. But in the other equilibrium, where the collective action problem is not overcome, condition (ii) fails to hold. For in that equilibrium, no matter how much A_1 and A_2 care about B 's ϕ -ing relative to the costs of intervention — no matter how great v_1 and v_2 — they each choose not to contribute and thus fail to intervene.

As this example illustrates, one reason that a group of agents A_1, \dots, A_m may not dominate an agent B is that collective action problems deprive them of a genuine ability to frustrate B 's choice. It may be *possible* for them to so act that B 's choice is frustrated — condition (i) may hold; but if they will fail to so act, irrespective of how strongly they prefer for B not to ϕ — if condition (ii) fails to hold — then they lack a genuine *ability* to frustrate B 's choice and do not dominate B .¹⁹

Collective action problems come in many varieties, some being more difficult to overcome than others. A group of agents may fail simply because they lack a mechanism for communicating and coordinating their actions. Our example

¹⁸Here we refer to a “profile of actions” because in many cases not every agent in a group need perform the same action for a successful intervention to occur. See footnote 18.

¹⁹Thus our analysis differs from that of List and Valentini (2016), which counts all such possibilities as genuine abilities.

shows that a group can fail in this way even in conditions that are conducive to collective action. The model makes the unrealistic simplifying assumption that it is common knowledge when all members of the group care more about preventing B from ϕ -ing than avoiding the cost of contributing to an intervention (the realized values of v_i are common knowledge). It is also common knowledge in our model that each must contribute to the intervention if they are to keep B from ϕ -ing: there is no possibility of “free-riding.” Even so, it does not follow that the members of the group will rationally act so as to intervene whenever they prefer for B not to ϕ : despite unrealistic assumptions that stack the deck in favor of collective action, they might still form self-fulfilling expectations that no one will contribute to the joint intervention.

In other more realistic cases, the deck is stacked *against* collective action. Normally, it will be unnecessary for every member of a group to contribute to an intervention for it to be successful, and so each individual may hope to “free ride,” reaping the benefits of an intervention without contributing to the effort.²⁰ Since the members of a generic collection of individuals, who are all strangers, will normally be unsure about each other’s preferences, nearly insurmountable obstacles stand in the way of coordinating expectations on exactly who will make the necessary contributions. In the real world, the most realistic equilibrium is thus the one where no one contributes to the intervention, however strongly each might prefer for B not to ϕ . When the members of a multitude find themselves in that kind of equilibrium, they lack a genuine *ability* to frustrate B ’s choice even though it is *possible* for them each to act so as to collectively frustrate B ’s choice.

²⁰Our game incorporates a special case of a Palfrey-Rosenthal contribution game in which a group with n members succeeds if and only if at least k members contribute, where $2 \leq k \leq n$ (Palfrey and Rosenthal, 1984). Even if the threshold k is less than n , there are equilibria in which the group overcomes its collective action problem, with exactly k members of the group contributing, but those equilibria are not very plausible for the reasons explained in the main text.

None of this is to deny that *some* groups have mechanisms for coordinating their actions, and will thus have abilities to frustrate individuals' choices. The point is just that for *generic* collections of agents, the mere fact that each agent cares a great deal about stopping some individual from ϕ -ing is no reason by itself to expect that they will undertake the actions necessary to stop that individual from ϕ -ing. In contrast to single-agent cases, one cannot reliably infer that a group has an ability to frustrate an individual's choice from the mere fact that it is possible for the group to do so.

Among those special groups that have solved their collective action problems, and thus have abilities to frustrate individuals' choices, such abilities, like those of a single agent, may still be sufficiently constrained so as to not count as dominating. In line with our earlier analysis, we will say that the ability of A_1, \dots, A_m to frustrate B 's choice is sufficiently constrained if the combinations of types who would rationally choose to act so as to frustrate B 's choice are ignorable. Specifically, let \mathcal{T}_i denote A_i 's possible types, so that $\mathcal{T}_1 \times \dots \times \mathcal{T}_m$ denotes the set of all possible combinations of types for the agents A_1, \dots, A_m . Let $\mathcal{P} \subset \mathcal{T}_1 \times \dots \times \mathcal{T}_m$ denote the set of type combinations who would rationally choose contributions that, in aggregate, make for a successful intervention in B 's choice. Then the ability of A_1, \dots, A_m to frustrate B 's choice is sufficiently constrained if \mathcal{P} is ignorable. Like before, a subset of $\mathcal{T}_1 \times \dots \times \mathcal{T}_m$ is ignorable if its becoming common knowledge that the types of A_1, \dots, A_m do not lie in this set would have no practical consequences.²¹

²¹Note that whether a group's ability to frustrate is constrained turns on the ignorability of *profiles* of types, not on the separate ignorability of each individual member's type taken in isolation. For example, suppose that a group of agents A_1, \dots, A_m vote on whether to retain an employee, B , who will be fired unless they *all* vote to retain her. For any particular agent A_i , its becoming common knowledge that A_i is not the type who would vote against B may have no impact on anyone's behavior—the set of A_i 's types who would vote against B may be ignorable. At the same time, its becoming common knowledge that *none* of the agents A_1, \dots, A_m is the type to vote against B might very well affect the behavior of B or others—the set of all profiles of types in which at least one agent would vote against B might be non-ignorable.

The following examples illustrate the above ideas.

- a) **Perfectly competitive markets.** If a monopolist A formed a preference over a consumer B 's purchasing decisions and cared enough (think of paternalistic firms running "company towns"), it would intervene in B 's choices, refusing to sell or charging B a price she cannot afford. It would dominate her. But in a perfectly competitive market, no firm A or group of firms A_1, \dots, A_m will dominate any consumer B . Given her budget, B will face a menu of affordable bundles of goods and services, and no firm or group of firms will have an effective ability to frustrate her choice. It is true that, if each firm reduced its output, the market price would rise, preventing B from choosing the bundle of goods and services she otherwise would. But even if each firm for some reason formed a preference over B 's choice, no firm would unilaterally reduce its output, because unilateral action would have a negligible effect on the market price, and no firm would expect any other to follow suit. The logic of collective action explains why they lack an ability to frustrate B 's choice in a competitive market. Of course, markets are unlikely to remain competitive without fair practice laws that ban collusion and price-fixing agreements. But the significance of such laws is not that they make it *impossible* for all the firms in the market to jointly intervene in the individual's choice by restricting output. Rather, their role is to prevent the firms from pursuing solutions to their collective action problem.
- b) **Consumer boycott.** Suppose that each member of a group of consumers A_1, \dots, A_m prefers for a firm B not to engage in some practice, and they all care enough about B 's practice to contribute to a boycott provided the others do as well. Their collective action problem might be overcome by a consumer advocacy group (since this sort of coordination, unlike

price-fixing, is not legally barred). If so, and if their intervention would induce B to end the targeted practice, then on our analysis they have a genuine ability to frustrate B 's choices. Their ability may not be sufficiently constrained—the types of consumers who would so intervene may be non-ignorable. In that case, the consumers dominate the firm's choice. Ordinarily, this does not concern us, since we don't care about the freedom of corporations. However, there might be problematic cases, as for example boycotts of stores that voluntarily desegregated in the civil rights era.

- c) **Spontaneous mob.** It is possible for a set of strangers in a restaurant to block a prospective customer from entering the restaurant if enough of them act in concert. But in a healthy society they will normally lack the ability to do so owing to collective action problems. Imagine each restaurant patron, upon seeing B approach, prefers that she not enter. Even if each would be willing to contribute to a joint effort to block B 's entrance—say, by congregating near the entryway—none will contribute, anticipating that unilateral action will be ineffective. Condition (ii), necessary for the group to have a real ability to frustrate B 's choice to ϕ , will fail to hold. Social norms may further exacerbate the collective action problem by creating an expectation that any attempt to unilaterally test the waters will incur social disapproval.

Say that a group A_1, \dots, A_m is a *team* if it has the effective ability to frustrate B 's choices, that is, if both conditions (i) and (ii) hold, while a group is a *mere multitude* if, owing to some collective action problem, condition (ii) fails to hold. Mere multitudes cannot dominate anyone. Even though it is often possible for the members of a mere multitude to so act that an individual would be frustrated in her choices—even though condition (i) often holds for various collections of

agents—collective action is not a trivial accomplishment. Condition (ii) will fail for generic collections of agents.

We can now directly respond to the challenge motivating this paper. Simpson (2017) entertains a condition for a group’s having an ability that is, on its face, similar to ours.²² Yet he claims that his condition will be satisfied for generic collections of individuals, while we claim that ours will be satisfied only in special cases, where groups have mechanisms to solve their collective action problems and the resulting abilities to frustrate the individual’s choices are insufficiently constrained. To see the reasons we come to a different conclusion, consider Simpson’s example of “a Dalit (‘Untouchable’) in a rural area of India breaking a caste norm” in the presence of other villagers. The villagers, by joining together, could prevent the infraction, but they have each come to repudiate caste norms. Do they dominate him in this case? Simpson claims that republicans must say that they do if all that keeps them from interfering in the Dalit’s choices is their moral repudiation of caste norms, a merely “internal” psychological constraint. In that case, they have an unconstrained power to interfere, by coordinating their actions. But if the Dalit is dominated in such a case—which Simpson believes is not atypical in its essential features—then “republican freedom is impossible” (Simpson, 2017, p. 33–34).

On our account, the crucial questions are, first, whether the other villagers would intervene if they preferred for the Dalit not to violate the caste norm and cared enough relative to the costs of intervening; and, second, whether the type profiles satisfying this condition are ignorable. Simpson’s construction of the example does not answer these questions. To see why, consider two variations on the example.

²² “Imagine that each member of G were to will to invade A, by coordinating. To the extent possible, hold fixed everything else, including the wills of all other agents. Then ask whether they invade A. If so, then the members of G have the power to invade A, otherwise not” (Simpson, 2017, p. 40).

In the first, not only does each villager repudiate caste norms (as Simpson supposes) but everyone believes that everyone repudiates caste norms. Now if each villager were privately to have a change of heart and come to prefer that the Dalit not violate the norm, none would expect anyone else to join him in enforcing the caste norm. For, by hypothesis, everyone believes that everyone repudiates caste norms. If it is costly to contribute to such a joint intervention, none of the villagers would rationally choose to do so, anticipating that their effort would be in vain.²³ Thus, the first necessary condition for domination fails to hold; the group lacks a genuine ability to intervene in the Dalit's choice, owing to the collective action problem.

Contrast this situation with one in which, as a matter of common knowledge, people endorse caste norms and are generally willing to join with others to enforce them. That common knowledge would plausibly sustain collective action: if each villager preferred for the Dalit not to violate the norm, then, anticipating that others share their preference and will join in the effort, they would each contribute to the joint intervention to enforce the norm. And if it is commonly believed that the typical villager is willing to do this, then it would of course be consequential if it somehow became common knowledge that the actual villagers happen not to be willing to join in such interventions; such combinations of types would be non-ignorable.²⁴ Under those specific assumptions, republicans quite rightly hold that the Dalit is dominated, even if, as it happens, he is in the presence only of atypical villagers who have repudiated

²³Might the joint intervention arise under these assumptions through an unintended cascade, where first one villager moves against the Dalit even though he has little reason to expect the intervention will succeed; then, observing the first mover, a second villager, now sufficiently confident that the intervention will succeed, joins in; then a third, and so on? If that is what would happen given that each villager preferred for B not to violate the caste norm, then the group would have an ability to intervene. But this profile of types in many cases might be ignorable. For example, the type of player A_1 who is willing to contribute to such an intervention, even though (as in our version of the example) he doubts anyone will follow suit, is plausibly ignorable.

²⁴More precisely, the profiles of types that induce such interventions are non-ignorable.

caste norms and are not the types to contribute to joint interventions to enforce them.

4 Popular control

For the most part, a well-ordered republic secures our freedom from domination in everyday interactions with fellow citizens through the effective rule of law and basic public welfare provisions, and it will be the state that provides these services under modern conditions. But if the state is powerful enough to secure citizens' non-domination in their private relationships with each other, there is a danger that those public officials who direct the state will dominate citizens. What protects our freedom from this threat?²⁵

Referring back to our basic model, suppose that A is a public official and B an ordinary citizen. Let ϕ -ing refer to the exercise of some legally protected basic liberty, say, publishing an article critical of the government. Before B chooses whether to ϕ or not, A chooses whether to intervene. Here an intervention might be instructing the police to arrest people who publish seditious articles. The public official will assuredly have the ability to frustrate B 's choice whether to ϕ : if A intervenes, B will choose not to publish a seditious article, satisfying condition (i); and A will choose to intervene if he prefers for B not to publish seditious articles and cares enough, satisfying condition (ii). It follows on the republican view that B will not be free to publish articles critical of the government unless the public official's ability to so intervene is suitably constrained. It will be suitably constrained, on our analysis, if and only if the types of public official who would rationally intervene are all ignorable.

So far in our paper we have mostly considered examples in which agents are

²⁵We focus here on the danger that public officials will dominate citizens, becoming their "rulers" instead of public servants. Alternatively, one might worry about the possibility that the state itself, as a corporate agent, will dominate citizens.

constrained by the law. But in the case of especially high-level public officials, the ultimate source of constraint might have to be different. The basic liberties may be written down in law, constitutional or otherwise, but what guarantees that powerful top public officials will respect those laws? The republican answer is the citizens themselves, empowered with the means to control public officials and virtuously disposed to do their part in exercising those powers when necessary.²⁶ Civic virtue and popular control are thus both central to the republican project.

The people are not a proper group agent, however. If they have an ability to control their public officials, it must be in the nature of a collective ability they hold as a plurality of agents. In section 3 we defended an analysis of what it means for a plurality of agents to have the collective ability to frustrate another agent's choices, and showed that on this analysis domination is not ubiquitous. We must now explain why it leaves open the possibility that citizens may have the ability to control public officials, preventing them from turning into rulers who dominate them.

Suppose that citizens have control over their public officials in this sense: if sufficiently many citizens preferred for their public officials to respect and enforce the law, and if those citizens cared enough about the issue, then they would choose actions—voting out corrupt elected officials, protesting their abuses, petitioning, and so on—that would in turn induce the relevant public officials to respect and enforce the law.²⁷ And suppose, further, that citizens are “virtuous” in the following sense: they not only prefer for their public officials to respect

²⁶We need not take any stand on exactly which public officials need to be under popular control (see footnote 4 above). The question we are addressing is whether, assuming popular control is *sometimes* necessary for republican freedom, we can avoid the conclusion that domination is ubiquitous while explaining the feasibility of whatever sort of popular control is required.

²⁷See Ingham (2019) for an explanation of how multiple majorities could have control in this sense, at the same time, even if citizens' preferences do not aggregate into a coherent “popular will.”

the law, but they also care enough about the rule of law that they are prepared to exercise their control over public officials, and choose the actions that induce their public officials to respect the laws when necessary. Finally, suppose that citizens' control over public officials is "robust" in the following sense: these actions induce the desired behavior in all but ignorable types of public officials, e.g., types of public officials who do not care about losing office or suffering even worse repercussions, like imprisonment, for corrupt behavior. Then the public official's ability to frustrate the citizen's choice whether to publish the article will be suitably constrained: only ignorable types of the public official, who do not care about suffering electoral defeat or other repercussions, will deviate from the law and intervene in the citizen's choice.

Importantly, the citizens' collective ability to control public officials in this way is the fruit of the right institutional environment — one that institutionalizes mechanisms of popular control — and the right political culture — one that gives citizens confidence that enough of their fellow citizens will join in collective action to resist officials' abuses of power, and instils in them a sense of civic duty to do their bit when enough others do theirs. Together, the right institutional environment and corresponding political culture make republican freedom sustainable "in equilibrium." This is not a trivial accomplishment. Only in a well-ordered republic do the usual barriers to collective action become surmountable, so that a critical mass of citizens becomes a latent team ready to exercise popular control. It may have been *possible* for the people living under traditional European monarchies to control their public officials through an implicit threat of revolution, but absent coordination mechanisms they normally lacked the *ability* to do so, and successful revolutions were accordingly rare. The people in such societies were mere multitudes.

5 Conclusion

Republicans hold that people are dominated if they are exposed to others' unconstrained abilities to frustrate their choices. To explain the sense in which a *group* could dominate an individual, we need an account of what it means for a group to have an ability. Critics of the republican position have exploited the ambiguity in this notion, arguing that if citizens, as a group, have the ability to control public officials, then republicans must also concede that individuals face innumerable groups with unconstrained abilities to frustrate their choices.

In reply, we have given an account of the conditions under which groups have abilities, and we have used it to formulate a novel interpretation of freedom as non-domination. A group lacks an ability to frustrate an individual's choice if—owing to collective action problems—its members would fail to act in the ways necessary to frustrate her choice, no matter how strong their shared preferences over her choice. Since generic collections of agents face collective action problems, they do not dominate anyone; they are mere multitudes, not latent teams. But—turning to the other horn of the dilemma—citizens may still be said to control public officials in a well-designed republic, where political institutions and civic culture help them surmount collective action problems.

We conclude with a question concerning popular control and non-domination. Popular control is one mechanism by which public officials, contemplating interventions that deviate from existing law, can be suitably constrained. But it is unrealistic to suppose that even the most comprehensive legal system could fully determine the exercise of public authority, were this even desirable. Considerable areas of discretion will necessarily remain. Some of this will take the form of agency discretion, as when central banks have the authority to set interest rate policies. The most important form of discretionary authority, however,

will be legislative: determining what the laws are going to be in the first place.

Significantly, popular control as such does not eliminate discretion so much as relocate it. To the extent that the people determine how public officials exercise their discretionary authority, it would seem that the said discretion really lies with popular majorities. Conversely, to the extent that the ability of popular majorities to direct discretionary authority is incomplete, any residual discretion necessarily falls to the relevant public officials. Someone will have to decide whether gambling should be outlawed or not, for example, and if it's not the people themselves who decide, then it will be public officials (legislators), and vice versa. Thus it is incumbent on republicans to provide an account of the conditions — if any — under which the exercise of public discretionary authority does not compromise our freedom.

One possibility is simply to concede that discretionary public authority entails some degree of domination. Since discretionary authority cannot be fully eliminated from public life, the relevant questions will be who should exercise how much of it in which domains. The answers will involve a complex balancing of the value of freedom from domination against other goods served by discretion. For example, should we prioritize freedom from domination by making legislation difficult and rare, even at the cost of inflexibility and inefficiency, or the reverse?

Another familiar response is to argue that so long as certain conditions are met, the exercise of discretionary public authority by the people themselves (through their elected representatives) over themselves does not constitute domination in the relevant sense. Pettit appeals, by analogy, to the intuition that if Andrea gives the key to her alcohol cupboard to Bob, with instructions not to give it back except on twenty-four hours notice, then Bob does not dominate her when he refuses to deviate from her instructions (Pettit, 2012, pp. 57–58).

Given the tangential relationship of these issues to the main topic of this paper, we do not here take a position on the merits of these or other possible responses. We merely flag them as important topics for future work.

References

- Berlin, Isaiah. 1969. *Four Essays on Liberty*. Oxford University Press.
- Blackstone, William. 1979. *Commentaries on the Laws of England*. Vol. 4. University of Chicago Press.
- Dowding, Keith. 2011. “Republican Freedom, Rights, and the Coalition Problem.” *Politics, Philosophy & Economics* 10(3): 301–322.
- 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. 2019. *Rule by Multiple Majorities: A New Theory of Popular Control*. Cambridge University Press.
- Johnson, James. 2014. “Models among the Political Theorists.” *American Journal of Political Science* 58(3):547–560.
- Kolodny, Nikolo. 2019. Being under the Power of Others. In *Republicanism and Democracy*, ed. Yiftah Elizar and Geneviève Rousselière. Cambridge University Press.
- Laborde, Cécile and John Maynor. 2009. *Republicanism and Political Theory*. John Wiley & Sons.
- List, Christian and Laura Valentini. 2016. “Freedom as Independence.” *Ethics* 126(4): 1043–1074.
- List, Christian and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- Lovett, Frank. 2010. *A General Theory of Domination and Justice*. Oxford University Press.
- Machiavelli, Niccoló. 1998. *The Prince*. University of Chicago Press.

- Palfrey, Thomas R. and Howard Rosenthal. 1984. "Participation and the Provision of Discrete Public Goods: A Strategic Analysis." *Journal of Public Economics* 24(2): 171–193.
- Pettit, Philip. 1997. *Republicanism: A Theory of Freedom and Government*. Oxford University Press.
- Pettit, Philip. 2005. The Domination Complaint. In *Political Exclusion and Domination*, ed. Melissa S. Williams and Stephen Macedo. New York University Press.
- Pettit, Philip. 2012. *On the People's Terms: A Republican Theory and Model of Democracy*. Cambridge University Press.
- Pettit, Philip. 2014. *Just Freedom: A Moral Compass for a Complex World*. W.W. Norton & Company.
- Posner, Eric A. and E. Glen Weyl. 2018. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Princeton University Press.
- Sharon, Assaf. 2016. "Domination and the Rule of Law." *Oxford Studies in Political Philosophy* 2: 128–55.
- Sidney, Algernon. 1996. *Discourses concerning Government*. Liberty Fund.
- Simpson, Thomas W. 2017. "The Impossibility of Republican Freedom." *Philosophy & Public Affairs* 45(1):27–53.
- Skinner, Quentin. 1998. *Liberty before Liberalism*. Cambridge University Press.
- Skinner, Quentin. 2002. A Third Concept of Liberty. In *Proceedings of the British Academy* 117: 237–268.
- Viroli, Maurizio. 2002. *Republicanism*. Hill & Wang.

Supplementary Material for Online Appendix

Original model

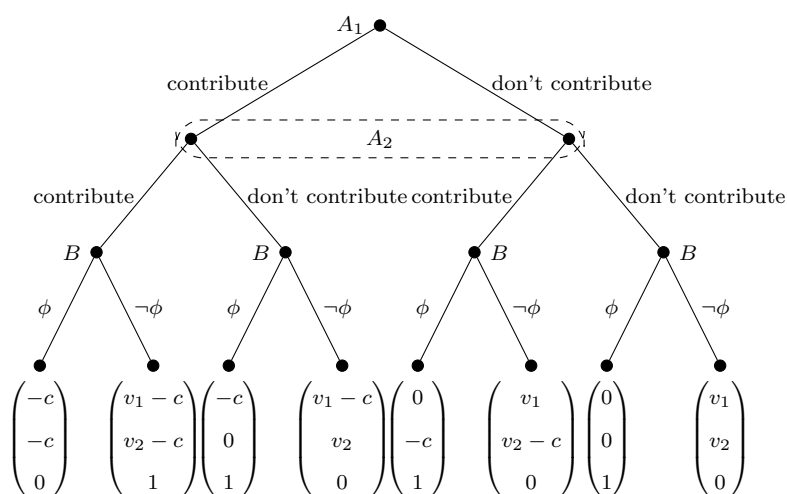
We present here, for ease of reference, the original version of the model of a latent team, from section 3 of the manuscript, before presenting an extension with incomplete information.

An agent B chooses whether to perform an action ϕ or not. Before she makes her choice, two agents, A_1 and A_2 , simultaneously choose whether to contribute to a joint effort of intervening in B 's choice. Specifically, if, but only if, both of the agents choose to contribute, the intervention is successful and the consequences of B 's ϕ -ing are modified. For example, the two agents may be the joint owners of a business and B their employee, and “contributing to a joint effort of intervening” may refer to agreeing to work together to draft a new workplace policy according to which employees are fired if they choose to ϕ . Alternatively, A_1 and A_2 may be two strangers who, acting together, could intervene to prevent B from traveling freely through town.²⁸

To model counterfactual variation in the preferences of A_1 and A_2 , we imagine “nature,” at the beginning of the game, choosing their types: the payoffs v_1 and v_2 that they receive when player B chooses not to ϕ . Here, we assume that players A_1 and A_2 each observe nature's choices, that is, they each know both their own type and the other player's type. In the next section, we assume instead that they each know their own type but are uncertain of the other player's type. Throughout, we assume that B does not observe their types, but has beliefs represented by some probability distribution over nature's choices. There is no need to specify the distribution; whatever B 's beliefs, the only sequentially rational strategy is ϕ unless A_1 and A_2 have each contributed.

²⁸In these examples, the relevant agents are individuals, but in principle they might be group agents, as when a plurality of business corporations team up for some joint purpose.

The figure below depicts the situation arising after the determination of the players' types. Players A_1 and A_2 simultaneously choose whether to contribute. Having made their choices, player B then chooses whether to ϕ . In the vectors of payoffs after each branch, the first number is A_1 's payoff, the second is A_2 's, and the third is B 's. A strategy for player A_1 or A_2 is a function that assigns to each possible pair (v_1, v_2) of realized types the player's action, contribute or not.



A strategy for player A_i is a function that specifies the action he takes — contribute, don't contribute — as a function of the pair (v_1, v_2) that nature chooses. A strategy for player B is a function that specifies the action she takes — ϕ , not- ϕ — as a function of the actions taken by A_1 and A_2 .

The game has infinitely many perfect Bayesian equilibria. In each equilibrium, player B 's strategy is to ϕ unless A_1 and A_2 have each chosen to contribute; as noted above, this is the only sequentially rational strategy for B , whatever B 's beliefs about the distribution from which nature chooses v_1 and v_2 . In every equilibrium, A_1 and A_2 use the same strategy. We focus on two equilibria that are plausible focal points.²⁹ In one, A_1 and A_2 never choose

²⁹As an example of one of the infinitely many equilibria we ignore, suppose A_1 and A_2 each

to contribute. In the other, each uses a strategy of contributing if and only if $v_1, v_2 \geq c$.

Revised model with incomplete information

We now consider a revised model that shows that our conclusion in the main text, made on the basis of the model from the previous section, is not just a consequence of our simplifying assumption that A_1 and A_2 observe the other's type. The conclusion, recall, was that the two players might overcome their collective action problem and have an ability to frustrate B 's choice, but they might also not; one cannot conclude that they will act to frustrate B 's choice merely from the assumption that they are rational and care a lot about frustrating B 's choice. In the original model, that conclusion followed because there was an equilibrium in which they each contribute when they place sufficient value on frustrating B 's choice, but also an equilibrium in which neither contributes, no matter how much value they attach to frustrating B 's choice. We will show that the same conclusion holds of the revised model where A_1 and A_2 are uncertain about the other's type.

Assume that players A_1 and A_2 each observe their own type but not the other's, and, as before, B observes neither player's type. A strategy for player A_1 is now a function mapping a realized value of v_1 to an action—contribute or not—and a strategy for player A_2 is a function mapping a realized value of v_2 to an action—contribute or not. Player B 's strategy set remains the same as before. Assume v_1 and v_2 are distributed identically and independently according to a distribution with cumulative distribution function F .

We characterize two perfect Bayesian equilibria in pure strategies. Sequential rationality requires, as before, that player B choose to ϕ unless both A_1 and A_2 have contributed to a successful intervention. Consider now strategies for

use a strategy of contributing if and only if $v_1, v_2 \geq 2c$.

players A_1 and A_2 with the following “cutoff” form: contribute if and only if $v_i \geq \bar{v}$. Assume the players use these strategies, and consider the payoff to player A_i of following it (for each $i = 1, 2$). If A_i does not contribute, his payoff is 0, while if he contributes, his expected payoff (given that A_j ($j \neq i$) and player B are using the postulated strategies) is $(1 - F(\bar{v}))v_i - c$. Thus, following the strategy is optimal if

$$v_i \geq \frac{c}{1 - F(\bar{v})}$$

Thus, the postulated strategy will be part of a perfect Bayesian equilibrium if there is a number \bar{v} such that the inequality holds if and only if $v_i \geq \bar{v}$. Under certain assumptions about c and F , such a number will exist. For example, if F is the uniform distribution and $c = 1/4$, then such an equilibrium exists with $\bar{v} = 1/2$ as the threshold.

There is also an equilibrium in which A_1 and A_2 each choose not to contribute irrespective of their type (and B 's strategy is the same as before).

In the first equilibrium, A_1 and A_2 will have an ability to frustrate B 's choice, as they will choose to contribute to a successful intervention if and only if they each prefer for B not to ϕ and care enough (if $v_1, v_2 \geq \bar{v}$). In the latter, they will lack this ability.

Thus, the model that incorporates uncertainty about others' types produces the same conclusion: in one equilibrium, A_1 and A_2 are a latent team, who will cooperate to achieve a shared aim of preventing B from ϕ -ing if they each prefer for her not to ϕ and care enough; while in another equilibrium, they are a mere multitude, hamstrung by the collective action problem.