On dynamic deterrence in collective protests*

Philippe Jehiel† Mathieu V. Leduc‡

January 20, 2026

Abstract

We study a dynamic model of social protest formation where activists asynchronously decide to join a movement, while a government can deploy repressive action at a chosen time. As the protest expands, punishment eventually ceases to be dissuasive due to a "safety in numbers" phenomenon. We show that the government has incentives to delay repression as long as it remains dissuasive, thereby providing a bound on protest size. When agents imperfectly observe protest participation, the government may further delay repression to obscure the protest size, inadvertently allowing participation to expand beyond the point where punishment remains dissuasive and deterrence collapses.

Keywords: Social Protests, Dynamic Coordination, Deterrence, Imperfect Monitoring
JEL codes: D00; C72; C73; D70; D82

*We thank seminar participants at the Lisbon Meetings in Game Theory and Applications #13 (Lisbon) and the European Summer Symposium in Economic Theory - ESSET 2024 (Gerzensee). We thank, namely, Bruno Strulovici, Thomas Wiseman, Nicola Persico, Sanjeev Goyal, Jan Zapal, Bryony Reich, Georg Nöldeke, Jesus Sanchez Ibrahim, Arda Gitmez, Andrea Prat, Arjada Bardhi, Yingni Guo, Marina Agranov, Meg Meyer, Leeat Yariv, Tom Palfrey and David Myatt.

†Email: philippe.jehiel@gmail.com

‡(Corresponding author) Email: mattvleduc@gmail.com.

1 Introduction

In many situations, a government must choose not only if, but also when to take repressive action against a social protest movement. This repressive action may be politically controversial or resource-draining and thus the government may not be able to take it more than once. Moreover, it is uncertain whether this repressive action will actually succeed in dismantling a protest movement. For instance, in the early phase of the Arab Spring in Egypt in 2011, as demonstrations in Cairo's Tahrir Square grew in numbers many observers expected the Mubarak government to employ decisive repressive measures to curb the movement. However, the regime hesitated to act, in part due to uncertainty about the potential outcomes and the domestic and international costs of using force on such a scale. Similarly, during the 1989 Tiananmen square protests in China, the leadership initially refrained from repression despite mounting demonstrations.

Since repressive action can be viewed as terminal—meaning that once undertaken, the government loses the strategic ability to act again—it may be tempted to delay it. Anticipating this, agents may consider the threat of repression to be non-credible and decide to join the movement. Yet, at a certain threshold, when the government becomes cornered and its survival is perceived to be under threat, the credibility of repression can be restored, effectively setting a limit on the expansion of the protest movement.

In this article, we study a model of social protests, where civil society activists try to organize a movement, and a government tries to dissuade them from doing so using the threat of repressive action. When attempting to build such a political protest movement, civil society activists must often make the decision to join asynchronously. Moreover, only when sufficiently many other activists have manifested themselves—or are expected to do so in the near future—does it become worth joining the movement, leading to a dynamic coordination problem. In addition to this dynamic coordination problem, the activists (henceforth, the agents) also face opposition from a government (henceforth, the Principal) with diverging interests. The Principal may try to dissuade the agents from coordinating by taking retaliatory actions against those who join the movement.

Specifically, we study a dynamic model where agents with shared interests are, in turn, each given an opportunity to take a binary action (i.e. joining the political protest movement or not). The marginal benefit of joining the movement is increasing in the number of agents who have (and who will) join it, thus defining a dynamic game of strategic complements among the agents. On the other side, a Principal is given the opportunity to take a repressive action against the agents who engaged in the adversarial activity and we assume she derives a benefit (no matter how small) from repressing offenders, which can be related to countering the actions of the protesters or to potentially halting the growth of the movement. The Principal can only take repressive action on a single occasion, as it could be too politically costly to do it more than once. However, the Principal can choose that time, as in a stopping game.

The Principal also has limited policing resources and can only punish a finite, (possibly small) number of agents. The punishment hence becomes diluted by the size of the protest movement, meaning that the more agents join it, the less likely they are to be punished.

We first show that in the absence of the Principal, the ability of the agents to dynamically coordinate depends on how patient they are relative to the movement's critical mass, i.e. the mass after which it becomes beneficial for any further agent to join. We call this critical mass the movement's *internal coordination threshold*. Indeed, when agents are sufficiently patient, we show by a subgame perfection argument that in any equilibrium, agents all choose to join the movement. This effectively selects a unique equilibrium behavior on the part of the agents.

We next introduce the Principal who can choose if and when to take the (terminal) repressive action. We show that as she can choose that action a single time, and the success of the repression is at best uncertain, she has an incentive to delay taking this option, thereby potentially allowing the protest movement to grow. However, the movement can only grow to reach a certain finite size. Indeed, a non-trivial consideration is whether the Principal is willing to punish—given that such an action is terminal and may not succeed in repressing the movement—and thus whether her punishment threat is credible. In fact, since she could be unable to act again thereafter, punishing may not always be the optimal way forward for her. However, if not punishing would allow enough agents to be in the pool of offenders, so that the expected punishment of future agents becomes sufficiently diluted that the latter would no longer fear it, then the Principal would find punishment optimal (and it would hence be credible in the eyes of the agents).[1]

This allows us to provide an upper bound on how many agents can join the protest movement in any equilibrium. This upper bound depends on what we call the *dissuasion threshold*, i.e. the size of the movement above which the probability of being punished becomes so diluted that agents no longer fear punishment. Interestingly, we note that in any equilibrium, no agent engages in the adversarial activity when the internal coordination threshold is larger than the dissuasion threshold. In other words, when sufficiently many agents need to join the movement for it to become worthwhile, then no agent chooses to join in any equilibrium, anticipating that the government will take the repressive action before agents can even reach that critical mass.

In some environments, agents may not be able to perfectly monitor the actions of other agents and thus to know the actual size of the protest movement. This has non-trivial implications as, with imperfect monitoring, the Principal may have additional reasons to further delay the use of repressive force as the movement grows, so as to keep the size of the movement hidden from agents. However, this can in turn allow the movement to grow beyond the point where it will be possible for the government to maintain the credibility of punishment, thereby making its growth unstoppable. This extra insight illustrates a key role played by the perfect monitoring assumption in our main deterrence result, and it is suggestive of why the Principal may benefit from publicizing (through an independent channel) the information she holds about the protest size.

Our findings contrast with those obtained in the equivalent simultaneous-actions coordination game, in which the size of the Principal's budget (i.e. her ability to punish a large number of agents) is key to ensuring that, in all equilibria, agents are deterred from joining the movement. Obviously, such a simultaneous-actions game lacks the richer dynamic structure, which drives the incentives to delay punishment and the punishment's credibility. In the dynamic setting, a larger budget can indeed incentivize the Principal to further delay the repressive action, as it remains dissuasive for a longer time.

Our analysis of the dynamic game brings what we believe to be novel insights into the literature on deterrence and the credibility of using punishment threats, as pioneered by Schelling (1960). That is, we derive conditions under which the Principal can dispense with explicit commitment devices to deter collective adversarial actions. More specifically, our results establish that when the protest size is commonly observed and reaches a certain threshold, the Principal's threat of using her punishment becomes credible and this in turn deters protesters from reaching that threshold with no need for explicit commitment power on the Principal's side.

We also examine a number of extensions. Namely, we examine a setting where the Principal can, with some probability, halt the growth of protest movement upon taking repressive action (Section 4.1). This introduces a non-monotonic role for the Principal's budget, as her incentives to intervene

[1] This is the one place where we use our assumption that the Principal derives a direct benefit (no matter how small) from punishing offenders, and we can alternatively dispense with this assumption if we assume, as we do in the extensions section (Section 4.1), that repression can automatically halt the protest with some possibly small probability.

earlier or later become conflicted. We also show that the ability to act multiple times can allow for better deterrence (Section 4.4). For instance, if a government were able to spread its resources across time rather than using them only once, its incentives to delay punishment would be reduced. We also extend our model to stochastic decision times, which allows us to conclude that the speed at which the Principal can react to the agents' actions is critical to achieving deterrence (Section 4.5). Indeed, a Principal with large punishment resources, but which reacts slowly, can fail to deter agents from joining the protest movement, as the latter then have time to build up a movement that is sufficiently large to dilute the expected punishment. We also examine the case when the Principal can commit to a particular punishment strategy that can individually target agents (Section 4.6). This case lends itself well to social media activism, for instance, where online posts are time stamped, or to settings involving statements made by politicians or celebrities. We show that the Principal is then allowed to react much more slowly and can still deter the agents from engaging in adversarial activity.

The paper is structured as follows. In Section 2, we describe the dynamic model and payoffs. In Section 3, we characterize the equilibrium behavior of the agents and of the Principal. We state our main results and also compare our dynamic setting to a game where agents act simultaneously. In Section 4, we discuss extensions and robustness checks. Detailed proofs are relegated to Section 5.

1.1 Comparison with related literature

Our paper is related to the literature on dynamic collective action games. See Battaglini and Palfrey (2025), Battaglini and Palfrey (2024) for recent references which, like the other papers in that literature, do not have the analog of a Principal as in our setting. In particular, Battaglini and Palfrey (2025) model dynamic collective actions, but in the absence of an opposing government. They consider private information about heterogeneous costs of contribution in a collective action public good game. Time represents a friction due to discounting, creating a tension between the option value of delay (free-riding) and the necessity of early commitment to signal viability. For finite populations, they show that success is inherently probabilistic and path dependent, as early contributions can signal a lower-cost population. For large populations, they show that success is guaranteed if the required threshold level of participation grows slowly enough relative to the overall population.

In contrast, our model emphasizes the role of a Principal trying to deter agents from joining a collective action. Our result that, without the Principal and in the presence of publicly known contribution costs, agents are able to coordinate on the outcome that is efficient for them is also reminiscent of the work of Gale (1995), who developed a similar insight in a different context of private provision of public goods with asynchronous decision times. Our results in the presence of the Principal have no counterpart in the literature as far as we know.[2]

A recent contribution modeling the interplay between social activists and a government is Correa (2025), who models a dynamic collective action game where actions are repeated in every time period, and thus capture not only joining but also remaining in or leaving the movement. This model also allows to describe the dynamics of the movement's gradual decay. In our setting with irreversible actions, the dynamics rather emphasizes the formation of a social movement and how large it can grow when repression is possible. In Correa (2025), the government's action is to either wait or make concessions to the movement, in the hope that it will dissolve, leading to a war of attrition phenomenon. In our case, the government's action is rather to take repressive action, and it is this threat of punishment that can prevent a movement from exceeding the threshold at which the government would credibly take that repressive action. Specifically, our theory suggests how a terminal repressive action, combined with

[2]Bueno de Mesquita, Myatt, Smith, and Tyson (2024) study a game of participation in a collective action in the presence of punishment, but in a static context.

diluted punishment, can incite a government to delay its response, thereby allowing a protest movement to reach a certain size, and also suggests settings in which repression can take place or in which the movement can grow undeterred.

While some work supports the idea of strategic substitutes for protest participation (Cantoni et al. (2019)), many recent models also assume strategic complements. As in Correa (2025), we model agents' actions as exhibiting strategic complements. In combination with the "safety in numbers" coming from the dilution of punishment, this leads to a natural setting to model certain strategic aspects of collective protests under the threat of governmental repression.

Recent empirical work on social protests include Cantoni et al. (2024), who analyze global protest data (1.2 million events across 218 countries from 1980-2020) to document some patterns regarding political mobilization. They namely find that movements exhibit rapid geographical spikes and slower declines, with the possibility of crackdowns (or the threat thereof) shortening their persistence in more authoritarian societies. This is consistent, in our model, with protest movements capping at a smaller size when the Principal can either halt the movement when taking repressive action (as in the extension presented in Section 4.1), or act multiples times instead of only once (as in the extension presented in Section 4.4). Indeed, in more authoritarian societies the nature of repression is more likely to be able to quell a social movement than in more democratic ones. Moreover, limiting repressive action to a single instance is better suited to a government operating in a democratic regime, where repression is politically costly and thus delayed until the movement has reached a larger size and time span. In a less democratic regime, a government can arguably take repressive action multiple times, as it faces a smaller political cost of doing so.

Our setting, and namely aspects related to the imperfect observability of actions, is also related to a literature on designing extensive forms for strategically interacting agents (namely, Nishihara (1997), Doval and Ely (2020), Salcedo (2017), Chen, Ghersengorin, and Petersen (2024) or Koh, Sanguanmoo, and Uzui (2024)). In this literature, equilibrium outcomes depend critically on the endogenous design of timing, information, and observability, illustrating that suitably chosen extensive forms and information structures can expand or refine the set of sustainable outcomes beyond those predicted by standard formulations.

Our extension that models stochastic decision times (Section 4.5) is somewhat similar to that adopted in the literature on revision games (Kamada and Kandori (2020)), in which players' ability to change their actions is modeled in a stochastic fashion using Poisson distributions. Other related models include Burdzy, Frankel, and Pauzner (2001), Frankel and Pauzner (2000) or Calvo (1983). However, to our knowledge, that literature has not considered the kind of games involving both a Principal and agents as in our setting.[3]

Less directly related to our model, one could mention static approaches of coordination games allowing for selection based on incomplete information. See, in particular, the global games approach of Carlsson and Van Damme (1993), Morris and Shin (1998, 2001) or more recently Persico (2023). There is also an approach called Poisson games, introduced in Myerson (1998, 2000), where there is uncertainty about the number of players (drawn according to a Poisson distribution, hence the name of this approach), although the game itself is static.[4] These are obviously different perspectives from the one we develop here, which is based on the dynamic nature of decision making rather than asymmetric or incomplete information. Finally, our paper is peripherally related to the literature on community enforcement (e.g. Kandori (1992), Takahashi (2010), Kandori and Obayashi (2014)).

[3] Calcagno, Kamada, Lovo, and Sugaya (2014) consider revision coordination games and obtain the selection of the Pareto-dominant equilibrium similarly as in Gale (1995). However, there is no analog of a Principal in their setting, which is the main focus of our study.

[4] See also Frankel (2023) for a treatment of participation games.

2 Model

2.1 Setting

We study a setting where at different times, different agents (e.g. civil society activists) get the chance to take a costly action and their payoff depends on the total number of agents who choose this action. In effect, they play a dynamic coordination game. There are T agents getting the chance to play the game at times $t = 1, 2, ..., T$, where we allow T to be infinite. At such a time, the t-th agent will get the chance to choose an action $a_t \in \{0, 1\}$. This action represents whether the agent joins a social protest movement.

A Principal (e.g. a government or a police force)[5] with interests diverging from those of the agents also gets the chance to police the population of agents at each time t, immediately after the agent. This Principal holds a budget $B \in \mathbb{N}^+$ and at each such time t, she can choose the repressive action. If she does, she can no longer act thereafter. Formally, at each time where she can still act, the Principal can choose an action $a_{P,t} \in \{0, 1\}$ where $a_{P,t} = 1$ means that the Principal punishes (e.g. arrests) up to B agents picked uniformly at random among those who previously chose $a_s = 1$ (for $s \leq t$), and she can no longer act thereafter. $a_{P,t} = 0$ means that the Principal does not take punitive action.[6]

Let the *full history* of play at time t be denoted by[7]

$$h_t = (\{a_j\}_{j \leq t}, \{a_{P,j}\}_{j \leq t}).$$

We call $\mathcal{H}$ the set of possible histories.

Now call $h_{P,t} = (\{a_j\}_{j \leq t}, \{a_{P,j}\}_{j \leq t-1})$ a *Principal's history* at time t. This is the history observed by the Principal at time t, upon making her decision. We call $\mathcal{H}_P$ the set of possible Principal's histories. A Principal's strategy is then $\sigma_P : \mathcal{H}_P \to \Delta(\{0, 1\})$. That is, at some action time t, based on a history of past play $h_{P,t}$ observed by the Principal, she can then choose to take repressive action (i.e. $a_{P,t} = 1$), if she has not done it before. She can alternatively decide not to take repressive action (i.e. $a_{P,t} = 0$). She can also choose a mixed strategy.

Likewise, call $h_{A,t} = (\{a_j\}_{j \leq t-1}, \{a_{P,j}\}_{j \leq t-1})$ an *agents' history* at time t. This is the history observed by the agent at time t, upon making his decision. We call $\mathcal{H}_A$ the set of possible agents' histories. The agents' strategy is then $\sigma_A : \mathcal{H}_A \to \Delta(\{0, 1\})$. That is, at some action time t, based on a history of past play $h_{A,t}$ observed by agent t, he can then choose to take action $a_t = 1$ or $a_t = 0$, and he can also randomize.

2.2 Payoffs

Given action times $t \in \{1, 2, ..., T\}$ for the agents and Principal, let us denote the agents' and Principal's action profiles as $\vec{a} = (a_1, a_2, ..., a_T)$ and $\vec{a}_P = (a_{P,1}, a_{P,2}, ..., a_{P,T})$.

It will be useful to also define the *running* action profiles at time t for the agents and for the Principal as $\vec{a}_t$ and $\vec{a}_{P,t}$. Here $\vec{a}_t = \{a_{1,t}, a_{2,t}, ..., a_{T,t}\}$, where $a_{j,t} = a_j \in \{0, 1\}$ if $j \leq t$ and thus agent j has already acted. By default, $a_{j,t} = 0$ if $j > t$ and the agent has not yet acted. Likewise $\vec{a}_{P,t} = \{a_{P,1,t}, a_{P,2,t}, ..., a_{P,T,t}\}$ with $a_{P,k,t} = a_{P,k} \in \{0, 1\}$ if $k \leq t$ and, by default, $a_{P,k,t} = 0$ if $k > t$.

[5]In the extensions section (Section 4) and throughout the paper, we also discuss other applications of the model, where the Principal can assume other roles than that of a government (e.g. a lobby, an NGO, etc.).

[6]More generally, we could allow the Principal to choose to target those who chose $a = 1$ or those who chose $a = 0$, but our payoff specification will make the latter suboptimal for the Principal.

[7]Note that, while the setting studies a stopping problem, we choose to record all actions $\{a_{P,j}\}_{j \leq t}$ by the Principal in a history, instead of just the actual stopping time. This is because the setting also lends itself to extensions with multiple repressive actions, as studied in Sections 4.4 and 4.5.

Note that we will sometimes denote by $n_t = \sum_j a_{j,t}$ the size of the social protest movement at time t. It is the sum of the running actions of the agents at time t. We will also denote by $n_{-i,t} = \sum_{j \neq i} a_{j,t}$ the size of the movement without agent i, which is the sum of the running actions of agents other than i at time t.

2.2.1 Agents' payoffs

At any time t, agent i receives a flow payoff

$$\tilde{\pi}_{i,t}(a_{i,t}, \vec{a}_{-i,t}, \vec{a}_{P,t}) = v(a_{i,t}, n_{-i,t}) - \kappa \cdot a_{i,t} - C \cdot \phi_{i,t}. \tag{1}$$

In the above equation, $v : \{0,1\} \times \mathbb{N} \to \mathbb{R}$ is the benefit function, which is increasing in both own running action $a_{i,t}$ and in the sum of the running actions of other agents $n_{-i,t} = \sum_{j \neq i} a_{j,t}$. Thus, at time t an agent benefits from the actions of all the agents who chose $a_j = 1$ up to time t. $\kappa \geq 0$ is the intrinsic cost to agent i of taking action $a_i = 1$ and we allow for the possibility that this cost be null. C is the punishment cost felt by agent i if he is punished by the Principal and $\phi_{i,t} \in \{0,1\}$ is the indicator that agent i has been punished by time t, i.e. $\phi_{i,t} = \mathbb{1}_{\{\exists \tau \leq t : \tilde{\phi}_{i,\tau}=1\}}$, where $\tilde{\phi}_{i,\tau} = 1$ if agent i is punished at time τ and $\tilde{\phi}_{i,\tau} = 0$ otherwise. Without loss of generality, we let $v(0,0) = 0$.

$\tilde{\phi}_{i,\tau}$ is then a Bernouilli random variable such that if the Principal chooses to take repressive action at time τ (i.e. $a_{P,\tau} = 1$), and if agent i has chosen action $a_i = 1$ at some time $i \leq \tau$, then

$$\tilde{\phi}_{i,\tau} = \begin{cases} 1 & \text{w.p. } \min(\frac{B}{n_\tau}, 1) \\ 0 & \text{w.p. } \max(\frac{n_\tau - B}{n_\tau}, 0). \end{cases}$$

where $n_\tau = \sum_j a_{j,\tau}$ the size of the social protest movement at time τ.

We emphasize that $\tilde{\pi}_{i,t}$ being a flow payoff, the punishment cost (if the agent has been punished in the past) is felt in every period and not only once.[8]

Moreover, we assume that $\Delta v(n_{-i,t}) = v(1, n_{-i,t}) - v(0, n_{-i,t})$, the individual net benefit of choosing $a = 1$ when $n_{-i,t}$ other agents choose $a = 1$, is increasing in $n_{-i,t}$ capturing a strategic complementarity in the interaction between agents. We also assume that there is some $N < T$ such that $\Delta v(n_{-i,t}) > \kappa$ for $n_{-i,t} \geq N - 1$, thereby implying that when N or more agents choose $a = 1$, it is individually best to choose $a = 1$ (in the absence of the Principal's intervention). We will refer to N as the *internal coordination threshold*, i.e. the size above which a social movement becomes attractive and can take hold, in the absence of any risk of punishment.

For some insights to be developed next, we also assume that $v(1,0) < \kappa$, in which case the static version of the game in the absence of the Principal has the structure of a coordination game, with one equilibrium in which all agents choose $a = 0$ and one equilibrium in which all agents choose $a = 1$. We also assume that $v(1, T-1) - v(0, T-1) - \kappa < C$ so that an agent always suffers from being punished, irrespective of how many other agents have chosen action $a_i = 1$. These properties of v are summarized in the following assumption. When there are infinitely many agents, we will assume that $\lim_{n_{-i,t} \to \infty} \Delta v(n_{-i,t})$ is finite and denote it by $\Delta v(\infty)$.

Assumption 1 (Properties of benefit function) *(i)* $v(0,0) = 0$. *(ii) Let* $\Delta v(n_{-i,t}) = v(1, n_{-i,t}) - v(0, n_{-i,t})$. $\Delta v(n_{-i,t})$ *is increasing in* $n_{-i,t}$, *with* $\Delta v(N-1) > \kappa$ *for some minimal* $N \in \mathbb{N}_+$ *with* $N < T$, *and* $\Delta v(T-1) - \kappa < C$. *We call* N *the* internal coordination threshold.

The forward-looking, discounted realized payoff at time t is then

[8] Paying the punishment cost for all future periods can be interpreted, for instance, as being put in prison.

$$\pi_{i,t}(a_i, \vec{a}_{-i}, \vec{a}_P) = \sum_{s=t}^{T} \delta_A^{s-t} \tilde{\pi}_{i,s}(a_{i,s}, \vec{a}_{-i,s}, \vec{a}_{P,s}), \tag{2}$$

where $\delta_A \in (0,1)$ is an agent's discount factor.

At his decision time $t = i$, agent i will thus choose a strategy $\sigma_A^*(h_{A,i})$ to maximize his expected payoff $\mathbb{E}[\pi_{i,i}(a_i, \vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,i}]$, given the Principal's strategy, the other agents' strategy, and an agents' history of play at time i.

2.2.2 Principal's payoff

The Principal's flow payoff at time t is

$$\tilde{\pi}_{P,t}(\vec{a}_{P,t}, \vec{a}_t) = -n_t + R(B)\psi_{s,t}, \tag{3}$$

where as already mentioned $n_t = -\sum_j a_{j,t}$ is the size of the protest movement at time t. Moreover, $R(B) > 0$ is an intrinsic benefit of taking the repressive action, which scales with the budget B (formally, $R : \mathbb{N} \to \mathbb{R}_+$). Finally, $\psi_{s,t} = \mathbb{1}_{\{(a_{P,s,t}=1) \bigcap (a_{j,s}=1, \forall j \in \{j:\tilde{\phi}_{j,s}=1\})\}}$ is the indicator that the Principal has taken repressive action at some time $s \leq t$ and targeted punishment at agents among those who had already joined the movement by time s.

We see, from the first term of Eq. (3), that the Principal suffers permanent disutility from all the agents who have chosen action $a_j = 1$ in the past, capturing her interests that diverge from those of the agents. Moreover, from the second term of Eq. (3), we see that she enjoys a permanent benefit $R(B)$ from having punished agents who had chosen action $a_j = 1$ before the repressive action took place.[9]

The Principal's forward-looking, discounted realized payoff at time t is then

$$\pi_{P,t}(\vec{a}_P, \vec{a}) = \sum_{s=t}^{T} \delta_P^{s-t} \tilde{\pi}_{P,s}(\vec{a}_{P,s}, \vec{a}_s), \tag{4}$$

where $\delta_P \in (0,1)$ is the Principal's discount factor.

The Principal will thus choose a strategy σ_P^* that maximizes her expected payoff $\mathbb{E}[\pi_{P,t}(\vec{a}_P, \vec{a})|\sigma_P, \sigma_A, h_{P,t}]$, given the agents' strategy and a Principal's history of play at time t.

2.2.3 Interpreting payoffs and assumptions

On the Principal acting only once and having a finite budget B

Taking punitive action against a social protest movement can be highly controversial and politically risky for a government. This motivates our assumption that the Principal can only take repressive action once, as doing so repeatedly could make the government so unpopular that it risks falling if it were to act again. This will also allow us to better analyze the incentives of the government to delay punishment. Additionally, in many cases, it seems reasonable to assume that the Principal lacks the resources to punish every agent in a large social movement, which is why we introduce a budget constraint, B, typically smaller than the maximal size of the movement, T. We believe this setting lends itself well to a range of applications in the political economy of social protest movements. Our model will explore how the Principal can strategically use these limited resources to effectively deter a collective action. In an extension, in Section 4.4 we consider the case in which the Principal can act at multiple different

[9] Several variants on this second term could be considered without altering the analysis. What is essential is that the Principal has an intrinsic extra preference (no matter how small) for punishing those who chose $a = 1$ rather than those who chose $a = 0$ as this inventivizes her to punish those in the former pool rather than those in the latter, thereby justifying our above formulation.

times, each time reducing her budget by one increment. While we note that in this alternative scenario the punishment cost C and the budget B play a different role, we observe that our basic insights about the ability of a government and the extent to which it can deter a protest movement remain unchanged compared to the basic scenario in which the Principal can act only once.

On random punishment

Another assumption we make in the payoffs is that the Principal punishes a randomly selected set of faulty agents. This implies that the Principal does not use information about the timing or order in which agents joined a social movement to determine punishment. This assumption is valid in many political economy applications. Indeed, a government may not have the resources or infrastructure to track the precise timing of when agents joined a movement, especially in large and decentralized movements. In the absence of surveillance technology that can track an individual's involvement over time, the Principal would likely make decisions based on available evidence of participation in the movement, without regard to the timing of that participation. Given this, using random selection for punishment may be a practical necessity. As will become clear, this random punishment will underpin what we call the "dilution" of punishment, i.e. the fact that given finite government resources, the repressive action becomes less of a deterrent as more agents join the social protest.

At a formal level, our formulation implies a form of imperfect recall for the Principal, to the extent that she acts after every agent and thus with perfect recall she could keep track of who did what and when in the past. Yet, there is another interpretation of our assumption, not based on memory imperfections. Specifically, assume that the Principal does not implement herself the strategy and that she uses a third party (a police force, say) to do it. Our random punishment assumption fits the scenario in which this third party would not have access to the entire historical record, which would be known only to the Principal and other intelligence branches, and it may be too costly or complicated to transmit this precise information.

Finally, we also note that agents being homogeneous, a Principal does not have a particular reason to punish specific protest members. Thus, the credible use of punishment targeted at specific agents rather than others would require some commitment ability. Such considerations are discussed as extensions in Section 4.6.

On the intrinsic benefit $R(B)$ derived from punishing

We noticed that the term $R(B)$ represents the Principal's intrinsic (i.e. exogenous) satisfaction gained from taking the repressive action. While $R(B) > 0$ may seem like a restrictive assumption, it turns out to be unnecessary to generate our results. Indeed, the reason this assumption is needed in the basic model presented so far is that taking the repressive action only punishes some current members of the protest movement and has no effect on the ability of the Principal to deter future agents from joining the protest movement after the repressive action has been taken. A strictly positive exogenous benefit $R(B) > 0$ is then needed to make the repressive action optimal for the Principal, as she would otherwise be indifferent between choosing $a_P = 1$ and $a_P = 0$. In a model extension presented in Section 4.1, we allow the repressive action to effectively halt, with some probability, the growth of the movement. This then creates, endogenously, an expected benefit from taking action $a_P = 1$, which eliminates the need for the exogenous benefit $R(B) > 0$, and can even allow this parameter to be negative (i.e., a "cost" $R(B) < 0$, instead of a benefit).

The reason we chose to restrict $R(B)$ to be strictly positive in the basic model is that such a simpler model allows us to clearly illustrate a central game theoretic insight, since this net benefit that the Principal gains from punishing protesters (be it exogenous or endogenous) is key to ensuring the *credibility* of punishment, and thus in effectively deterring agents from coordinating. Indeed, in a context in which the Principal would merely incur a cost when taking punitive action, she would have an incentive to constantly delay punishment until the next offense, which would then allow agents to coordinate.

Nevertheless, it is easy to find realistic interpretations for the strictly positive parameter $R(B)$ in the basic model. Suppose, for example, that $R(B)$ has a linear form such that $R(B) = rB$, with $r \in (0,1)$. In the context in which social movements take the form of making public announcements (say in favor or against some policy), this could be understood as follows: When the Principal punishes an agent, she can partially "annul" that agent's action in the future, resulting in its net effect now being $-1+r$ instead of -1. For example, this could involve the Principal issuing a public statement that contradicts what a political activist had previously said. While this may not convince other activists (as it does not feed back into their payoff functions), it can influence the wider society, which is beneficial to the Principal. In this sense, the collective action taken by a set of up to T agents can be seen as a subset of the broader population, most of whom will never be directly involved in the collective action, but whose opinions can be influenced by it. Thus, beyond merely discouraging activists from organizing by raising the risk of punishment, the Principal's repressive action can also serve to (partially) undo the social movement's influence on society at large, which is a benefit for the Principal.

It is interesting to note that our model can apply to contexts beyond just social movements. For example, in the case of school bullying, a student who has been bullied by others can derive satisfaction from taking revenge on those who harassed him. In this context, punishing someone after being insulted can be understood as partially repairing a wound to one's honor, for example.[10]

3 Equilibrium analysis

We will analyze the Subgame Perfect Nash Equilibria of the game described in Section 2.

3.1 Successful agent coordination

We will first briefly examine the agents' behavior in the absence of a Principal. The dynamic nature of the game allows us to select a unique equilibrium behavior for the agents. Namely, if the agents are patient enough, then in the absence of a Principal, they always succeed in coordinating on action $a = 1$. This is formalized in the following lemma.

Lemma 1 (Successful agent coordination) *In the game without a principal, there exists $\overline{\delta}_A \in (0,1)$, such that when $\delta_A > \overline{\delta}_A$, any equilibrium involves $a_i^* = 1$ for all i.*

Note that, in contrast with a static version of the game outlined in Section 3.4, the dynamics can allow us to select a unique equilibrium behavior on the part of the agents.

To gain some intuition into Lemma 1, note that in the absence of a Principal (and thus with no chance of ever being punished), agents always have an interest in choosing $a = 1$. Indeed, when agents are sufficiently patient, and thus the future actions of subsequent agents are not too severely discounted by the discount factor δ_A^{s-t}, then by choosing $a_1 = 1$, agent 1 precipitates a subgame in which agents $i = 2, ..., N-1$ also choose $a_i = 1$, as it then becomes strictly dominant for agent N (and all subsequent agents) to choose $a_N = 1$. It is thus strictly dominant for all agents to choose $a_i = 1$. In other words, the early agents effectively have an incentive to initiate a herding behavior by the subsequent agents, so as to reach the internal coordination threshold N. This allows agents to coordinate dynamically.

[10]From a different perspective, one could also argue that a code of honor according to which one directly benefits from punishing offenders is evolutionary adaptive, to the extent that it allows to deter in the first place (as applies in general but also in our context, as we next show).

This result that agents are able to coordinate on the outcome that is efficient for them in a dynamic coordination game is reminiscent of the work of Gale (1995), who provided a similar insight, but in the different context of the private provision of public goods with asynchronous decision times.[11]

3.2 Successful deterrence by the Principal

We next consider the case in which the Principal can act immediately after each agent.

3.2.1 An upper bound on the number of agents who can coordinate

If the Principal could commit to punishing agents whenever at least one of them has chosen $a = 1$, then doing so would be optimal for the Principal as she would then deter any agent from choosing $a = 1$ in the first place. However, such a commitment ability is not available to the Principal. As it turns out, it cannot be credible for the Principal to threaten to use her punishment device in all such circumstances.

To see this concretely, assume that the Principal's punishment budget allows her to punish only one agent (i.e. $B = 1$). Suppose by contradiction that, no matter how many agents have already chosen $a = 1$, the Principal were to always find it optimal to punish the offenders. Then, in the subgame where the first agent has chosen $a = 1$, it is not optimal for the Principal to punish when C satisfies the condition $\Delta v(T-1) - \kappa < \frac{C}{2}$, which holds for large enough C.

Indeed, if the Principal chooses to punish, she would no longer be able to punish in the future and all subsequent agents would then choose $a = 1$ (by Lemma 1).

By contrast, if the Principal chose not to punish at $t = 1$, she would ensure that no other agent will ever choose $a = 1$ in the future. This is so because if a second agent were to choose $a = 1$, he would expect to be punished with probability $\frac{1}{2}$. However, the condition $\Delta v(T-1) - \kappa < \frac{C}{2}$ implies that this second agent (and any subsequent agent) would strictly prefer not to choose $a = 1$. This yields the desired contradiction, as in the subgame considered above, the Principal would strictly prefer not to punish whenever only the first agent has chosen $a = 1$.

The next proposition takes into account the credibility constraint of the punishment threat and it establishes an upper bound on the number of agents who can possibly choose $a = 1$ in any equilibrium assuming the number of agents is infinite.

First, we introduce an important definition that will allow us to capture this credibility constraint.

Definition 1 (Dissuasion threshold) *Let $T = \infty$. Given a punishment cost C felt by agents and a Principal's budget B, we call* dissuasion threshold *the integer*[12] *m such that $\frac{BC}{m} < \Delta v(\infty) - \kappa < \frac{BC}{m-1}$.*

This *dissuasion threshold* is the critical size of a social movement, above which punishment becomes too diluted to be dissuasive. We note that it depends on the agent's benefit function v, as well as on the Principal's budget B and on the severity of the punishment C.[13]

Proposition 1 (Partial deterrence) *Let $T = \infty$ and let m be the dissuasion threshold. There exist $\overline{\delta}_A \in (0,1)$ and $\overline{\delta}_P \in (0,1)$ such that if $\delta_A > \overline{\delta}_A$ and $\delta_P > \overline{\delta}_P$, then in any equilibrium there are at most $m-2$ agents choosing $a_i^* = 1$.*

[11]Note also that in our case, the participation cost κ is homogeneous and publicly known. In contrast, Battaglini and Palfrey (2025) show that, when participation costs are heterogeneous and private, successful coordination becomes path dependent and can fail to happen since low initial participation can signal a high cost structure, which may discourage further agents from participating. We do not focus on such a setting so as to emphasize what comes next: The role of the Principal in deterring coordination among agents.

[12]Since it follows by Assumption 1 that $\Delta v(\infty) - \kappa < BC$, and since $\frac{BC}{m}$ is decreasing in m, it follows that there can only be a single m that satisfies the condition $\frac{BC}{m} < \Delta v(\infty) - \kappa < \frac{BC}{m-1}$. We may also impose the genericity assumption that $\frac{BC}{\Delta v(\infty)-\kappa}$ is not an integer, which then garanties that there exists m satisfying these two strict inequalities.

[13]Note that we chose $T = \infty$ to obtain a simpler expression for this dissuasion threshold, as will be discussed later.

Proposition 1 highlights how both finite policing resources and the size of the social movement at a given time shape the credibility and effectiveness of punishment. To avoid taking the (terminal) repressive action—and becoming unable to act thereafter—the government may prefer to withhold punishment as long as it remains a credible deterrent in the future—i.e., as long as the movement has not grown so large that punishment becomes too diluted and hence not dissuasive any more. Recognizing this, agents will not take the threat of punishment seriously until the movement has expanded sufficiently, such that punishment could soon lose its deterrent effect. At this point, they realize that the government will have no option but to act, and the threat of punishment becomes credible. This limits the size of the social movement in equilibrium.

More specifically, Proposition 1 establishes a link between the long-run potential benefit of the collective protest as measured by $\Delta v(\infty) - \kappa$, the cost of being punished as measured by C, the risk of being punished as determined by the budget B—which are all primitives of our model—and an upper bound on the number of agents who can possibly choose $a = 1$ in any Subgame Perfect Nash Equilibrium.

Intuitively, the upper bound is derived so that two conditions are satisfied. On the one hand, if two more agents beyond this bound were to join the protest, the diluted punishment that results from the randomized implementation would be perceived as less damaging than the prospect of joining the protest under the assumption that all remaining agents will join it as well (i.e., $\frac{BC}{m} < \Delta v(\infty) - \kappa$). As a result, the punishment would be ineffective in deterring the action $a = 1$ from then on, which would make the Principal prefer triggering the repressive action immediately (since she benefits from punishing offenders as $R(B) > 0$, and the sooner the better as $\delta_P < 1$). On the other hand, if one more agent were to join the protest, the punishment (which is diluted, but less than when two more agents join) would be more harmful than the long-run benefit of choosing $a = 1$ when all remaining agents choose $a = 1$ (i.e., $\Delta v(\infty) - \kappa < \frac{BC}{m-1}$). This agent would then find it optimal to choose $a = 0$ if he expected the punishment to be triggered. Putting these two observations together yield Proposition 1.

Note that in Proposition 1, we have assumed that $T = \infty$. This is to make the expression of the benefit of joining the protest—when all remaining agents are expected to join—the same irrespective of the current time t. The same result would otherwise hold for finite T.[14]

3.2.2 On the role of the budget B and the punishment cost C felt by agents

While it might seem that a larger budget B and a larger punishment cost C felt by the agent are always beneficial to the Principal, an interesting aspect suggested by Proposition 1 is that they are not necessarily more dissuasive and do not necessarily lead to a smaller number of agents choosing $a = 1$ in equilibrium.[15]. As it turns out, the magnitude of BC can be viewed as playing a dual role. On the one hand, the budget B and the punishment cost C serve the role of deterring the next agent from choosing $a = 1$ when the Principal is expected to take the punishment action. For that purpose, a large BC is more dissuasive and thus beneficial to the Principal. On the other hand, a small enough BC guarantees that at some point (i.e. when enough agents have already chosen $a = 1$), the expected punishment will no longer be sufficient to deter subsequent agents, after one additional agent chooses $a = 1$. This smaller BC is then beneficial to the Principal, as it makes the strategy of punishing at this particular point in time (rather than later) optimal for her, and hence *credible*. Proposition 1 puts these two considerations together by relating—through the dissuasion threshold m—the budget B and the punishment cost C to

[14] If a finite T were considered instead, the bound on the remaining number of agents choosing $a_i = 1$ would depend on the number of periods left, $T - t$, as m would depend on $T - t$. The same conclusion as in Proposition 1 would hold replacing $\Delta v(\infty)$ by $\Delta v(T - t + 1)$ and saying that from t onwards, there can be at most $m - 2$ agents choosing $a_i^* = 1$.

[15] To see this more explicitly, consider a case in which $N = 1$ (say, because $\kappa = 0$). When $\frac{BC}{2} < \Delta v(\infty) - \kappa < BC$, we know by Proposition 1 that no agent will choose $a = 1$ in equilibrium. However, when $\frac{BC}{m} < \Delta v(\infty) - \kappa < \frac{BC}{m-1}$ for some $m > 2$ (and thus a larger BC), it is easily verified that one can construct an equilibrium in which the first $m - 2$ agents choose $a = 1$. Thus, a larger BC is not always good for the Principal.

the number of agents who can possibly choose $a = 1$ in equilibrium.

We now formalize the relationship between the dissuasion threshold m and BC in the following lemma.

Lemma 2 (Comparative statics in B and C) *The dissuasion threshold m is non-decreasing in BC.*

Lemma 2 thus illustrates the role of the Principal's budget and the punishment cost felt by an agent have on credible dissuasion. The larger they are, the more a Principal who can take repressive action only once will be tempted to delay taking it, since the latter remains credible for longer.[16] In the case of social protest movements, the punishment is typically arrest and possibly imprisonment, which means that C can be quite high. Even if governments have the resources to punish a substantial number of people, protest movements can reach a large size before any repressive action is expected to be taken. This is evident, namely, in western societies, where taking repressive action is controversial and the government can thus indeed have an incentive to delay it for as long as possible, until it would cease to be dissuasive.

It is important to note that the model developed so far only allows the possible repressive action to dissuade agents through the threat of punishment. In reality, a government's ressources could be devoted not only to "punishment" in the form of imprisonment or lawsuits, but also to dismantling the very organization of a social movement to prevent its future growth. For instance, the repressive action could also focus on terminating the social protest movement by dissolving its leadership structure, seizing its assets and financial resources, etc. The basic model presented so far does not account for that, to allow us to focus only on the purely dissuasive aspect of a repressive action. However, in Section 4.1, we extend the model to allow the Principal's budget to also be used to terminate the protest movement, with some probability $p(B)$ increasing in B. We show that the budget B can then have a non-monotonic role: Incentivizing the Principal to delay punishment so that the protest movement can reach a certain positive size (when B is low enough), and incentivizing her to act very early and thus allowing her to dissuade all agents from ever joining the protest movement (when B is high enough). This is due to the countervailing effects of the budget on dissuasion (as studied so far) and on disrupting the social movement (as permitted in this extended model).

3.2.3 Successful deterrence of all agents

Proposition 1 establishes an upper bound on how many agents can choose $a = 1$ in equilibrium and this upper bound is solely determined by $\frac{\Delta v(\infty)-\kappa}{BC}$. When this upper bound is smaller than the number of agents N required for coordination to be worthwhile (absent the punishment cost) as considered in Assumption 1, we can strengthen Proposition 1 by establishing that, in any equilibrium, no agent chooses $a = 1$.

Proposition 2 (Threshold criterion for full deterrence) *Let $T = \infty$, let m be the dissuasion threshold, and N the internal coordination threshold. If $m - 2 < N$, there exist $\overline{\delta}_A \in (0, 1)$ and $\overline{\delta}_P \in (0, 1)$ such that if $\delta_A > \overline{\delta}_A$ and $\delta_P > \overline{\delta}_P$, then in any equilibrium all agents choose $a_i^* = 0$.*

When the internal coordination threshold N is sufficiently larger than the dissuasion threshold m (i.e. when $N > m - 2$), then agents are dissuaded from joining the social protest movement, as its size will never reach the point where it could take off anyway. Indeed, by Proposition 1, we know that there can be at most $m - 2$ agents choosing $a = 1$. But if $m - 2 < N$, the conditions of Assumption 1 imply that it is not worth choosing $a = 1$ for any agent if at most $m - 2$ other agents are to choose $a = 1$.[17]

[16] Although a very different application, it is interesting to note the parallel with the context of nuclear deterrence, where such considerations have been central (see Schelling (1960)). The idea of a very large C is also obviously plausible with nuclear weapons, which preserves their dissuasive nature, even when their use is constantly delayed. Combining this with the "terminal" nature of using such weapons also creates a similar incentive to delay using them, until all appears to be lost.

[17] It is interesting to draw a parallel with Battaglini and Palfrey (2025). In their setting with private information about

3.3 On successful agent coordination with imperfect monitoring

An important assumption made in the main model is that agents can observe the actions previously chosen by other agents. We note that our results carry over as long as the agents can observe the total number of agents who previously chose $a = 1$. This is what allows agents to know if the Principal will take action $a_P = 1$ or delay, and thus makes the deterrent effective when the dissuasion threshold is about to be reached. In this section, we illustrate how the possibility of a successful protest is affected when agents imperfectly observe this number while still assuming that the Principal perfectly observes it (as sounds natural if we have in mind that there are state agencies that are explicitly devoted to collecting such information). However, we naturally assume that agents observe the Principal's repressive action when it takes place.

More precisely, we will illustrate in this monitoring scenario that a protest with nearly all agents choosing $a = 1$ can possibly arise in equilibrium, in contrast to our findings in Propositions 1 and 2.

Multi-stage games with imperfect private monitoring are known to be hard to analyze because of the induced complexity as a history gets longer (the private history gets more and more complex as time elapses). We will illustrate our insight in the context in which what is observed by agents is commonly observed by all agents as well as the Principal (who, in addition, observes the number of agents who previously chose $a = 1$). Suppose that agents observe the calendar time t, the Principal's actions $a_{P,t}$, and a binary signal s_t about the current size of the protest movement. That is, the history commonly observed by agents at the beginning of period t, just before agent t gets to choose his action, is $h_{A,t} = (\{s_j\}_{j\leq t}, \{a_{P,j}\}_{j\leq t-1})$. The principal, however, observes as before the history $h_{P,t} = (\{s_j\}_{j\leq t}, \{a_j\}_{j\leq t}, \{a_{P,j}\}_{j\leq t-1})$.

We now introduce the signal structure. Let $n_{t-1} = \sum_{s<t} a_s$ be the number of agents who have chosen to join the movement before time t. The agent observes a signal as follows

$$s_t = \begin{cases} 1 & \text{w.p. } g(n_{t-1}) \\ 0 & \text{w.p. } 1 - g(n_{t-1}), \end{cases}$$

where $g : \mathbb{N} \to [\underline{g}, \overline{g}]$, with $0 < \underline{g} < \overline{g} < 1$. Here $g(n_{t-1})$ is a non-decreasing function of n_{t-1}.[18] The more agents have already joined the movement, the more likely it is that they observe such a signal $s_t = 1$.

We begin by observing that the Principal will consistently face an incentive to postpone repressive action, so long as she holds sufficiently optimistic expectations regarding the behavior of future agents and her benefit $R(B)$ is not too large. Since repression is a terminal move—after which the government can no longer intervene—there is thus a temptation to wait, in the hope that future agents will indeed act in accordance with her expectations.

This is formalized in the following lemma.

Lemma 3 (Delayed punishment when one future agent behaves well enough) *Let $h_{P,t}$ be a history on which the Principal has not already taken repressive action. Call $\mu_{t'} = \mathbb{E}[a_{t'}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]$ the probability that the agent takes action 1 at time t', given that the Principal does not take repressive action at time t. There exist $\underline{R} > 0$, $\overline{\delta}_P \in (0, 1)$, $\overline{T} > 0$, and an upper bound $\underline{\mu}(\delta_P, R(B)) \in (0, 1)$ such that, when $T > \overline{T}$, $\delta_P > \overline{\delta}_P$, $R(B) < \underline{R}$, and if $\mu_{t'} < \underline{\mu}(\delta_P, R(B))$, then the Principal chooses strategy $\sigma_P(h_{P,t}) = 0$ for all $t < t'$. Moreover, $\underline{\mu}(\delta_P, R(B))$ is increasing in δ_P and decreasing in $R(B)$.*

participation costs, agents also have the equivalent of an internal coordination threshold (N) and successful coordination is guaranteed—as the total number of agents grows arbitrarily large (i.e., $T \to \infty$)—if the ratio N/T decreases sufficiently fast. As previously discussed, their setting does not involve a Principal trying to dissuade the agents. In our case, it is rather the ratio of the dissuasion threshold—which arises precisely because of the presence of a Principal—to the internal coordination threshold, m/N, that must be small enough for the Principal to prevent the agents from coordinating.

[18] The particular case in which g is constant corresponds to agents not observing any informative signal at all.

Naturally, by induction, Lemma 3 also applies if more than one agent is expected to behave sufficiently well in the future.

With imperfect monitoring, the Principal has greater incentives to refrain from taking repressive action, as this allows her to conceal the true size of the movement from the agents. Specifically, if the social movement has expanded to a point where the effectiveness of punishment would be significantly diminished—i.e., diluted to the extent that it no longer serves as a deterrent—the Principal is motivated to delay repressive measures if she expects that future agents will behave sufficiently well. Doing otherwise would then allow all future agents to update their strategy in favor of joining the movement, since they know that this repressive action is terminal and hence no longer dissuasive thereafter. However, this Principal's strategy of delaying the repressive action can then enable the movement to grow arbitrarily large. Indeed, sooner or later agents will guess that the dissuasion threshold has been reached and will thus decide to join the movement with high probability, thereby forcing the Principal to take the repressive action as she does not expect the agents to behave well enough in the future. This action being terminal, it allows the movement to grow arbitrarily large thereafter. Thus, by trying to "buy time", the government's punishment becomes less credible and it can thus ultimately fail to contain the movement.

This is formalized in the next proposition.

Proposition 3 (Imperfect monitoring) *Consider the informational environment characterized by the imperfect monitoring, by the agents, of the protest movement's size. For any $\eta > 0$, there are $\overline{\delta}_A \in (0,1)$, $\overline{\delta}_P \in (0,1)$, $\overline{T} > 0$ and $\underline{R} > 0$ such that when $\delta_A > \overline{\delta}_A$, $\delta_P > \overline{\delta}_P$, $T > \overline{T}$ and $R(B) < \underline{R}$, there exist equilibrium strategies σ_A^* and σ_P^* such that, with probability greater than $1 - \eta$, at least $T - 1$ agents choose $a = 1$.*

An interesting implication of Proposition 3 is that, to limit the growth of the social protest movement, the Principal would have an interest in credibly communicating its size n_{t-1} to the agents at all times. In such a (perfect monitoring) case, the Principal could achieve the result of Proposition 1 (or Proposition 2) and contain the movement. This would allow it to convince the agents that the movement is still small enough (i.e. not exceeding the dissuasion threshold) so that the repressive action can still act as a deterrent.

However, this would have to be done through some communication technology that would be mediated outside the control of the Principal to ensure a kind of commitment that would not hold if the Principal were in charge of it. As this may not always be feasible, imperfect monitoring of the size of a movement is likely often present in reality and this suggests a reason why protest movements can succeed.

To better understand Proposition 3, we can consider the following equilibrium as an example. Let the first $m - 2$ agents choose action $a = 1$, while agents $m - 1$ and m play mixed strategies. The Principal plays a mixed strategy at time $m - 1$, since she expects agent m to behave well enough in the future (Lemma 3) and thus hesitates to take the repressive action. This uncertainty about whether the repressive action will be taken then justifies the hesitation of agent $m - 1$ (and his mixed strategy). Due to imperfect monitoring, agent m is then unsure of whether the size of the protest movement is $n_{m-1} = m - 2$, in which case the dominant strategy would be to choose action $a = 0$, or whether it is $n_{m-1} = m - 1$, in which case the dominant strategy would be to choose action $a = 1$ as the punishment is sufficiently diluted. This uncertainty in turn makes agent m hesitate to join the movement (and thus choose a mixed strategy). Note that in such an equilibrium, the Principal can actually take the repressive action on path—either if the $(m-1)$-th agent chooses $a = 1$, or if the m-th agent chooses $a = 1$—which then allows all subsequent agents to choose action $a = 1$ thereafter. Under the parametric conditions of Proposition 3, it can be shown that in such an equilibrium, all agents except agent $m - 1$ choose action $a = 1$ with arbitrarily high probability, and thus the Principal is effectively unable to contain the protest movement.

Real-world applications consistent with this result include cases in which a government implements repressive measures and yet fails to contain the protest movement, which subsequently continues to expand.[19]

3.4 Contrast with a game where agents act simultaneously

Our dynamic analysis yields additional important insights that cannot be captured in a static game where all agents act simultaneously and the principal can only act in one period. Obviously the incentives to delay punishment cannot be captured by such a model, nor can the impact of agents joining the movement dynamically on the credibility of a repressive governmental action.

To illustrate this, we will now examine the equivalent simultaneous-actions coordination game, in which agents must act at the same time, and the Principal gets the chance to take the punitive action immediately after.

Consider T agents, each of whom can choose an action $a \in \{0, 1\}$ at time 1. The Principal then observes the action profile $\vec{a}$ and, immediately after, chooses whether to take repressive action, i.e. her action is $a_P \in \{0, 1\}$, where $a_P = 1$ means that she takes repressive action, in which case a set of up to B agents is selected uniformly at random to be punished, whereas $a_P = 0$ means she does not take repressive action. Payoffs are realized at time 1.

Agent i has payoff

$$\pi_i(a_i, \vec{a}_{-i}, a_P) = v(a_i, \sum_{j \neq i} a_j) - \kappa \cdot a_i - C \cdot \phi_i \tag{5}$$

where $\phi_i = 1$ if agent i is in the set of agents punished by the Principal and $\phi_i = 0$ otherwise. All agents are homogeneous. Call $\sigma_A \in \Delta(\{0, 1\})$ an agent strategy.

The Principal has payoff

$$\pi_P(a_P, \vec{a}) = -\sum_j a_j + R(B) \cdot \psi, \tag{6}$$

where $\psi = \mathbb{1}_{\{a_P=1 \bigcap a_j=1 \ \forall j \ s.t. \ \phi_j=1\}}$. We call $\sigma_P : \{0, 1\}^T \to \Delta(\{0, 1\})$ the Principal's strategy, which is a mapping from an agents' actions profile $\vec{a}$ to a probability measure over taking the repressive action or not.

We assume $v(1, 0) - v(0, 0) < \kappa < v(1, T - 1) - v(0, T - 1)$ so that the coordination problem among agents is not trivial. In such a game, there could be multiple equilibria. Namely a zero-contribution equilibrium with $a_i = 0$ for all i, a full contribution equilibrium with $a_i = 1$ for all i as well as mixed strategy equilibria. Equilibrium selection here will depend on the size of the Principal's budget. Namely, when $B = 0$, the Principal is effectively absent and this corresponds to a standard coordination game among agents only. When $1 \leq B < T$, the best the Principal could do after observing agents taking action $a = 1$ would be to punish up to B randomly-selected such agents, each agent being selected with probability $\min(\frac{B}{\sum_j a_j}, 1)$. A sufficient condition to obtain a unique, zero-contribution equilibrium ($a_i = 0$ for all i) here is that the Principal's budget be large enough, since the expected marginal payoff of joining would be too low while the probability of being punished is too large. This is summarized in the following proposition.

[19] It is to be noted that many real-life examples, such as the Arab Spring protests of 2011 in Egypt, show that repressive action can be taken and yet be unsuccessful in stopping a protest movement's growth. In the model extension introduced in Section 4.1, the Principal can halt the growth of the movement with probability $p(B)$ upon choosing the repressive action. Under imperfect monitoring, protest activity may therefore continue to expand despite the use of repression along the equilibrium path, or alternatively be quelled with positive probability. While this extension better captures the range of outcomes observed in real-world protest dynamics, presenting the results within the baseline model (as in Proposition 3, where $p(B) = 0$) allows us to more clearly highlight the logic of deterrence and the additional incentives to delay repression that arise from imperfect monitoring.

Proposition 4 (Equilibria with simultaneous actions) *Consider the game where T agents act simultaneously.*

(I) In the absence of the Principal (or when $B = 0$), there exist multiple equilibria. These include, namely, a no-contribution equilibrium where agents choose $a_i^* = 0$ for all i, a full-contribution equilibrium where agents choose $a_i^* = 1$ for all i, as well as a symmetric mixed-strategy equilibrium.

(II) In the presence of the Principal (when $B \geq 1$), when $B/T < \frac{\Delta v(T-1)-\kappa}{C}$, then there is always an equilibrium in which agents choose $a_i^* = 1$ for all i. For all equilibria to require $a_i^* = 0$, for all i, we need that $B/T > \frac{\Delta v(T-1)-\kappa}{C}$.

Thus the size B of the Principal's budget (her ability to punish a large number of agents) is key to equilibrium selection in this model where agents act simultaneously. Since under our assumptions, $\frac{\Delta v(T-1)-\kappa}{C}$ is bounded away from 0, we conclude that the Principal would need a budget B that also grows very large as T gets large to be sure to deter any $a_i = 1$ in equilibrium. This is to be contrasted with our findings in the dynamic version of the game, where the budget B did not play such a role in deterring any $a_i = 1$ in equilibrium under the conditions of Proposition 1. Indeed, a higher B (holding C fixed) meant the Principal had higher incentives to delay repressive action and thus allowed the protest movement to grow larger before the threat of punishment became credible (Lemma 2).

4 Discussion and extensions

We now discuss some extensions.

4.1 On using government resources to dissolve the movement

The role of the repressive action in the model presented so far has been to contain a social protest movement through pure dissuasion, i.e. the fear of punishment, which could take the form of arrests or lawsuits directed at some of its participants.

As discussed in Section 3.2.2, the repressive action could also involve dismantling the very organizational structure of a social movement so as to prevent it from growing further. For instance, the repressive action could also focus on terminating the social protest movement by dissolving its leadership structure, seizing its assets and financial resources, etc. The model presented in the main part of the paper did not account for that, so as to focus only on the purely dissuasive aspect of a repressive action.

We can however easily augment our model to allow the budget to also be used to quell the movement. Let $p(B) \in [0, 1]$ be the probability that, upon taking the repressive action, the government succeeds in stopping the growth of the movement. In that event, no further agents will ever join the movement in the future, i.e. $a_t = 0$ for all $t > s$ if the Principal took the repressive action at time s. We let $p(B)$ be increasing in B, to model the idea that a Principal endowed with higher resources is more likely to be effective in quelling a protest movement.

In such a case, the exogenous (intrinsic) benefit $R(B) > 0$ of taking the repressive action is no longer needed, as the endogenous expected benefit of stopping the movement makes the punishment credible in the eyes of the agents when $m - 2$ agent have taken action $a = 1$. Indeed, even if $R(B)$ is negative (i.e. an intrinsic "cost" of taking the repressive action), the overall expected benefit of stopping the growth of the movement can still exceed that cost, and thus make the repressive action optimal and hence credible. This renders this more general model amenable to a wider range of real-world applications. Finally, if $R(B) > 0$ and when $P(B)$ is high enough, the Principal can become sufficiently confident that the repressive action will stop the movement that she finds it optimal to take repressive action *at the first* *occasion*, i.e. as soon as one agent has joined the movement. Anticipating this, no agent will then join the protest movement in any equilibrium.

We thus have the following proposition.

Proposition 5 (Equilibrium when the Principal can stop the growth of the protest movement) *Let $p : \mathbb{N} \to [0, 1]$, be increasing, one to one and onto. (i) If $p(B) > 0$, then there exists $\overline{R} < 0$ such that Propositions 1 and 2 also hold for $R(B) > \overline{R}$. (ii) If $R(B) > 0$, there exists $p_{th} \in (0, 1)$ such that when $p(B) \geq p_{th}$, then the Principal finds it optimal to choose $a_{P,t} = 1$ at the first t such that $a_t = 1$. Thus, in any equilibrium, all agents choose $a_i^* = 0$.*

An interesting implication of this setting is that the effect of the budget B on the number of agents who can join the protest movement can now potentially be non-monotonic and discontinuous. As B increases, the dissuasion threshold m increases (as per Lemma 2), which incentives the Principal to delay the repressive action, ceteris paribus. However, the incentives of the Principal also depend on the interplay with the functions $R(B)$ and $p(B)$. As $p(B)$—which is also increasing in B—keeps rising, the anticipated benefit of quelling the protest movement then becomes high enough so that the Principal chooses to quell it at the first opportunity, thus limiting its size to a maximum of 0.

This is formalized in the next corollary to Proposition 5.

Corollary 1 (Comparative statics in B) *For any $R(B) > 0$, there exists $\underline{B} \in \mathbb{N}_+$, such that the maximal size of the protest movement, $\overline{n}$, is non-negative over $B \in [1, \underline{B})$, and is 0 when $B \geq \underline{B}$.*

4.2 A punishment budget versus a punishment cost

In the extended model presented in Section 4.1, we saw that the Principal could even pay an exogenous cost $R(B) < 0$ when punishing and still have an incentive to take the repressive action, since the expected endogenous benefits of stopping the growth of the protest movement could exceed that cost.

In contrast, in the main model without the ability to stop the growth of the protest movement, we do not explicitly model a cost of punishing paid by the Principal, as we assumed that $R(B) > 0$. This exogenous benefit, coupled with a punishment budget B allocated to countering a social movement, enabled the Principal to deter collective action.

Note that, in the main model we could micro-found the budget B in different ways. For instance, suppose that $R(B) = rB$, where $r > 0$ is the intrinsic marginal benefit of punishing an agent. Suppose also that the Principal faces an increasing marginal cost $c(x)$ of punishing the x-th agent, such that $c(x) < r$ for $x \leq B$, while $c(x) > r$ for $x > B$. Then our analysis indicates that the principal would have a net intrinsic benefit $r - c(x) > 0$ for punishing up to B agents, while it would incur a net cost $c(x) - r > 0$ thereafter. Our previous results still hold in this scenario, as they only required that the intrinsic benefit be strictly positive for each of up to B punished agents.

Conversely, if punishing each agent always carried a net cost $c(x) - r > 0$ for all $x \geq 1$ (and hence $R(B) < 0$, in the main model), our analysis suggests that the Principal would always be inclined to delay punishment, and thus would be unable to counter the collective action effectively. An implication of this is that, for the purpose of deterring collective actions, the Principal having to pay a net cost or a "price" for each punished agent would prevent her from achieving deterrence, even if she were allowed to punish an arbitrary number of agents. On the other hand, a pre-allocated, fixed policing budget B associated with a small net punishment benefit, and which thus simply puts a cap on the number of agents that can be punished, would allow her to achieve deterrence. This follows because, as our analysis reveals, the latter but not the former allows the Principal to make punishment credible, in some circumstances, and hence achieve a form of commitment ability. In this sense, constraining the enforcer to follow a prescribed set of policing rules may be more effective for deterrence than granting her discretion over the number of agents to arrest.

4.3 On patience and the importance of dilution

In our model, we have assumed that the Principal as well as the agents were forward-looking and patient. We have also assumed that the Principal, when taking a repressive action, could not consider the order in which agents picked $a = 1$, which forced her to use a randomization device to punish agents who had joined the protest. In turn, this induced a dilution effect as the size of the protest was growing. Since, by Assumption 1, $\Delta v(T-1) - \kappa < C$ and thus the punishment, when suffered by an agent, was always greater than his marginal benefit of joining the movement, it is this dilution effect that in turn allowed us to obtain an inversion of the incentive to join the protest as its size was crossing a threshold that we referred to as the dissuasion threshold (Definition 1). This led to our partial deterrence result (Proposition 1).

While our assumption that agents are patient is natural in many applications, it is also interesting to analyze the case in which agents are impatient, and we will consider the limiting case in which they are very impatient (i.e., very low δ_A). Our main observation is that, when the punishment cost C is not too large, we can obtain an inversion of incentives to join the protest—which leads to a partial deterrence result—even if the Principal uses the order in which agents joined the protest to always target, say, the most recent ones for punishment. In such a case, the punishment is not diluted as the size of the protest grows larger.

To illustrate this, assume the Principal can punish only one agent (i.e. $B = 1$) and that she targets the last agent who picked $a = 1$. Also assume that the intrinsic cost of joining the movement is zero, i.e. $\kappa = 0$. Remember that $\Delta v(n)$ measures the (immediate) marginal benefit of joining the protest when n other agents are already in, which is essentially the only benefit the agent will care about if he is very impatient. We will still assume that $\Delta v(n)$ is increasing in n, but will no longer assume that $\Delta v(T-1) < C$. Instead, for a given punishment cost C, let $n^* > 2$ be defined so that $\Delta v(n^* - 1) > C > \Delta v(n^* - 2)$. In this case, the n^*-th agent to join the protest cannot be deterred given that the immediate benefit of joining exceeds the punishment cost C (i.e. $\Delta v(n^* - 1) > C$). This agent will thus choose $a = 1$ no matter what he expects about the Principal's punishment strategy. Now consider the $(n^* - 1)$-th agent. If he were to join the protest, it would be clear for the Principal that all subsequent agents will choose $a = 1$, as we have just shown. The Principal would then strictly prefer taking repressive action now (given that $\delta_P < 1$). This in turn implies that the $(n^* - 1)$-th agent would receive a net (immediate) payoff of $\Delta v(n^* - 2) - C$ by joining the protest and, given that $\Delta v(n^* - 2) - C < 0$, he prefers not to. This simple argument establishes that there cannot be more than $n^* - 2$ agents joining the protest in any Subgame Perfect Nash Equilibrium. That is, we have a partial deterrence result.

When agents are impatient (i.e. when δ_A is very low), it is the increasing immediate benefit of joining the protest when its size grows that drives the partial deterrence result and no dilution in the punishment is needed for this to arise in this case. In our model with patient agents (i.e. when δ_A is higher), the situation is different because they mostly care about the ultimate rather than the current size of the protest, i.e. they care more about $\Delta v(\infty)$ than about $\Delta v(n)$. Since the ultimate size is the same for all agents, no inversion of incentive can arise if the last agent who joined the protest is consistently punished. The dilution is then essential to drive our partial deterrence result when agents are patient, which we believe is the more natural benchmark.

4.4 On sequential punishment

While in the main part of the paper, we focused on a terminal repressive action, it is interesting to see what happens when the Principal can act on more than one occasion. To see this, consider a setting where the budget can be spread over time, instead of being used only once. That is, B is now the number of times the principal can act and punish an agent. Thus, the Principal initially holds a budget $B_0 = B \in \mathbb{N}^+$ and at each time t, she can choose the punishment option, in which case the budget is decreased by one increment. When the budget reaches 0, the Principal can no longer act.

Formally, at each such time, as long as the budget permits it, i.e., as long as $B_t \geq 1$, the Principal can choose an action $a_{P,t} \in \{0, 1\}$ where $a_{P,t} = 1$ means that the Principal punishes one of those agents who previously chose $a_s = 1$ (for $s \leq t$), and where $a_{P,t} = 0$ means that the Principal does not take punitive action. When $a_{P,t} = 1$, the budget is reduced by one unit, i.e. $B_{t+1} = B_t - 1$ and one of the agents who chose $a_s = 1$, for $s \leq t$, is picked uniformly at random to be punished.[20]

The only modification we need to make to the setting of Section 2.2 is in the Principal's flow payoff at time t, which is

$$\tilde{\pi}_{P,t}(\vec{a}_{P,t}, \vec{a}_t) = -n_t + R(B)\psi_{s,t}, \tag{7}$$

but where now $\psi_{s,t} = \sum_{j,s\leq t} \tilde{\psi}_{j,s}$, with $\tilde{\psi}_{j,s} = \mathbb{1}_{\{a_{j,s}=1 \bigcap \tilde{\phi}_{j,s}=1\}}$, indicating that the Principal reaps a permanent benefit $R(B)$ every time she punishes an agent who had previously joined the protest. Indeed, $\psi_{s,t}$ is now the running sum of the instances of punishment of agents who had joined the movement ($a_{j,s} = 1$) before they were punished ($\tilde{\phi}_{j,s} = 1$).[21] Note that here, $R(B)$ could plausibly be made a decreasing function of B, which aligns with the idea that the Principal's budget is split across the agents, although this does not affect our results.

We also modify the definition of the dissuasion threshold, to account for the fact that the Principal punishes *a single* protest member, uniformly at random, every time she chooses action $a_P = 1$.

Definition 2 (Dissuasion threshold under sequential punishment) *Let $T = \infty$. Given a punishment cost C felt by agents, we call the* dissuasion threshold *the integer $\tilde{m}$ such that $\frac{C}{\tilde{m}} < \Delta v(\infty) - \kappa < \frac{C}{\tilde{m}-1}$.*

We then have the following lemma.

Lemma 4 *The dissuasion threshold $\tilde{m}$ under sequential punishment is no greater than the dissuasion threshold m in the main model where the Principal can only take repressive action once, i.e. $\tilde{m} \leq m$.*

We finally have the following proposition.

Proposition 6 (Partial and full deterrence under sequential punishment) *Consider the sequential punishment setting where the Principal can take action $a_P = 1$ up to B_0 times. (i) Let $T = \infty$ and $\tilde{m}$ be the dissuasion threshold. Also let $B_0 \geq 1$. There exist $\overline{\delta}_A \in (0, 1)$ and $\overline{\delta}_P \in (0, 1)$ such that if $\delta_A > \overline{\delta}_A$ and $\delta_P > \overline{\delta}_P$, then in any equilibrium there are at most $max(0, \tilde{m} - 1 - B_0)$ agents choosing $a_i^* = 1$. (ii) If, additionally, $\tilde{m} < N + 1 + B_0$, then in any equilibrium all agents choose $a_i^* = 0$.*

[20] Note that this now implies that a particular agent can possibly be punished multiple times. Indeed, in this modified setting, $\phi_{i,t}$ remains as defined in Section 2.2—i.e. $\phi_{i,t} = \{\exists\tau \leq t : \tilde{\phi}_{i,\tau} = 1\}$, where $\tilde{\phi}_{i,\tau} = 1$ if agent i is punished at time τ and $\tilde{\phi}_{i,\tau} = 0$ otherwise—but is now the event that agent i has been punished, *possibly multiple times*, by time t. We could adjust the probability of being selected for punishment so that the Principal targets agents without replacement, but this would simply complicate the exposition without affecting our results.

[21] By contrast, recall that in Eq. (3), $\psi_{s,t}$ was instead expressed as $\psi_{s,t} = \mathbb{1}_{\{(a_{P,s,t}=1) \bigcap (a_{i,s}=1, \forall i \in \{i: \tilde{\phi}_{i,s}=1\})\}}$, which is the indicator that the Principal has taken the terminal repressive action at some time s (which was specified to be $s \leq t$) and targeted punishment at agents among those who had already joined the movement by time s. The Principal reaped a total permanent benefit of $R(B)$ when taking this terminal repressive action, whereas in Equation (7), the principal reaps such a benefit for every agent she punishes (hence the summation).

Proposition 6 suggests a potential role for the size of the Principal's budget B_0 in deterring agents from choosing $a = 1$. Indeed, we see that being able to act multiple times allows the Principal to increase its dissuasiveness for the agents preceding the dissuasion threshold $\tilde{m}$. Not only can the dissuasion threshold $\tilde{m}$ be smaller than when the Principal can only take repressive action once, i.e. $\tilde{m} \leq m$, but the incentive to delay punishment is reduced by how many additional times the Principal can take repressive action. She can thus reduce the number of agents joining the movement by even more than the number of times she can act.

But in some natural cases, the role of the budget turns out to be insignificant. Indeed, when the number of agents required for coordination to be worthwhile (i.e., the internal coordination threshold N) is large enough, we can again strengthen the result by establishing that, in any equilibrium, no agent chooses $a = 1$, irrespective of the size of the Principal's budget. Thus, under the conditions of Proposition 6(ii), even with a single "bullet" the Principal can deter *all* agents from coordinating on action $a = 1$. Fixing the punishment cost C, these conditions hold as long as the minimum number of agents needed to make coordination worthwhile is large enough (i.e. $N > \tilde{m} - 2$).

Note that the effect of the budget B is richer than in the main model with one-time repression. Indeed, in this sequential model where the budget can be spread over time, the dissuasion threshold $\tilde{m}$ is actually defined independently of B (Definition 2). A larger budget then becomes more effective in limiting the size of the movement. On the other hand, in the main model with one-time repression, a larger budget meant that punishment remained credible for longer, as the dissuasion threshold m was non-decreasing in B (Lemma 2) and thus the Principal had a greater incentive to delay repression. This allowed more agents to join the movement.

4.5 On the importance of reaction time in deterring collective action

Our framework also lends itself well to stochastic decision times, which provide further insights.

Consider the sequential punishment model presented in Section 4.4, but modified so that at each time t, agent t gets a chance to act with probability $q_A \in [0, 1]$ and the Principal gets a chance to act (immediately after the agent) with probability $q_P \in [0, 1]$. This can allow us to model the "speed" at which agents get the chance to join the protest movement (i.e. q_A) and, most importantly, the speed at which the Principal can *react* to the agents' actions (i.e. how large q_P is relative to q_A).

Then, if the Principal reacts slowly enough (i.e., if q_P is low enough relative to q_A), agents succeed in dynamically coordinating and the protest movement grows arbitrarily large, since agents feel that the probability with which they may be punished is low enough. This is due both to the fact that the Principal is unlikely to act quickly after an agent has joined the movement (and hence the expected punishment is further discounted) and because this then further dilutes the expected punishment among more agents who can join the movement between the time an agent takes action $a = 1$ and the time the Principal gets to react. On the other hand, if the Principal reacts sufficiently quickly (i.e., if q_P is large enough relative to q_A), we recover the previous result capping the size of the protest movement. This is summarized in the following proposition, which is a generalization of Proposition 6.

Proposition 7 (Reaction time and deterrence) *Consider the sequential punishment setting where the Principal can take action $a_P = 1$ up to B_0 times, and suppose the agent and the Principal each get a chance to act with probabilities q_A and q_P, respectively, at each time t. (i) For any $T > N$, there exist $\overline{\delta}_A \in (0, 1)$, $\overline{q}_A \in (0, 1)$ and $\underline{q}_P(q_A) \in (0, 1)$, such that when $\delta_A > \overline{\delta}_A$, $q_A > \overline{q}_A$ and $q_P < \underline{q}_P(q_A)$, then in any equilibrium all agents choose $a_i^* = 1$.*
Now let $T = \infty$, $\tilde{m}$ be the dissuasion threshold, and $B_0 \geq 1$. Then: (ii) For any $q_A \in (0, 1]$, there exist $\overline{\delta}_A \in (0, 1)$, $\overline{\delta}_P \in (0, 1)$ and $\overline{q}_P(q_A) \in (0, 1)$, such that when $\delta_A > \overline{\delta}_A$, $\delta_P > \overline{\delta}_P$ and $q_P > \overline{q}_P(q_A)$, then

in any equilibrium there are at most $max(0, \tilde{m} - 1 - B_0)$ *agents choosing* $a_i^* = 1$*; (iii) If, additionally,* $\tilde{m} < N + 1 + B_0$*, then in any equilibrium all agents choose* $a_i^* = 0$.

Note that in Proposition 6, and throughout the article, it is assumed that $q_A = 1$ and $q_P = 1$, and thus agent t always gets a chance to act at any time t, and the Principal can act immediately after. Proposition 6 is indeed a special case of Proposition 7.

This result supports the view that in plausible scenarios, the Principal may be more effective in deterring adversarial collective actions by reacting quickly rather than by being able to take retaliatory measures against a large number of agents.[22] It shows that having a large budget B_0 while reacting too slowly cannot allow the Principal to deter agents.[23]

4.6 On the effect of personalized punishments

4.6.1 Modified setting

In the setting with sequential actions and random opportunities to act of Section 4.5—as well as in the rest of the paper—we have assumed that when the Principal triggers a punishment, she anonymously targets those agents who have previously chosen $a = 1$. This can be rationalized on the ground that at each time the Principal is called to play, she is only informed of the set of agents who previously chose $a = 1$ and nothing else, or is not allowed to use agents' labels in her punishment strategy.[24]

As an alternative scenario, consider now the case in which the Principal would perfectly observe the identities of the agents and would be allowed to target the punishment at any particular agent. This may fit better with applications to politics, where high profile public figures are well known and can be individually targeted. The Principal, who here could represent a lobby or an NGO, could be scrutinizing the public statements made by politicians, and could also be capable of targeting a particular politician for punishment. The punishment, in this case, could take the form of a lawsuit or the issuance of critical statements directed at that particular politician.

Likewise, an application to online social media activism lends itself well to this setting. Indeed, the agents (i.e. civil society activists) make statements on social media at different times, hoping that other agents will support the same view in the future so that a movement can be sustained. Their posts being time stamped, the Principal could then use information about the timing of their actions to target them specifically.

Calling $\mathcal{H}_P$ the set of possible Principal's histories, the Principal's strategy is then $\sigma_P : \mathcal{H}_P \rightarrow \Delta(\{\emptyset, \mathbb{N}_+\})$. That is, at some action time τ, based on a Principal's history of past play $h_{P,\tau}$, the Principal can choose to punish any particular agent $i \in \mathbb{N}_+$. She can also decide not to take punitive action (i.e. not to choose any agent to punish, $\emptyset$). She can also randomize.

Here $\tilde{\phi}_{j,\tau} = 1$ if $a_{P,\tau} = j$, i.e. agent j is punished at time τ, and $\tilde{\phi}_{j,\tau} = 0$ if $a_{P,\tau} \neq j$. The payoffs are otherwise defined as in the setting of Section 4.5.

[22] In an application to crime and policing, it is interesting to note a parallel with the famous "broken window theory" (e.g. Corman and Mocan (2005)), in which the police (the Principal, in this case) wants to react quickly even when a minor crime is committed, as this signals to the criminals (the agents, in this case) that she has a high q_P relative to q_A. In this famous theory, the actions of the police must also be visible, which is the case in our model as $a_{P,t}$ is observed by all agents.

[23] It may be interesting to analyze how the maximum size of the protest movement can vary with the Principal's budget B_0 and with her reaction speed (which is proportional to q_P). While a full analysis is non-trivial and left for future research, our results (obtained when q_P is low or high) suggests a form of substitutability in the sense that a high q_P may substitute for a high B_0 (at least in some range of q_P) for the purpose of reducing the protest movement's size. Indeed, Proposition 7 tells us that the reaction speed limits the dilution of punishment, as fewer agents can join the movement before the Principal gets a chance to act, and thus helps limit the size of the movement. Likewise it tells us explicitly that a larger budget also helps limit the number of agents who join the movement.

[24] We discussed real-life settings in which this is likely to occur in Section 2.2.3. Formally, this implies a form of forgetfulness, as the Principal is not assumed to keep track of these sets for each of the previous times in which she was called to play. Alternatively, assuming all information about past play is available to the Principal, it can be viewed as formalizing a kind of Markovian-like refinement of SPNE based on symmetry and the payoff-relevance criterion.

4.6.2 Main insights

In this scenario, we note the following: Irrespective of whether the Principal reacts quickly or not to the agents' actions, there is always an equilibrium in which all agents are deterred from choosing $a = 1$. To see this, assume that $B_0 = 1$, $T = \infty$, and consider a tentative equilibrium in which the Principal punishes the *first* agent who chose $a = 1$ (if there is one) when she has an opportunity to move. In such a situation, no agent would be willing to be the first to choose $a = 1$, and as a result all agents would choose $a = 0$. It is readily verified that the Principal's strategy is part of an equilibrium, since if a first agent were to choose $a = 1$, all subsequent agents would also choose $a = 1$ (since the Principal would exhaust her budget by punishing that first agent). Punishing the first agent (instead of any other agent who chose $a = 1$) when she gets a chance to act would then be (weakly) optimal for the Principal.

While there is a possibility that no agent chooses $a = 1$ in equilibrium, we now note that when punishments can be personalized, there are also other equilibria in which this is not the case. To illustrate this most simply, assume that absent the punishment, $a = 1$ is a dominant strategy for the agents (i.e., $N = 1$ in Assumption 1). Then punishing the l-th agent who chooses $a = 1$, with the first $l - 1$ agents choosing $a = 1$ and all subsequent agents choosing $a = 0$, would also be an equilibrium.

One may also note that the Principal would ensure the best outcome for herself if she could *commit* to always punishing the first agent who chooses $a = 1$. The only requirement here is that either the Principal acts sufficiently often (i.e. $q_P > \overline{q}_P^c$, for some $\overline{q}_P^c \in (0,1)$), or that the agents do not discount the future too much (i.e. $\delta_A > \overline{\delta}_A^c$ for some $\overline{\delta}_A^c \in (0,1)$). This then ensures that an agent will fear punishment even if it comes much later than the time at which he acted. But the *relative* reaction speed of the Principal to the agents (as measured by the requirement that $q_P > \overline{q}_P(q_A)$ in Proposition 7) is now irrelevant. This is in sharp contrast with the analysis in Section 4.5 (and in the main model, where q_P was assumed to be 1, and thus the reaction speed was maximal).

This is formalized in the following definition and proposition.

Definition 3 (First offender punishment strategy) *σ_P is called a* first offender punishment strategy *if the Principal punishes the first agent who took action $a = 1$, as soon as the Principal has an opportunity to react. That is, let h_{P,τ_k} be a Principal's history where i is such that $a_i = 1$ and $a_j = 0$ for all $j < i$, and $\tau_{k-1} < i < \tau_k$, where $\tau_{k-1}, \tau_k \in \mathbb{N}_+$ are two consecutive times where the Principal has the opportunity to act. Then $\sigma_P(i|h_{P,\tau_k}) = 1$, where $\sigma_P(i|h_{P,\tau_k})$ denotes the probability of selecting agent i for punishment after history h_{P,τ_k}.*

Proposition 8 (Equilibrium with commitment and personalized punishment) *Consider the sequential punishment setting where the Principal can take action $a_P = 1$ up to B_0 times, and suppose the agent and the Principal each get a chance to act with probabilities q_A and q_P, respectively, at each time t.*[25] *Let $T = \infty$ and let the Principal commit to a first offender punishment strategy σ_P^*. There exist $\overline{\delta}_A \in (0,1)$ and $\overline{q}_P^c(\delta_A) \in (0,1)$ such that for* any *budget size $B_0 \geq 1$, if $\delta_A > \overline{\delta}_A$ and $q_P > \overline{q}_P^c(\delta_A)$, then the equilibrium involves agents choosing $a_i^* = 0$ for all i. Moreover, $\overline{q}_P^c$ does not depend on q_A.*[26]

In Proposition 8, q_P must only be high enough so that an agent has a large enough chance of being punished *in the not-too-distant future.* $\overline{q}_P^c(\delta_A)$ is thus completely independent of the agents' activity rate

[25]Note that the result of this Proposition also follows in the single repression setting presented in the main part of the paper. We chose to present this result in the model where the Principal can punish at multiple different times because it seems to fit better applications where the latter could target agents individually.

[26]Note that here $\delta_A > \overline{\delta}_A$ is just a standard patience requirement to ensure that agents find it optimal to choose $a = 1$ in the absence of a Principal, for the given $q_A > 0$. Also note that we could also state Proposition 8 with a more stringent requirement on δ_A rather than q_P. Indeed, given any $q_P > 0$, if agents are patient enough, they will fear a punishment even if it comes very far into the future. Thus, there exists $\overline{\delta}_A^c \in (0,1)$ such that the same equilibrium outcome holds when $\delta_A > \overline{\delta}_A^c$, for any $q_P > 0$.

q_A and the principal does *not* need to react quickly to the agents' actions. However, it still depends on how much the agents value the future and thus on their discount rate δ_A.

The threat of a single bullet (i.e. any $B_0 \geq 1$) can discipline an entire population and the Principal does not even need to react quickly anymore.[27]

4.7 Presence of fearless agents

Consider again the setting of Section 4.5, with sequential actions and random opportunities to act, and assume for simplicity that $T = \infty$.

Suppose there are two types of agents: Rational and fearless, i.e. $\theta_i \in \{R, F\}$, and let the probability that an agent is fearless be $y \in (0, 1)$. The fearless agents have flow payoff

$$\tilde{\pi}^F_{i,t}(a_{i,t}, \vec{a}_{-i,t}, \vec{a}_{P,t}) = v(a_{i,t}, n_{-i,t}) - \kappa \cdot a_{i,t} \tag{8}$$

and thus do not fear punishment, just as if there were no Principal. The rational players have the payoff function as in Eq. (1), as before.

If the intensity q_A of the agents' activity is sufficiently high and if there is a sufficiently high fraction $y > \overline{y}(q_A)$ of all agents who are fearless (with $\overline{y}(q_A)$ decreasing in q_A), then the fearless agents will find it worthwhile to choose action $a = 1$, just like in a standard dynamic coordination game without a Principal. Indeed, they can always expect sufficiently many other fearless agents to choose action $a = 1$ after them, thus selecting (by subgame perfection) an equilibrium in which fearless agents always choose action $a = 1$.

Suppose an agent's type θ_i is private and thus not publicly observable. Then, if the Principal can announce and commit to a personalized punishment strategy, she can choose a first offender punishment strategy as in Section 4.6 (Definition 3). She will then have to punish the agent who chose $a = 1$ first. On the equilibrium path, this will surely be a fearless agent, but if the Principal did not punish him, this would incentivize other rational agents to choose $a = 1$ in the future (trying to be considered as fearless agents and thereby avoiding punishment). Thus the Principal will punish the fearless agents and deplete her budget. The second agent choosing $a = 1$ will also be a fearless agent, but the Principal will also have to punish him.[28] This will go on until her budget is completely depleted, at which point all agents will start choosing $a = 1$ since the Principal is no longer effectively active and the game becomes a standard dynamic coordination game. Thus, the presence of fearless agents can ultimately allow later rational agents to coordinate on action $a = 1$. The only way for the Principal to deter rational agents from choosing action $a = 1$ in this setting would be to have an infinite budget B_0, irrespective of her reaction speed. In a variant of this model, we could also suppose that there is only a finite number $N_F \in \mathbb{N}$ of fearless agents. In this case, the Principal would need to have a budget larger than the number of fearless agents (i.e. $B_0 > N_F$) in order to deter the rational agents from coordinating on action $a = 1$, once again illustrating the importance of the budget size when the agents' types (i.e. rational or fearless) are undetectable.

If θ_i is publicly observable, then a strategy by which only rational types can be punished can allow the Principal to preserve her budget and keep as much control over the agents as she can. Under such a strategy, the fearless agents are allowed to take action $a = 1$, but not the rational agents, and this would

[27] An interesting variation of this setting would be one where the Principal has limited commitment power, in the sense that she can commit only over a certain time window (say, a few periods ahead). This may however involve several non-trivial comparisons, such as how the commitment window compares to the timescale over which the Principal and the agents make their decisions, which itself depends on their respective "speeds" (i.e. q_A and q_P). We leave this analysis for future research.

[28] Indeed, going forward, this second agent becomes the next 'first agent' to choose action $a = 1$. A slight variation on Definition 3 is an *ordered punishment strategy*, by which agents are punished *in the order* in which they took action $a = 1$. Here the Principal could commit to such a strategy.

imply the coexistence of fearless agents choosing action $a = 1$ with rational agents choosing $a = 0$. Such a strategy would be implementable with any budget $B_0 > 0$.

The same insights apply when the Principal cannot use a personalized punishment strategy. Suppose that the conditions of Proposition 6 are satisfied, and that the Principal can condition her strategy on agent type (fearless or rational), but not on an agent's label i. Then if θ_i is private, she would choose a strategy that punishes all offenders, irrespective of their types, and inevitably exhaust her budget at some point. If θ_i is publicly observable, and if she can react quickly enough to the agents' actions (i.e. if q_P is sufficiently large relative to q_A), then she could choose a strategy by which she punishes the first $B_0 - 1$ agents, since she benefits from punishing offenders (recall that $R(B)$ is positive in her payoff function (cf. Eq. (7))). Perpetually keeping a budget of size $B_0 = 1$ thereafter would then allow her (at least ultimately) to credibly dissuade the rational agents from taking action $a = 1$—as she enjoys punishing offenders—and minimize the expected remaining number of offenders—which is precisely achieved by dissuading the rational types. Under such a strategy, all fearless agents thus take action $a = 1$ and rational agents (at least ultimately) take action $a = 0$, allowing for the coexistence of both offenders and non-offenders as before.

It is interesting to note that in this setting with fearless agents, the Principal would therefore benefit more from an improvement in the detection technology, which allows her to differentiate fearless from rational agents, than from an increase in her budget B_0. This again illustrates the greater importance of factors such as information and reaction speed in allowing the Principal to deter collective actions on the part of the agents.

5 Proofs

The following notation will be useful in the subsequent proofs.

Call $\Delta\pi_{t,t}(\vec{a}_{-t}, \vec{a}_P) = \pi_{t,t}(1, \vec{a}_{-t}, \vec{a}_P) - \pi_{t,t}(0, \vec{a}_{-t}, \vec{a}_P)$ and $\Delta v(\sum_{j\neq t} a_{j,t}) = v(1, \sum_{j\neq t} a_{j,t}) - v(0, \sum_{j\neq t} a_{j,t})$.

At time t, given some history $h_{A,t}$, a Principal's strategy σ_P and a strategy σ_A for the agents, agent t's expected marginal payoff from choosing $a_t = 1$ as opposed to $a_t = 0$ can be written as

$$\begin{aligned}\mathbb{E}[\Delta\pi_{t,t}(\vec{a}_{-t}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,t}] &= \sum_{s=t}^{T} \delta_A^{s-t}\Big(\mathbb{E}[\Delta v(\sum_{j\neq t} a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{A,t}] - \mathbb{E}[C\phi_{t,s}|a_t = 1, \sigma_P, \sigma_A, h_{A,t}]\Big) \quad (9)\\ &= \sum_{s=t}^{T} \delta_A^{s-t}\Big(\mathbb{E}[\Delta v(\sum_{j\neq t} a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{A,t}] - C\mathbb{P}\{\phi_{t,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,t}\}\Big)\end{aligned}$$

where the expectation on the right-hand side is taken over the a_j's.

Proof of Lemma 1 (Successful agent coordination).

Note that without a Principal the last term in Eq. (9), that is $\sum_{s=t}^{T} \delta_A^{s-t} C\mathbb{P}\{\phi_{t,s} = 1|a_t = 1, \sigma_P, \sigma_A, h_{A,t}\}$, is zero for any t.

Now recall from Assumption 1 that N is the number of agents who must choose $a = 1$ in order to make it worthwhile (in the absence of a Principal) for some agent i to choose $a = 1$. Thus, agent N will have positive expected marginal payoff of choosing $a = 1$ when the $N - 1$ previous agents have also chosen action $a = 1$, since $\Delta v(N - 1) - \kappa > 0$:

$$\mathbb{E}[\Delta\pi_{N,N}(\vec{a}_{-N}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,N}] = \sum_{s=N}^{T} \delta_A^{s-N}\Big(\mathbb{E}[\Delta v(\sum_{j\neq N} a_{j,s})|\sigma_P, \sigma_A, h_{A,N}] - \kappa\Big) > 0,$$

where $h_{A,N}$ is a history in which the $N - 1$ previous agents have also chosen action $a = 1$.

Next, consider agent 1's decision at $t = 1$. Note that if δ_A is high enough (the agent is sufficiently patient), then agent 1 will have positive expected marginal benefit of choosing action $a = 1$, since then he precipitates a subgame in which all agents will choose $a = 1$. Indeed, let $\overline{\delta}_A(N,T)$ be such that, given a fixed profile of actions $\vec{a}_{-1} = \vec{1}$ for the other agents, then

$$\sum_{s=1}^{T} \delta_A^{s-1} \Big(\mathbb{E}[\Delta v(\sum_{j \neq 1} a_{j,s}) | \sigma_P, \sigma_A, h_{A,1}] - \kappa \Big) > 0, \quad \forall \delta_A > \overline{\delta}_A(N,T).$$

Specifically, if $N = 2$, then by choosing $a_1 = 1$, agent 1 precipitates a subgame in which it becomes strictly dominant for agent 2 (and all subsequent agents) to choose $a_2 = 1$. Since agent 1 is sufficiently patient ($\delta_A > \overline{\delta}_A(N,T)$), he will never choose $a_1 = 0$, and thus the profile $a_i = 1$ for all i is part of any subgame perfect Nash equilibrium.

Likewise, if $N = 3$ and $\delta_A > \overline{\delta}_A(N,T)$, then by choosing $a_1 = 1$, agent 1 precipitates a subgame in which when agent 2 chooses $a_2 = 1$, then it becomes strictly dominant for agent 3 (and all subsequent agents) to choose $a_3 = 1$. Thus, in such a case, agent 2 will choose $a_2 = 1$ (since he is sufficiently patient) and it follows that agent 1 will never choose $a_1 = 0$ (since he is also sufficiently patient to wait for this to unravel). Therefore $a_i = 1$ for all i is part of any subgame perfect Nash equilibrium.

Thus, by induction, we have that for any N, when $\delta_A > \overline{\delta}_A(N,T)$, then $a_i = 1$ for all i is part of any subgame perfect Nash equilibrium. It is trivial to show that $\overline{\delta}_A(N,T)$ is increasing in N, as the more patient agent 1 must then be to enjoy this unraveling of subsequent agents choosing action $a = 1$. This concludes the proof.

Note that the backward induction argument used in this proof is similar to the one used in Theorem 1 of Gale (1995), although in the latter setting, agents can all act at any given time. In contrast, Gale (1995) shows that as the length between any two decision times becomes small enough, agents succeed in coordinating very rapidly. ■

Proof of Proposition 1 (Partial deterrence).

Suppose δ_A is large enough. This will ensure that agents value sufficiently the future actions of other agents. Suppose then that, at time $t = i'$, $m - 2$ agents have already chosen $a = 1$ and the Principal has not yet taken action $a_P = 1$. If an additional agent (i.e. agent i') chooses $a_{i'} = 1$, then the Principal will choose $a_{P,i'} = 1$. This is so because otherwise, if the Principal does not punish, all subsequent agents will choose $a = 1$ as $\frac{BC}{m} < \Delta v(\infty) - \kappa$ (with the correct anticipation that all other agents will choose $a = 1$).

To see this, suppose that the Principal has chosen not to punish, i.e. $a_{P,i'} = 0$. Consider then the next agent i who gets to act at a time $i = i' + 1$. Using Eq. (9), and when δ_A is large enough (and when T is sufficiently larger than i), this agent i's expected marginal payoff from choosing action $a_i = 1$ can be expressed as

$$\begin{aligned} \mathbb{E}[\Delta\pi_{i,i}(\vec{a}_{-i}, \vec{a}_P) | \sigma_P, \sigma_A, h_{A,i}] &= \sum_{s=i}^{T} \delta_A^{s-i} \Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa | \sigma_P, \sigma_A, h_{A,i}] - C\mathbb{P}\{\phi_{i,s} = 1 | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \Big) \\ &\geq \sum_{s=i}^{T} \delta_A^{s-i} \Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa | \sigma_P, \sigma_A, h_{A,i}] - \frac{BC}{m} \Big) \\ &> 0, \end{aligned}$$

where the first inequality follows from the fact that $\mathbb{P}\{\phi_{i,s} | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \leq \frac{B}{m}$. Indeed, if the principal later chooses to punish (i.e. take action $a_{P,\tau} = 1$) at some time $\tau \geq i$, agent i will be punished (i.e. $\tilde{\phi}_{i,\tau} = 1$) with probability $\frac{B}{m}$ if no further agent has yet taken action $a = 1$ after agent i, or with probability less than $\frac{B}{m}$ if further agents have taken action $a = 1$. It follows that $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\}$ cannot be greater than $\frac{B}{m}$.

The second inequality is established by the following argument: Since $\Delta v(\infty) - \kappa > \frac{BC}{m}$, then when sufficiently many agents choose $a = 1$ (or are expected to do so in the future), the expected net benefits exceed the expected punishment cost $\frac{BC}{m}$. Indeed, under the correct expectation that all future agents choose $a = 1$, then for all $\xi > 0$, there exists s' (with $T > s' > i$, when T is large enough) such that when $s > s'$, then $\big(\Delta v(T) - \kappa\big) - \mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{A,i}] < \xi$. It thus follows that there exist $\overline{T}(i) > 0$ and $\overline{\delta}_A > 0$, such that $\sum_{s=i}^{T} \delta_A^{s-i}\Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{A,i}] - \frac{BC}{m}\Big) > 0$ when $T > \overline{T}(i)$ and $\delta_A > \overline{\delta}_A$, under the correct expectation that all future agents will choose $a = 1$. That T can be larger than $\overline{T}(i)$ follows trivially with our assumption that $T = \infty$.

Thus, since $\mathbb{E}[\Delta\pi_{i,i}(\vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,i}] > 0$, the Principal prefers choosing $a_{P,i'} = 1$, anticipating that all subsequent agents will choose $a = 1$ anyway. She prefers to take repressive action because she obtains a positive benefit from choosing $a_P = 1$ when punishing those who chose $a = 1$ (i.e. $R(B) > 0$ in Eq. (3)). She prefers doing it at the earliest opportunity (i.e. at time i') because of discounting (i.e. $\delta_P \in (0, 1)$ in Eq. (4)).

Now given this, after $m - 2$ agents have chosen $a = 1$, the subsequent agent will not be willing to choose $a = 1$ as this will trigger a repressive action and $\Delta v(\infty) - \kappa < \frac{BC}{m-1}$ (and thus he fears the random punishment). Thus, there cannot be more than $m - 2$ agents choosing $a = 1$.

To see this, consider again the next agent i who gets to act at a time $t = i$ such that $i = i' + 1$. Using Eq. (9), this agent i's expected marginal payoff from choosing action $a_i = 1$ can be expressed as

$$\begin{aligned}
\mathbb{E}[\Delta\pi_{i,i}(\vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,i}] &= \sum_{s=i}^{T} \delta_A^{s-i}\Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{A,i}] - C\mathbb{P}\{\phi_{i,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\}\Big) \\
&\leq \sum_{s=i}^{T} \delta_A^{s-i}\Big(\Delta v(\infty) - \kappa - C\mathbb{P}\{\phi_{i,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\}\Big) \\
&< 0,
\end{aligned}$$

where the first inequality follows from $\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s})|\sigma_P, \sigma_A, h_{A,i}] \leq \Delta v(\infty)$. To establish the second inequality, first note that, $\mathbb{P}\{\phi_{i,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} = \frac{B}{m-1}$. Since, $\Delta v(\infty) - \kappa - \frac{BC}{m-1} < 0$, this in turn implies that $\Delta v(\infty) - \kappa - C\mathbb{P}\{\phi_{i,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} < 0$, thus establishing the second inequality above.

Therefore, the $(m - 1)$-th agent (as well as all subsequent ones) does not choose $a = 1$.

Now we show that, when the Principal is patient enough, she will not choose $a_P = 1$ at any time $t < i'$. Indeed, in such a case, the benefits of taking repressive action are outweighed by the following growth of the protest movement. To see this, suppose a history $h_{P,t}$ where no more than $m - 3$ agents have chosen $a = 1$, i.e. the size of the movement is $n_t \leq m - 3$. Then the Principal's expected benefit of deviating and choosing $a_{P,t} = 1$, as opposed to following a strategy σ_P where she does not choose $a_P = 1$ until $m - 1$ agents have joined the protest movement is

$$\begin{aligned}
\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P, \vec{a})|\sigma_P, \sigma_A, h_{P,t}] &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 1, \sigma_P, \sigma_A, h_{P,t}] - \mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]\Big) \\
&= \sum_{s=t}^{T} \delta_P^{s-t}\Big(R(B) + \mathbb{E}[-n_s|a_{P,t} = 1, \sigma_P, \sigma_A, h_{P,t}] - \mathbb{E}[-n_s|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]\Big) \\
&\leq \sum_{s=t}^{T} \delta_P^{s-t}\Big(R(B) - (n_t + s - t) + (m - 2)\Big),
\end{aligned}$$

where the inequality follows from the fact that $\mathbb{E}[n_s|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}] \leq (m-2)$ when following strategy σ_P. For any $R(B) > 0$, if T is sufficiently larger than t, then there exists s, with $t < s < T$, such that $R(B) - (n_t + s - t) + (m-2) < 0$. Thus, there exists $\overline{\delta}_P \in (0,1)$ such that when $\delta_P > \overline{\delta}_P$ and when T is sufficiently larger than t (which follows trivially from our assumption that $T = \infty$), then $\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P,\vec{a})|\sigma_P,\sigma_A,h_{P,t}] < 0$.

Hence, the Principal will not take repressive action when she is still able to dissuade agents from joining the movement.

From the above, it follows that in any equilibrium, no more than $m-2$ agents choose action $a_i^* = 1$, and the Principal does not take repressive action on the equilibrium path. ∎

Proof of Lemma 2 (Comparative statics in B and C). $\frac{BC}{m}$ is decreasing in m, while $\Delta v(\infty) - \kappa$ is finite and constant, with $\Delta v(\infty) - \kappa < BC$ by Assumption 1. It follows that there can only exist a unique integer $m \geq 2$ such that $\frac{BC}{m} < \Delta v(\infty) - \kappa < \frac{BC}{m-1}$. Moreover, as $\frac{BC}{m}$ is increasing in BC, it follows that this m is non-decreasing in BC. ∎

Proof of Proposition 2 (Threshold criterion for full deterrence).

Proposition 2 is a corollary of Proposition 1. By Proposition 1, we know that there can be at most $m-2$ agents choosing $a = 1$. Given that $m - 2 < N$, no agent can find choosing $a = 1$ profitable. Indeed, by Assumption 1, fewer than $N-1$ other agents choosing $a = 1$ does not make the collective action worthwhile (i.e. $\Delta v(n) < \kappa$ for $n < N-1$). Here, this is ensured by the condition $m - 2 < N$. ∎

Proof of Lemma 3 (Delayed punishment when one future agent behaves well enough).

Suppose a history $h_{P,t}$ where $a_t = 1$. The Principal's expected payoff gain from choosing $a_{P,t} = 1$ instead of $a_{P,t} = 0$ is

$$\begin{aligned}
\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P,\vec{a})|\sigma_P,\sigma_A,h_{P,t}] &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t}=1,\sigma_P,\sigma_A,h_{P,t}] - \mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}]\Big) \\
&= \sum_{s=t}^{T} \delta_P^{s-t}\Big(R(B)(1-\mathcal{A}_s) - (n_t + s - t) + \mathbb{E}[n_s|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}]\Big) \\
&= \sum_{s=t}^{T} \delta_P^{s-t}\Big(R(B)(1-\mathcal{A}_s) - (s-t) + \mathbb{E}[\sum_{j>t} a_{j,s}|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}]\Big)
\end{aligned}$$

where, in the first equality, $\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t},\sigma_P,\sigma_A,h_{P,t}] = \mathbb{E}[-n_s + R(B)\sum_j a_{P,j,s}|a_{P,t},\sigma_P,\sigma_A,h_{P,t}]$, and where we later denote by $\mathcal{A}_s = \mathbb{E}[\sum_j a_{P,j,s}|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}]$ the probability that the repressive action has been taken by time s conditional on $a_{P,t} = 0$.

Note that $(1-\mathcal{A}_s) \in [0,1]$, since $\sum_j a_{P,j,s} \in \{0,1\}$ for any s (as the Principal can use action $a_{P,j}$ at only one time j). It follows that the term $R(B)(1-\mathcal{A}_s)$ is bounded, and converges to 0 as $R(B) \to 0$.

On the other hand, for $s \geq t' > t$, the term

$$-(s-t) + \mathbb{E}[\sum_{j>t} a_{j,s}|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}] < 0$$

when there is some t' such that $\mu_{t'} = \mathbb{E}[a_{t',s}|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}] < 1$. That is, when agent t' is expected to be sufficiently well behaved.

Combining these two observations, it follows that there exist $\underline{R} > 0$ and $\underline{\mu}(\delta_P, R(B)) \in (0,1)$, such that

$$R(B)(1-\mathcal{A}_s) - (s-t) + \mathbb{E}[\sum_{j>t} a_{j,s}|a_{P,t}=0,\sigma_P,\sigma_A,h_{P,t}] < 0$$

for all s large enough when $R(B) < \underline{R}$ and $\mu_{t'} < \underline{\mu}(\delta_P, R(B))$. It also follows that there exist $\overline{T} > 0$ and $\overline{\delta}_P \in (0,1)$ such that

$$\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P, \vec{a})|\sigma_P, \sigma_A, h_{P,t}] < 0$$

when $T > \overline{T}$ and $\delta_P > \overline{\delta}_P$. Thus, a sufficiently patient Principal chooses not to take repressive action as long as a future agent is expected to be sufficiently well behaved and the benefit $R(B)$ is small enough.

Note that $\underline{\mu}(\delta_P, R(B))$ is decreasing in $R(B)$, with $\underline{\mu}(\delta_P, R(B)) \to 1$ as $R(B) \to 0$, and $\underline{\mu}(\delta_P, R(B))$ is increasing in δ_P. Thus, the Principal has a more demanding requirement for expected good agent behavior when she benefits more from punishing, and she has a less demanding requirement for expected good agent behavior when she is more patient. ■

Proof of Proposition 3 (Imperfect monitoring).

We can construct an equilibrium with the following properties.

Agents play a strategy σ_A such that:

- On any agent history $h_{A,t}$ with $t \leq m-2$, then $\sigma_A(h_{A,t}) = 1$.
- On any agent history $h_{A,t}$ with $t = m-1$ and where the principal has not yet chosen $a_P = 1$, then $\sigma_A(h_{A,t}) \in (0,1)$.
- On any agent history $h_{A,t}$ with $t = m$ and where the principal has not yet chosen $a_P = 1$, then $\sigma_A(h_{A,t}) \in (0,1)$.
- On any agent history $h_{A,t}$ with $t \geq m+1$ and where the principal has not yet chosen $a_P = 1$, then $\sigma_A(h_{A,t}) \in (0,1]$.
- On any agent history $h_{A,t}$ where the principal has already chosen $a_P = 1$, then necessarily $\sigma_A(h_{A,t}) = 1$.

The Principal plays a strategy σ_P such that:

- On any Principal history $h_{P,t}$ with $t \leq m-2$, then $\sigma_P(h_{P,t}) = 0$.
- On any Principal history $h_{P,t}$ with $t = m-1$ and $a_t = 1$, and where she has not already chosen $a_P = 1$, then $\sigma_P(h_{P,t}) \in (0,1)$.
- On any Principal history $h_{P,t}$ with $t \geq m$ and $a_t = 1$, and where she has not already chosen $a_P = 1$, then $\sigma_P(h_{P,t}) = 1$.
- On any Principal history where she has already chosen $a_P = 1$, then necessarily $\sigma_P(h_{P,t}) = 0$, as repressive action can only be taken once.

The proof consists of four steps. In Steps 1, 2 and 3, we will suppose that signals s_t are *not* observed. That is, agents and Principal histories are simply $h_{A,t} = (\{a_{P,s}\}_{s\leq t-1})$ and $h_{P,t} = (\{a_s\}_{s\leq t}, \{a_{P,s}\}_{s\leq t-1})$. We will prove the existence of the previously-defined equilibrium in this setting.

Then, in Step 4, we will assume that signals s_t are observed. That is, agents and Principal histories are then $h_{A,t} = (\{s_s\}_{s\leq t}, \{a_{P,s}\}_{s\leq t-1})$ and $h_{P,t} = (\{s_s\}_{s\leq t}, \{a_s\}_{s\leq t}, \{a_{P,s}\}_{s\leq t-1})$. We will show that the previously-defined equilibrium is robust to the presence of such signals.

Step 1 (Showing existence of mixed strategies):

In what follows, we will denote the previously-defined mixed strategies at times $t = m-1$ and $t = m$ by $\sigma_{A,m-1}$, $\sigma_{A,m}$, and $\sigma_{P,m-1}$.

Specifically, we denote by $\sigma_{A,m-1}$ the strategy $\sigma_A(h_{A,t})$ on any agent history $h_{A,t}$ at $t = m-1$ and where the principal has not yet chosen $a_P = 1$. On such a history, the agent knows that if he acts, he will be the $(m-1)$-th agent to join the protest.

We denote by $\sigma_{A,m}$ the strategy $\sigma_A(h_{A,t})$ on any agent history $h_{A,t}$ with $t = m$ and where the principal has not yet chosen $a_P = 1$. On that history, the agent knows that there is a chance the movement has reached size $n_{t-1} = m-1$, and thus that the dissuasion threshold may have been passed.

Finally, we denote by $\sigma_{P,m-1}$ the Principal strategy $\sigma_P(h_{P,t})$ on any Principal history $h_{P,t}$ with $a_{A,t} = 1$ at $t = m-1$, and where she has not previously chosen $a_P = 1$. On such a history, the Principal knows that the agent has just caused the movement to cross the dissuasion threshold, and we will show that she will be indifferent (i.e. hesitant) to taking repressive action.

Let us thus first verify that, assuming the above strategies σ_A and σ_P are fixed on all other histories, there indeed exist mixed strategies $\sigma^*_{P,m-1}$, $\sigma^*_{A,m-1}$, and $\sigma^*_{A,m}$ that are part of the best-response set.

Step 1a (At time $t = m-1$, the agent can be made indifferent when the Principal chooses a mixed strategy $\sigma^*_{P,m-1}$):

Denote by $\Delta\mathbb{E}[\pi_{m-1,m-1}]$ the expected benefit of choosing action $a_{m-1} = 1$ over $a_{m-1} = 0$, for agent $i = m-1$ at time $t = m-1$, conditional on the agent history $h_{A,m-1} = (\{a_{P,j}\}_{j\leq m-2})$, and where the principal has not yet chosen $a_P = 1$.

We then have

$$\begin{aligned}\Delta\mathbb{E}[\pi_{m-1,m-1}] \quad = \quad & \sigma_{P,m-1}\cdot\mathbb{E}[\pi_{m-1,m-1}|a_{m-1}=1, a_{P,m-1}=1] \\ & +(1-\sigma_{P,m-1})\cdot\mathbb{E}[\pi_{m-1,m-1}|a_{m-1}=1, a_{P,m-1}=0] - \mathbb{E}[\pi_{m-1,m-1}|a_{m-1}=0, a_{P,m-1}=0].\end{aligned}$$

We see that $\Delta\mathbb{E}[\pi_{m-1,m-1}] > 0$ when $\sigma_{P,m-1}$ is 0,[29] and $\Delta\mathbb{E}[\pi_{m-1,m-1}] < 0$ when $\sigma_{P,m-1}$ is 1. Since $\Delta\mathbb{E}[\pi_{m-1,m-1}]$ is continuous in $\sigma_{P,m-1}$, it follows that there exists $\sigma^*_{P,m-1} \in (0,1)$ such that $\Delta\mathbb{E}[\pi_{m-1,m-1}] = 0$, and thus agent $m-1$ is indifferent between choosing $a = 1$ or $a = 0$.

Step 1b (At time m, the agent can be made indifferent when the previous agent chose a mixed strategy $\sigma^*_{A,m-1}$ at time $t = m-1$):

Denote by $\Delta\mathbb{E}[\pi_{m,m}]$ the expected benefit of choosing action $a_m = 1$ over $a_m = 0$, for agent $i = m$ at time $t = m$, conditional on the agent history $h_{A,m} = (\{a_{P,j}\}_{j\leq m-1})$, and where the principal has not yet chosen $a_P = 1$.

We then have

$$\begin{aligned}\Delta\mathbb{E}[\pi_{m,m}] \quad = \quad & \sigma_{A,m-1}\cdot\mathbb{E}[\pi_{m,m}|a_m=1, a_{m-1}=1, a_{P,m-1}=0] \\ & \quad +(1-\sigma_{A,m-1})\cdot\mathbb{E}[\pi_{m,m}|a_m=1, a_{m-1}=0, a_{P,m-1}=0] \\ & -\Big(\sigma_{A,m-1}\cdot\mathbb{E}[\pi_{m,m}|a_m=0, a_{m-1}=1, a_{P,m-1}=0] \\ & \quad +(1-\sigma_{A,m-1})\cdot\mathbb{E}[\pi_{m,m}|a_m=0, a_{m-1}=0, a_{P,m-1}=0]\Big) \\ = \quad & \sigma_{A,m-1}(\Delta\mathbb{E}[\pi_{m,m}|a_{m-1}=1, a_{P,m-1}=0]) \\ & \quad +(1-\sigma_{A,m-1})(\Delta\mathbb{E}[\pi_{m,m}|a_{m-1}=0, a_{P,m-1}=0]). \qquad (10)\end{aligned}$$

where $\Delta\mathbb{E}[\pi_{m,m}|a_{m-1}, a_{P,m-1}] = \mathbb{E}[\pi_{m,m}|a_m = 1, a_{m-1}, a_{P,m-1}] - \mathbb{E}[\pi_{m,m}|a_m = 0, a_{m-1}, a_{P,m-1}]$.

We see that $\Delta\mathbb{E}[\pi_{m,m}] > 0$ when $\sigma_{A,m-1}$ is 1, and $\Delta\mathbb{E}[\pi_{m,m}] < 0$ when $\sigma_{A,m-1}$ is 0. Since $\Delta\mathbb{E}[\pi_{m,m}]$ is continuous in $\sigma_{A,m-1}$, it follows that there exists $\sigma^*_{A,m-1} \in (0,1)$ such that $\Delta\mathbb{E}[\pi_{m,m}] = 0$, and thus

[29] Indeed, when $\sigma_{P,m-1} = 0$, agent $m-1$ does not get punished and thus he benefits from choosing action $a = 1$ (when δ_A and T are sufficiently large and when sufficiently many future agents also choose $a = 1$, which will be made clear later).

agent m is indifferent between choosing $a = 1$ or $a = 0$. The intuition is that agent m is unsure of whether the movement has already passed the dissuasion threshold, since he is unsure of whether $n_{m-1} = m - 2$ or $n_{m-1} = m - 1$. Thus, he is indifferent between joining or not, even if he knows that the Principal will take repressive action if he does. Indeed, the mixed strategy played by the previous agent causes the expected dilution of the punishment to be just enough to make him indifferent.

Step 1c: (The Principal can be made indifferent at time $t = m - 1$, when agent m chooses a mixed strategy $\sigma^*_{A,m}$):

Denote by $\Delta\mathbb{E}[\pi_{P,m-1}]$ the expected benefit of choosing action $a_{P,m-1} = 1$ over $a_{P,m-1} = 0$, for the Principal at time $t = m - 1$, conditional on the Principal history $h_{P,m-1} = (\{a_j\}_{j \le m-1}, \{a_{P,j}\}_{j \le m-2})$, and where the principal has not yet chosen $a_P = 1$. When agent $m - 1$ has chosen $a_{m-1} = 1$, we then have

$$\begin{aligned} \Delta\mathbb{E}[\pi_{P,m-1}] &= \mathbb{E}[\pi_{P,m-1}|\sigma_{A,m}, a_{P,m-1} = 1] - \mathbb{E}[\pi_{P,m-1}|\sigma_{A,m}, a_{P,m-1} = 0] \\ &= R(B) + \delta_P(\mathbb{E}[\pi_{P,m}|\sigma_{A,m}, a_{P,m-1} = 1] - \mathbb{E}[\pi_{P,m}|\sigma_{A,m}, a_{P,m-1} = 0]). \end{aligned} \tag{11}$$

We note that $\Delta\mathbb{E}[\pi_{P,m-1}] > 0$ when $\sigma_{A,m} = 1$, as punishing now is better than punishing later (under the expectation that all future agents choose $a = 1$). On the other hand, if $R(B)$ is sufficiently small, then by Lemma 3, $\Delta\mathbb{E}[\pi_{P,m-1}] < 0$ when $\sigma_{A,m} = 0$, as the Principal refrains from taking the repressive action when she expects a future agent to behave well enough. Moreover, since $\Delta\mathbb{E}[\pi_{P,m-1}]$ is continuous in $\sigma_{A,m}$, it follows that there exists $\sigma^*_{A,m} \in (0, 1)$ such that $\Delta\mathbb{E}[\pi_{P,m-1}] = 0$, and thus the Principal is indifferent between choosing $a_{P,m-1} = 1$ or $a_{P,m-1} = 0$ at time $m - 1$, when agent $m - 1$ has taken action $a_{m-1} = 1$.

From Steps 1a, 1b, and 1c we conclude that, assuming the above strategies σ_A and σ_P are fixed on all other histories, then there exist mixed strategies $\sigma^*_{P,m-1}$, $\sigma^*_{A,m-1}$, and $\sigma^*_{A,m}$ that are part of the best-response set.

Step 2 (Showing the optimality of pure strategies at times $t \le m - 2$):

We can now verify if σ_A and σ_P, at times $t \le m-2$, are optimal given σ_A and σ_P at times $t \ge m-1$.

From Lemma 3, the Principal will find $\sigma^*_{P,t} = 0$ to be optimal at times $t \le m - 2$, if $\sigma^*_{A,m-1}$ is low enough. To show that $\sigma^*_{A,m-1}$ can indeed be made low enough, first recall that $\frac{BC}{m} < \Delta v(\infty) - \kappa < \frac{BC}{m-1}$. For any $\epsilon > 0$, we can then choose C such that $\Delta v(\infty) - \kappa = \frac{BC}{m-1} - \epsilon$, and thus make the long-term net marginal benefit (absent the punishment cost) as close as we want to the expected punishment cost when $m - 2$ agents have previously joined the movement.

Second, note that as the expected net marginal benefit of choosing $a = 1$ (absent the punishment cost) becomes arbitrarily close to $\Delta v(\infty) - \kappa$ for high enough δ_A and high enough T, under the correct expectation that all future agents choose $a = 1$ with arbitrarily high probability (which will be shown in Step 3 when $R(B) \to 0$), it then follows that the expected marginal benefit for agent m of choosing $a = 1$ when $m - 2$ past agents have joined the movement, i.e.

$$\mathbb{E}[\pi_{m,m}|a_m = 1, a_{m-1} = 0, a_{P,m-1} = 0] - \mathbb{E}[\pi_{m,m}|a_m = 0, a_{m-1} = 0, a_{P,m-1} = 0],$$

can be made arbitrarily close to 0. On the other hand, the expected marginal benefit for agent m of choosing $a = 1$ when $m - 1$ past agents have joined the movement, i.e.

$$\mathbb{E}[\pi_{m,m}|a_m = 1, a_{m-1} = 1, a_{P,m-1} = 0] - \mathbb{E}[\pi_{m,m}|a_m = 0, a_{m-1} = 1, a_{P,m-1} = 0],$$

remains bounded above 0 since $\frac{BC}{m} < \Delta v(\infty) - \kappa$. It follows that, to make agent m indifferent (and thus to make $\Delta\mathbb{E}[\pi_{m,m}] = 0$ in Eq. (10) of Step 1b), the mixed strategy $\sigma^*_{A,m-1}$ chosen by agent $m-1$ can be made arbitrarily close to 0.

Third, such a low $\sigma^*_{A,m-1}$ means that agent $m-1$ is expected by the Principal to be sufficiently well-behaved, which then implies by Lemma 3 that there necessarily exists an upper bound $\underline{\mu}(\delta_P, R(B)) \in (0,1)$ (as well as other parameter bounds) such that $\mu_{m-1} < \underline{\mu}(\delta_P, R(B))$, and thus $\sigma^*_P(h_{P,t}) = 0$ is optimal at times $t \leq m-2$. It then in turn follows that agents will also find $\sigma^*_A(h_{A,t}) = 1$ to be optimal at times $t \leq m-2$.

We can thus conclude that there exist $\overline{\delta}_A \in (0,1)$, $\overline{T} > 0$, $\overline{\delta}_P \in (0,1)$, and $\underline{R} > 0$, such that when $\delta_A > \overline{\delta}_A$, $T > \overline{T}$, $\delta_P > \overline{\delta}_P$, and $R(B) < \underline{R}$, then $\sigma^*_P(h_{P,t}) = 0$ and $\sigma^*_A(h_{A,t}) = 1$ are optimal at times $t \leq m-2$.

Step 3: (Showing the optimality of pure strategies for the Principal at times $t \geq m$):

First note that, if the Principal were to choose $a_{P,m} = 0$, then agent $i = m+1$ would face the same decision problem as agent m, however since the probability that $n_{t-1} = m-1$ is no smaller at $t = m+1$ than at $t = m$, then agent $m+1$ is at least as tempted to choose $a = 1$ as agent m, and thus necessarily $\sigma^*_A(h_{A,m+1}) \geq \sigma^*_{A,m}$. By induction, $\sigma^*_A(h_{A,t+1}) \geq \sigma^*_A(h_{A,t})$ for all $t \geq m$.

Second, from Eq. (11) in Step 1c, we also observe that $0 < \Delta\mathbb{E}[\pi_{P,m-1}] = R(B)$ when $\sigma_{A,m} = 1$. On the other hand, when $\sigma_{A,m} = 0$, $\Delta\mathbb{E}[\pi_{P,m-1}] < 0$ and bounded away from 0 if $R(B)$ is small enough. By continuity, it follows that as $R(B) \to 0$, then $\sigma^*_{A,m} \to 1$.

Third, this then implies that $\sigma^*_A(h_{A,t}) \to 1$, for all $t \geq m+1$, when $R(B) \to 0$. It follows that, as $R(B) \to 0$, the probability of any subgame, starting at $t = m$, in which some agents choose $a_i = 0$ goes to zero. It is thus optimal for the Principal to choose $a_{P,m} = 1$ when $a_m = 1$. By the same argument, at times $t > m$, it is also optimal for her to choose $a_{P,t} = 1$ on any Principal history $h_{P,t}$ where the agent has chosen $a_t = 1$ and the Principal has not already taken $a_P = 1$. Hence, $\sigma^*_P(h_{P,t}) = 1$ on any Principal history $h_{P,t}$ where $t \geq m$, the agent has chosen $a_t = 1$ and the Principal has not already taken $a_P = 1$.

Naturally, after the Principal has chosen $a_{P,t} = 1$ at some $t \geq m$, then all subsequent agents choose $a_t = 1$.

Thus, in this equilibrium, only agent $m-1$ may choose $a = 0$ with non-trivial probability. It follows that the probability that, in any equilibrium realization, more than a single agent chooses $a = 0$ can be made arbitrarily small as $R(B) \to 0$.

Thus, we conclude that for any $\eta > 0$, there exist $\underline{R} > 0$, $\overline{T} > 0$, $\overline{\delta}_A \in (0,1)$, and $\overline{\delta}_P \in (0,1)$, such that when $R(B) < \underline{R}$, $T > \overline{T}$, $\delta_A > \overline{\delta}_A$, and $\delta_P > \overline{\delta}_P$, the probability that at least $T-1$ agents chooses $a = 1$ is greater than $1 - \eta$.

Step 4 (Showing that the previous equilibrium survives in the presence of signals s_t about the size of the movement)

We now suppose that the agents observe signals s_t and thus that $h_{A,t} = (\{s_s\}_{s \leq t}, \{a_{P,s}\}_{s \leq t-1})$.

Note that for $t \leq m-1$, the signal is not informative as $\sigma_A(h_{A,t}) = 1$ for $t \leq m-2$ and the agent thus already knows that $n_{t-1} = t-1$.

For $t = m$, however, the agent can learn from the signal s_m. By Bayes' rule, his belief about $n_{t-1} = m-1$ can be expressed as

$$\mathbb{P}(n_{t-1} = m-1 | s_m = x, \sigma_{A,m-1}) = \frac{\mathbb{P}(s_m = x | n_{t-1} = m-1)\sigma_{A,m-1}}{\mathbb{P}(s_m = x | n_{t-1} = m-1)\sigma_{A,m-1} + \mathbb{P}(s_m = x | n_{t-1} = m-2)(1 - \sigma_{A,m-1})} \quad (12)$$

where $x \in \{0,1\}$, $\mathbb{P}(s_m = 1 | n_{t-1} = y) = g(y)$ and $\mathbb{P}(s_m = 0 | n_{t-1} = y) = 1 - g(y)$.

Thus, we can re-express Eq. (10) in Step 1b as

$$\begin{aligned}\Delta\mathbb{E}[\pi_{m,m}] &= \mathbb{P}(n_{t-1}=m-1|s_m=x,\sigma_{A,m-1})(\Delta\mathbb{E}[\pi_{m,m}|a_{m-1}=1,a_{P,m-1}=0]) \\ &\quad +(1-\mathbb{P}(n_{t-1}=m-1|s_m=x,\sigma_{A,m-1}))(\Delta\mathbb{E}[\pi_{m,m}|a_{m-1}=0,a_{P,m-1}=0]).\end{aligned}$$

From Eq. (12), $\mathbb{P}(n_{t-1}=m-1|s_m=x,\sigma_{A,m-1})$ is continuous in $\sigma_{A,m-1}$, with $\mathbb{P}(n_{t-1}=m-1|s_m=x,\sigma_{A,m-1})=1$ when $\sigma_{A,m-1}=1$, and $\mathbb{P}(n_{t-1}=m-1|s_m=x,\sigma_{A,m-1})=0$ when $\sigma_{A,m-1}=0$. By an argument analogous to Step 2, we can then find an equilibrium mixed strategy $\sigma^*_{A,m-1}$ that satisfies Lemma 3, which then implies that $\sigma^*_P(h_{P,t})=0$ for $t\leq m-2$, which in turn implies that $\sigma^*_A(h_{P,t})=1$ for $t\leq m-2$.

As the rest of Steps 1 and 3 are not affected by the observation of s_t, this concludes the proof. ■

Proof of Proposition 4 (Equilibria with simultaneous actions). Part (I):

In the absence of the Principal, $a^*_i=0,\forall i$, is a pure strategy equilibrium. Indeed, $\Delta\pi_i(\vec{a}_{-i})=v(1,0)-\kappa-v(0,0)<0$ by assumption when $\vec{a}_{-i}=\vec{0}$ and hence no agent i would want to deviate from $a^*_i=0$.

Likewise $a^*_i=1,\forall i$, is a pure strategy equilibrium. Indeed $\Delta\pi_i(\vec{a}_{-i})=v(1,T-1)-\kappa-v(0,T-1)>0$ by assumption when $\vec{a}_{-i}=\vec{1}$ and hence no agent i would want to deviate from $a^*_i=1$.

Now call $\sigma_A\in(0,1)\subset\Delta(\{0,1\})$ a symmetric (mixed) strategy followed by the agents. σ^*_A is an equilibrium strategy when it satisfies

$$\mathbb{E}[\Delta\pi_i(\vec{a}_{-i})|\sigma_A]=\mathbb{E}[v(1,\sum_{j\neq i}a_j)-\kappa-v(0,\sum_{j\neq i}a_j)|\sigma_A]=0.$$

Such a σ^*_A exists since $\mathbb{E}[\Delta\pi_i(\vec{a}_{-i})|\sigma_A]$ is continuous in σ_A and since by assumption $\mathbb{E}[\Delta\pi_i(\vec{a}_{-i})|\sigma_A=0]<0$ and $\mathbb{E}[\Delta\pi_i(\vec{a}_{-i})|\sigma_A=1]>0$.

Part (II):

The Principal will direct punishment only at agents who have chosen action $a=1$, since she gets no utility from punishing agents who chose $a=0$. Thus, $\mathbb{P}\{\phi_i=1|\sigma_P\}=\min(\frac{B}{\sum_j a_j},1)$ for an agent who has chosen $a_i=1$. We will show that under such a strategy, the situation where $a^*_i=1$ for all i in equilibrium cannot be ruled out when B/T (her budget relative to the number of agents) is low enough.

Let $1\leq B<T$. Consider a profile of agents' actions $a_i=1,\ \forall i$. Consider the strategy σ_P under which the Principal punishes with equal probability agents who have chosen $a=1$, so that $\mathbb{P}\{\phi_i=1|\sigma_P\}=B/T$. There exists $\beta\in(0,1)$ such that, when $B/T<\beta$, then $\Delta v(T-1)>\kappa$ (by assumption) and $\mathbb{P}\{\phi_i=1|\sigma_P\}=B/T<\beta$ for all i. Thus with the action profile $\vec{a}=\vec{1}$,

$$\begin{aligned}\mathbb{E}[\Delta\pi_i(\vec{a}_{-i})|\sigma_P] &= \Delta v(T-1)-\kappa-C\cdot\mathbb{E}[\phi_i|\sigma_P] \\ &= \Delta v(T-1)-\kappa-C\cdot\mathbb{P}\{\phi_i=1|\sigma_P\} \\ &> 0\end{aligned}$$

for all i when $B/T<\beta$ and hence all agents playing $a^*_i=1$ is an equilibrium.

This implies that $a^*_i=0$ cannot be the only agent behavior that can occur in equilibrium when B/T is small enough (i.e. when the Principal's budget is positive, but small relative the number of agents). We see that β is simply equal to $\frac{\Delta v(T-1)-\kappa}{C}$.

By a similar argument, there exists $\gamma\in(0,1)$, with $\gamma\geq\beta$, such that when $B/T>\gamma$, then the only equilibrium involves $a^*_i=0$ for all agents, since then each agent has a large enough probability

$\mathbb{P}\{\phi_i = 1|\sigma_P\}$ of being punished. Here, such a γ also corresponds to $\frac{\Delta v(T-1)-\kappa}{C}$. Indeed, in such a case

$$\begin{aligned}\mathbb{E}[\Delta\pi_i(\vec{a}_{-i})|\sigma_P] &= \Delta v(T-1) - \kappa - C \cdot \mathbb{P}\{\phi_i = 1|\sigma_P\} \\ &= \Delta v(T-1) - \kappa - C \cdot \frac{B}{T} \\ &< 0,\end{aligned}$$

which rules out an equilibrium where $a_i = 1$ for all i, but also rules out any equilibrium where some agents choose $a_i = 1$. Indeed, $\frac{\Delta v(T-1)-\kappa}{C} < \frac{B}{T}$ implies that $\frac{\Delta v(n-1)-\kappa}{C} < \frac{B}{n}$ for any $n < T$ and thus $\mathbb{E}[\Delta\pi_i(\vec{a}_{-i})|\sigma_P] < 0$ under any agents' actions profile. ■

Proof of Proposition 5 (Equilibrium when the Principal can stop the growth of the protest movement).

Part (i):

Let $h_{P,t}$ be a history on which $m-2$ agents have chosen $a = 1$ before time t and where $a_t = 1$. Then, the marginal expected benefit of taking repressive action can be expressed as

$$\begin{aligned}\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P, \vec{a})|\sigma_P, \sigma_A, h_{P,t}] &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 1, \sigma_P, \sigma_A, h_{P,t}] - \mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]\Big) \\ &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 1, \sigma_P, \sigma_A, h_{P,t}] - (-(m-1+s-t) + R(B)\mathcal{A}_s)\Big) \\ &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(p(B)(-(m-1) + R(B)) + (1-p(B))(-(m-1+s-t) + R(B)) \\ &\qquad -(-(m-1+s-t) + R(B)\mathcal{A}_s)\Big) \\ &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(p(B)(s-t) + R(B) - R(B)\mathcal{A}_s\Big) \\ &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(p(B)(s-t) + R(B)(1-\mathcal{A}_s)\Big).\end{aligned}$$

In the above, $\mathcal{A}_s = \mathbb{E}[\sum_j a_{P,j,s}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]$ is the probability that the repressive action has been taken by time s conditional on $a_{P,t} = 0$. Naturally, $\mathcal{A}_s \in [0,1]$, since $\sum_j a_{P,j,s} \in \{0,1\}$ for any s (as the Principal can use action $a_{P,j}$ at only one time j). Moreover, here $\mathcal{A}_t = 0$ since we condition on $a_{P,t} = 0$.

For any $\delta_P \in (0,1)$ and $p(B) > 0$, there exists $\overline{R} < 0$ such that for all $R(B) > \overline{R}$,

$$\sum_{s=t}^{T} \delta_P^{s-t}\Big(p(B)(s-t) + R(B)(1-\mathcal{A}_s)\Big) > 0.$$

Thus, $\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P, \vec{a})|\sigma_P, \sigma_A, h_{P,t}] > 0$ and it is optimal for the Principal to choose $a_{P,t} = 1$.

Punishment is thus credible and dissuasive in the eyes of the agents, when the $(m-1)$-th agent takes action $a = 1$, and such an agent would thus never take action $a = 1$ in any equilibrium. Propositions 1 and 2 thus hold under these conditions.

Part (ii):

Let $h_{P,t}$ be a history on which agent t is the first to have chosen $a = 1$. Then,

$$\begin{aligned}
\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P, \vec{a})|\sigma_P, \sigma_A, h_{P,t}] &= \sum_{s=t}^{T} \delta_P^{s-t}\Big(\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 1, \sigma_P, \sigma_A, h_{P,t}] - \mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]\Big) \\
&= \sum_{s=t}^{T} \delta_P^{s-t}\Big(p(B)(R(B) - 1) + (1 - p(B))(R(B) - 1 - (s - t)) \\
&\qquad -\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]\Big) \\
&= \sum_{s=t}^{T} \delta_P^{s-t}\Big(R(B) - 1 - (1 - p(B))(s - t) - \mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}]\Big) \\
&\geq \sum_{s=t}^{T} \delta_P^{s-t}\Big(R(B) - 1 - (1 - p(B))(s - t) - \mathcal{B}_s\Big) \\
&= , \sum_{s=t}^{T} \delta_P^{s-t}\Big(\mathcal{C}_s - (1 - p(B))(s - t)\Big)
\end{aligned}$$

where $\mathcal{B}_t = -1$ and $\mathcal{B}_s = -1 + R(B)$ for $s > t$, and thus where $\mathcal{C}_t = R(B)$ and $\mathcal{C}_s = 0$ for $s > t$.

The inequality follows from the fact that $\mathbb{E}[\tilde{\pi}_{P,s}|a_{P,t} = 0, \sigma_P, \sigma_A, h_{P,t}] \leq \mathcal{B}_s$, since $\mathcal{B}_s$ is the highest expected flow payoff the Principal could achieve by choosing $a_{P,t} = 0$ (which would happen if she chose $a_{P,t+1} = 1$, the agent chose $a_{t+1} = 0$, and she succeeded in stopping the growth of the movement at time $t + 1$).

Thus, for any $R(B) > 0$ and $\delta_P \in (0, 1)$, there exists $p_{th} > 0$ such that when $p(B) > p_{th}$, then $\Delta\mathbb{E}[\pi_{P,t}(\vec{a}_P, \vec{a})|\sigma_P, \sigma_A, h_{P,t}] > 0$ and the Principal finds it optimal to choose action $a_{P,t} = 1$ as soon as one agent has joined the protest movement. Given such a strategy, no agent will find it optimal to take action $a = 1$ and thus in any equilibrium, $a_i^* = 0$ for all agents. ■

Proof of Corollary 1 (Comparative statics in B). By Proposition 5(ii), the maximal size of the protest movement is $\overline{n} = 0$ when $p(B) \geq p_{th}$. Since $p : \mathbb{N}_+ \to [0, 1]$ is increasing, one to one and onto, then it follows that (generically) there exists $\underline{B} \in \mathbb{N}_+$ such that $p(\underline{B}) > p_{th} > p(\underline{B} - 1)$. Thus, $\overline{n} = 0$ when $B \geq \underline{B}$.

Necessarily, $\overline{n} \geq 0$ on $B \in [1, \underline{B})$. ■

Proof of Lemma 4. Note that the condition $\frac{BC}{m} < \Delta v(\infty) - \kappa < \frac{BC}{m-1}$ implies that m is the largest integer such that $\Delta v(\infty) - \kappa < \frac{BC}{m-1}$.

Since $\frac{C}{m-1} \leq \frac{BC}{m-1}$, then there are two possibilities: Either $\frac{C}{m} < \Delta v(\infty) - \kappa < \frac{C}{m-1}$, in which case $\tilde{m} = m$. Or, $\frac{C}{m-1} < \Delta v(\infty) - \kappa$, in which case the largest $\tilde{m}$ such that $\Delta v(\infty) - \kappa < \frac{C}{\tilde{m}-1}$ must be strictly smaller than m.

Thus, it must be that $\tilde{m} \leq m$. ■

Proof of Proposition 6 (Partial and full deterrence under sequential punishment). In the model presented in Section 4.5, at each time t, agent t gets a chance to act with probability $q_A \in [0, 1]$ and the Principal gets the chance to act (immediately after the agent) with probability $q_P \in [0, 1]$. Note that in Proposition 6, and in the rest of the article, it is implicitly assumed that $q_A = 1$ and $q_P = 1$, and thus agent t always gets a chance to act at any time t, and the Principal can act immediately after. Proposition 6 is thus a special case of Proposition 7. Indeed, in Proposition 7(ii)-(iii), for any $q_A \in (0, 1]$, there exists $\overline{q}_P(q_A) \in (0, 1)$ such that the result holds when $q_P > \overline{q}_P(q_A)$. This implies that Proposition 7(ii)-(iii) hold when $q_A = 1$ and $q_P = 1$, as in the case of Proposition 6.

Specifically, Proposition 6(i) follows from Proposition 7(ii), Proposition 6(ii) follows from Proposition 7(iii). ■

Proof of Proposition 7 (Reaction time and deterrence).

Consider the model presented in Section 4.5. At each time t, agent t gets a chance to act with probability $q_A \in [0,1]$ and the Principal gets a chance to act (immediately after the agent) with probability $q_P \in [0,1]$.

Part (i):

In Lemma 1, it was implicitly assumed that $q_A = 1$. Also note that $\mathbb{E}[\Delta\pi_{t,t}(\vec{a}_{-t}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,t}]$, the marginal benefit of choosing action $a = 1$ over action $a = 0$, as defined in Eq. (9), is also continuous in q_A. Thus, by continuity, there exists $\overline{q}_A \in (0,1)$ such that when $q_A > \overline{q}_A$, the result of Lemma 1 necessarily holds, and thus all agents choose $a_i^* = 1$ in any equilibrium.

Likewise, in Lemma 1, it was implicitly assumed that $q_P = 0$, since the Principal never had a chance to act. For that reason, as said at the beginning of the proof of Lemma 1, the last term in Eq. (9), that is $\sum_{s=t}^{T} \delta_A^{s-t} C\mathbb{P}\{\phi_{t,s} = 1|a_t = 1, \sigma_P, \sigma_A, h_{A,t}\}$, was zero for any t. Note that $\mathbb{P}\{\phi_{t,s} = 1|a_t = 1, \sigma_P, \sigma_A, h_{A,t}\}$, the probability that agent t will have been punished by time s, is continuous in q_P, the probability that the Principal gets a chance to act at any given time. By continuity, for any $\epsilon > 0$ there exists $\underline{q}_P(q_A) \in (0,1)$ such that when $q_P < \underline{q}_P(q_A)$, then $\sum_{s=t}^{T} \delta_A^{s-t} C\mathbb{P}\{\phi_{t,s} = 1|a_t = 1, \sigma_P, \sigma_A, h_{A,t}\} < \epsilon$. Hence, the same result as in Lemma 1 continues to hold, as $\mathbb{E}[\Delta\pi_{t,t}(\vec{a}_{-t}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,t}]$, the expected marginal payoff of any agent, remains positive for small enough q_P and thus all agents choose $a_i^* = 1$ in any equilibrium.

This concludes the proof of Part (i).

Part (ii):

Consider first the case where $B_0 = 1$.

Suppose that q_P is large enough given q_A. This ensures that between two consecutive agents' actions, the Principal has an opportunity to move with high probability. Suppose also that δ_A is large enough. This will ensure that agents value sufficiently the future actions of other agents. Suppose then that, at time $t = i'$, $\tilde{m} - 2$ agents have already chosen $a = 1$ and the Principal has not yet used her budget, i.e $B_{i'} = 1$. If an additional agent (i.e. agent i') chooses $a_{i'} = 1$, then the Principal will choose $a_{P,k} = 1$ as soon as she has an opportunity to move (i.e. at the first $k \geq i'$). This is so because otherwise, if the Principal does not punish, all subsequent agents will choose $a = 1$ as $\frac{C}{\tilde{m}} < \Delta v(\infty) - \kappa$ (with the correct anticipation that all other agents will choose $a = 1$).

To see this, suppose that the Principal has chosen not to punish, i.e. $a_{P,k} = 0$. Consider then the next agent i who gets to act at a time i such that $i' \leq k < i \leq k'$. Using Eq. (9), and when δ_A is large enough, this agent i's expected marginal payoff from choosing action $a_i = 1$ can be expressed as

$$\begin{aligned}
\mathbb{E}[\Delta\pi_{i,i}(\vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{A,i}] &= \sum_{s=i}^{T} \delta_A^{s-i} \Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa | \sigma_P, \sigma_A, h_{A,i}] - C\mathbb{P}\{\phi_{i,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \Big) \\
&\geq \sum_{s=i}^{T} \delta_A^{s-i} \Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa | \sigma_P, \sigma_A, h_{A,i}] - \frac{C}{\tilde{m}} \Big) \\
&> 0,
\end{aligned}$$

where the first inequality follows from the fact that $\mathbb{P}\{\phi_{i,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \leq \frac{1}{\tilde{m}}$. Indeed, if the principal later chooses to punish (i.e. take action $a_{P,\tau} = 1$) at some time $\tau \geq k'$, agent i will be punished (i.e. $\tilde{\phi}_{i,\tau} = 1$) with probability $\frac{1}{\tilde{m}}$ if no further agent has yet taken action $a = 1$ after agent i, or with probability less than $\frac{1}{\tilde{m}}$ if further agents have taken action $a = 1$. It follows that $\mathbb{P}\{\phi_{i,s} = 1|a_i = 1, \sigma_P, \sigma_A, h_{A,i}\}$ cannot be greater than $\frac{1}{\tilde{m}}$.

The second inequality is established by the following argument: Since $\Delta v(\infty) - \kappa > \frac{C}{\tilde{m}}$, then when sufficiently many agents choose $a = 1$ (or are expected to do so in the future), the expected net benefits exceed the expected punishment cost $\frac{C}{\tilde{m}}$. Indeed, under the correct expectation that all future agents choose $a = 1$, then for all $\xi > 0$, there exists s' (with $s' > i$) such that when $s > s'$, then $\big(\Delta v(\infty) - \kappa\big) - \mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa | \sigma_P, \sigma_A, h_{A,i}] < \xi$. It thus follows that there exists $\overline{\delta}_A > 0$, such that $\sum_{s=i}^{T} \delta_A^{s-i} \Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa | \sigma_P, \sigma_A, h_{A,i}] - \frac{C}{\tilde{m}} \Big) > 0$ when $\delta_A > \overline{\delta}_A$, under the correct expectation that all future agents will choose $a = 1$.

Thus, since $\mathbb{E}[\Delta \pi_{i,i}(\vec{a}_{-i}, \vec{a}_P) | \sigma_P, \sigma_A, h_{A,i}] > 0$, the Principal prefers choosing $a_{P,k} = 1$, anticipating that all subsequent agents will choose $a = 1$ anyway. She prefers to punish because she enjoys punishing those who chose $a = 1$ (i.e. $R(B) > 0$ in Eq. (3)). She prefers doing it at the earliest opportunity (i.e. at time k) because of discounting (i.e. $\delta_P \in (0, 1)$ in Eq. (4)).

Now given this, after $\tilde{m} - 2$ agents have chosen $a = 1$, the subsequent agent will not be willing to choose $a = 1$ as this will trigger a punishment and $\Delta v(\infty) - \kappa < \frac{C}{\tilde{m}-1}$. Thus, there cannot be more than $\tilde{m} - 2$ agents choosing $a = 1$.

To see this, consider again the next agent i who gets to act at a time $t = i$ such that $i' \leq k < i \leq k'$, but now assume that $a_{i'} = 0$ so that only $\tilde{m} - 2$ agents have already chosen $a = 1$. Using Eq. (9), and when q_P is large enough given q_A, this agent i's expected marginal payoff from choosing action $a_i = 1$ can be expressed as

$$\begin{aligned} \mathbb{E}[\Delta \pi_{i,i}(\vec{a}_{-i}, \vec{a}_P) | \sigma_P, \sigma_A, h_{A,i}] &= \sum_{s=i}^{T} \delta_A^{s-i} \Big(\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) - \kappa | \sigma_P, \sigma_A, h_{A,i}] - C\mathbb{P}\{\phi_{i,s} = 1 | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \Big) \\ &\leq \sum_{s=i}^{T} \delta_A^{s-i} \Big(\Delta v(\infty) - \kappa - C\mathbb{P}\{\phi_{i,s} = 1 | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \Big) \\ &< 0, \end{aligned}$$

where the first inequality follows from $\mathbb{E}[\Delta v(\sum_{j \neq i} a_{j,s}) | \sigma_P, \sigma_A, h_{A,i}] \leq \Delta v(\infty)$. To establish the second inequality, first note that, by an argument analogous to the one stated earlier, $\mathbb{P}\{\phi_{i,s} = 1 | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \leq \frac{1}{\tilde{m}-1}$. Moreover, $\mathbb{P}\{\phi_{i,s} = 1 | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} \to \frac{1}{\tilde{m}-1}$ as $q_P \to 1$, for any given fixed q_A. This implies that $\forall \xi > 0$, $\exists\ \overline{q}_P(q_A) \in (0, 1)$ such that when $q_P > \overline{q}_P(q_A)$, then $\mathbb{P}\{\phi_{i,s} = 1 | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} > \frac{1}{\tilde{m}-1} - \xi$. Since, $\Delta v(\infty) - \kappa - \frac{C}{\tilde{m}-1} < 0$, this in turn implies that $\Delta v(\infty) - \kappa - C\mathbb{P}\{\phi_{i,s} = 1 | a_i = 1, \sigma_P, \sigma_A, h_{A,i}\} < 0$ when $q_P > \overline{q}_P(q_A)$, thus establishing the second inequality above.

Therefore, the Principal reacting quickly enough to an agent's action (i.e. $q_P > \overline{q}_P(q_A)$) ensures that $\mathbb{E}[\Delta \pi_{i,i}(\vec{a}_{-i}, \vec{a}_P) | \sigma_P, \sigma_A, h_{A,i}] < 0$, and thus that the $(\tilde{m} - 1)$-th agent (as well as all subsequent ones) does not choose $a = 1$.

For $B_0 > 1$, continuing on the above, it is readily verified by backward induction on B_0 that after $max(0, \tilde{m} - 1 - B_0)$ agents have chosen $a = 1$, if an additional agent chooses $a = 1$, the Principal will choose to punish, thereby deterring any $a = 1$ after such an event.

Finally, an argument analogous to the one made at the end of the proof of Proposition 1 also ensures that a sufficiently patient Principal will not take action $a_P = 1$ when $\tilde{m} - 1 - B_0$ or fewer agents have taken action $a = 1$.

Note that the condition $q_P > \overline{q}_P(q_A)$ is obviously satisfied in Proposition 6, since it is assumed that $q_P = 1$ (i.e., the Principal always gets a chance to act at time t).

Part (iii):

By Part (ii), we know that there can be at most $max(0, \tilde{m} - 1 - B_0)$ agents choosing $a = 1$. Given that $\tilde{m} - 1 - B_0 < N$, no agent can find choosing $a = 1$ profitable. Indeed, by Assumption 1, fewer than $N - 1$ agents choosing $a = 1$ does not make the collective action worthwhile (i.e. $\Delta v(n) < \kappa$ for $n < N - 1$). Here, this is ensured by the condition $\tilde{m} - 1 - B_0 < N$. ■

Proof of Proposition 8 (Equilibrium with commitment and personalized punishment). Let σ_P be such that the Principal uses a first offender punishment strategy.

Given some agent i's action time i, the Principal gets a first chance to punish agent i at some random time $\tau_k = i+w$, where w is geometrically distributed over a support $\mathbb{N}$ and has probability mass function $(1-q_P)^w q_P$. Given any agent's history $h_{A,i}$ such that all previous agents $j<i$ choose $a_j = 0$, then we can write the last term on the righthand side of Eq. (9) as

$$\begin{aligned}\sum_{s=i}^{\infty} \delta_A^{s-i} C\mathbb{P}\{\phi_{i,s}=1|a_i=1,\sigma_P,\sigma_A,h_{A,i}\} &= \sum_{\tau_k=i}^{\infty}\Big(C\sum_{s=\tau_k}^{\infty}\delta_A^{s-i}\mathbb{P}\{\phi_{i,s}=1|a_i=1,\sigma_P,\sigma_A,h_{A,i}\}\Big)f(\tau_k)\\ &= \sum_{\tau_k=i}^{\infty}\Big(C\sum_{s=\tau_k}^{\infty}\delta_A^{s-i}\Big)f(\tau_k)\\ &= \sum_{\tau_{k_i}=i}^{\infty}\Big(C\frac{\delta_A^{\tau_k-i}}{1-\delta_A}\Big)f(\tau_k),\end{aligned}$$

where $f(\tau_k)$ denotes the probability mass function of τ_k. The first equality follows from the fact that $\mathbb{P}\{\phi_{i,s}=1|a_i=1,\sigma_P,\sigma_A,h_{A,i}\}=0$ for $s<\tau_k$, while the second equality follows from the fact that $\mathbb{P}\{\phi_{i,s}=1|a_i=1,\sigma_P,\sigma_A,h_{A,i}\}=1$ for $s\geq\tau_k$ under a first offender punishment strategy, as i will surely be punished at time τ_k.

As $\sum_{\tau_k=i}^{\infty}\frac{\delta_A^{\tau_k-i}}{1-\delta_A}f(\tau_k)\uparrow\frac{1}{1-\delta_A}$ when $q_P\to 1$, then $\forall\delta_A\in(0,1)$ and $\epsilon>0$, there exists $\overline{q}_P^c(\delta_A)\in(0,1)$ such that when $q_P>\overline{q}_P^c(\delta_A)$,

$$\sum_{s=i}^{\infty}\delta_A^{s-i}C\mathbb{P}\{\phi_{i,s}=1|a_i=1,\sigma_P,\sigma_A,h_{A,i}\}>\frac{1}{1-\delta_A}C(1-\epsilon). \tag{13}$$

Moreover, since

$$\begin{aligned}\sum_{s=i}^{\infty}\delta_A^{s-i}\mathbb{E}[\Delta v(\sum_{j\neq i}a_{j,s})-\kappa|\sigma_P,\sigma_A,h_{A,i}] &< \sum_{s=i}^{\infty}\delta_A^{s-i}\big(\lim_{n\to\infty}\Delta v(n)-\kappa\big)\\ &= \frac{1}{1-\delta_A}\cdot\big(\lim_{n\to\infty}\Delta v(n)-\kappa\big)\end{aligned}\tag{14}$$

and $\lim_{n\to\infty}\Delta v(n)-\kappa<C$ by Assumption 1, it then follows from Eqs. (13) and (14) that

$$\begin{aligned}\mathbb{E}[\Delta\pi_{i,i}(\vec{a}_{-i},\vec{a}_P)|\sigma_P,\sigma_A,h_{A,i}] &< \frac{1}{1-\delta_A}\cdot\big(\lim_{n\to\infty}\Delta v(n)-\kappa\big)-\big(\frac{1}{1-\delta_A}C(1-\epsilon)\big)\\ &< 0\end{aligned}$$

when $q_P>\overline{q}_P^c(\delta_A)$.

Hence no agent i wants to be the first to choose $a_i=1$. Denoting by $\sigma_A(0|h_{A,i})$ the probability of choosing action $a=0$ after history $h_{A,i}$ under strategy σ_A, we conclude that $\sigma_A(0|h_{A,i})=1$ is an optimal strategy after any such history $h_{A,i}$. Applying this reasoning by induction to all $i'>i$ yields that the unique equilibrium involves $a_i^*=0$ for all i.

It is immediate from the above that here, and in contrast to the proof of Proposition 7, the bound $\overline{q}_P^c(\delta_A)$ is independent of q_A.

Finally, note that the requirement that $\delta_A>\overline{\delta}_A$, for some $\overline{\delta}_A\in(0,1)$, is simply to ensure that, absent a Principal, agents will always find it optimal to choose $a=1$, as in Lemma 1 or Proposition 7. ■

References

Battaglini, M. and T. R. Palfrey (2024): "Organizing for Collective Action: Olson Revisited," *Journal of Political Economy*, 132, 2881–2936.

——— (2025): "Dynamic collective action and the power of large numbers," *(Working paper)*.

Bueno de Mesquita, B., D. P. Myatt, A. Smith, and S. A. Tyson (2024): "The Punisher's Dilemma," *The Journal of Politics*, 86, 395–411.

Burdzy, K., D. M. Frankel, and A. Pauzner (2001): "Fast equilibrium selection by rational players living in a changing world," *Econometrica*, 69, 163–189.

Calcagno, R., Y. Kamada, S. Lovo, and T. Sugaya (2014): "Asynchronicity and coordination in common and opposing interest games," *Theoretical Economics*, 9, 409–434.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of monetary Economics*, 12, 383–398.

Cantoni, D., A. Kao, D. Y. Yang, and N. Yuchtman (2024): "Protests," *Annual Review of Economics*, 16.

Cantoni, D., D. Y. Yang, N. Yuchtman, and Y. J. Zhang (2019): "Protests as strategic games: experimental evidence from Hong Kong's antiauthoritarian movement," *The Quarterly Journal of Economics*, 134, 1021–1077.

Carlsson, H. and E. Van Damme (1993): "Global games and equilibrium selection," *Econometrica: Journal of the Econometric Society*, 989–1018.

Chen, E., A. Ghersengorin, and S. Petersen (2024): "Imperfect recall and AI delegation," .

Corman, H. and N. Mocan (2005): "Carrots, sticks, and broken windows," *The Journal of Law and Economics*, 48, 235–266.

Correa, S. (2025): "Persistent protests," *American Economic Journal: Microeconomics*, 17, 321–357.

Doval, L. and J. C. Ely (2020): "Sequential information design," *Econometrica*, 88, 2575–2608.

Frankel, D. and A. Pauzner (2000): "Resolving indeterminacy in dynamic settings: the role of shocks," *The Quarterly Journal of Economics*, 115, 285–304.

Frankel, D. M. (2023): "Equilibrium Selection in Participation Games: A Unified Framework," Tech. rep., Mimeo, Melbourne Business School.

Gale, D. (1995): "Dynamic coordination games," *Economic theory*, 5, 1–18.

Kamada, Y. and M. Kandori (2020): "Revision games," *Econometrica*, 88, 1599–1630.

Kandori, M. (1992): "Social norms and community enforcement," *The Review of Economic Studies*, 59, 63–80.

Kandori, M. and S. Obayashi (2014): "Labor union members play an OLG repeated game," *Proceedings of the National Academy of Sciences*, 111, 10802–10809.

Koh, A., S. Sanguanmoo, and K. Uzui (2024): "Informational Puts," *arXiv preprint arXiv:2411.09191*.

Morris, S. and H. S. Shin (1998): "Unique equilibrium in a model of self-fulfilling currency attacks," *American Economic Review*, 587–597.

——— (2001): "Global games: Theory and applications," .

Myerson, R. B. (1998): "Population uncertainty and Poisson games," *International Journal of Game Theory*, 27, 375–392.

——— (2000): "Large poisson games," *Journal of Economic Theory*, 94, 7–45.

Nishihara, K. (1997): "A resolution of N-person prisoners' dilemma," *Economic theory*, 10, 531–540.

Persico, N. (2023): "A theory of non-democratic redistribution and public good provision," *(Working paper)*.

Salcedo, B. (2017): "Interdependent choices," Tech. rep., Working Paper.

Schelling, T. C. (1960): *The Strategy of Conflict*, Harvard university press.

Takahashi, S. (2010): "Community enforcement when players observe partners' past play," *Journal of Economic Theory*, 145, 42–62.