On the Power of Reaction Time in Deterring Collective Actions*

Philippe Jehiel†
Paris School of Economics
& University College London

Mathieu V. Leduc‡
Paris School of Economics

10 September 2024

Abstract

Multiple agents (e.g. policymakers, civil society activists) play a dynamic coordination game, each having an opportunity to act at a separate random time. A principal (e.g. a lobby/NGO, a government) can, at random times, take costly punitive actions to dissuade the agents. If the principal, irrespectively of her budget, reacts sufficiently quickly to the agents' actions, there is a unique equilibrium where agents are deterred from coordinating. Moreover, if the principal can ex ante commit to a punishment strategy, deterrence is ensured even when agents act much more quickly. This brings new perspectives on the regulation of lobbies and policing.

Keywords: Dynamic games, coordination games, equilibrium selection, reaction time
JEL codes: D00; C72; C73; D70

*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). Namely, 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 and Leeat Yariv. A Supplementary Online Appendix is available at `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4957235`

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

1 Introduction

In expressing support for a given policy, politicians (henceforth, policymakers) must often make public statements *asynchronously.* For example, they may be interviewed in the print or television media at random times. Likewise, on social media, appropriate opportunities to respond to a particular tweet or post may also come at random times, for each of them. Moreover, only when sufficiently many other policymakers have expressed support for a given policy—or are expected to do so in the near future—does it become worth expressing support for the policy, leading to a dynamic coordination problem.

In addition to this dynamic coordination problem, the policymakers inclined to supporting the policy may also face opposition from a lobby (or NGO) with diverging interests. The lobby may try to dissuade the policymakers by taking retaliatory action against those who express support for the policy.

Note that this basic setting can be applied to other situations. For instance, the same dynamic coordination problem is faced by civil society activists trying to build a movement. In this case, it is a government police force that may try to prevent this movement from forming. Although applications are multiple, in this article we will mainly focus on the policy makers/lobby example. This is for convenience and we will occasionally analyze other examples when appropriate.

It has often been argued that lobbies/NGOs must be well-funded in order to effectively influence policymakers. Indeed, dissuading policymakers from taking a position can be a costly undertaking that may involve lawsuits or the release of public statements against policymakers who have expressed support for the policy, and if many policymakers are engaged in the collective action, this may require enormous funding abilities for lobbies to be effective. Based on this observation, some argue that lobbies should be regulated by limiting their funding[1].

In this article, we examine this claim and will argue that a lobby's ability to react quickly to the policymakers' actions is a key determinant of the effectiveness of the lobbying activity, more so than its financial ability in our environment in which actions are taken sequentially at different times.

Specifically, we study a dynamic model where agents (policymakers) with shared interests are each given an opportunity, at a random time, to take a binary action (i.e. expressing support for a policy or not). This can model, for example, a random opportunity to give a television interview or to reply to a particular tweet or social media post. These random opportunities are driven by a Poisson process with a given arrival rate. The marginal benefit of expressing support for the policy is increasing in the number of policymakers who have (and who will) express support for it, thus defining a dynamic coordination game among policymakers. On the other side, a Principal (a lobby) with interests diverging from those of the agents is also given random opportunities to take retaliatory actions against the agents, in an effort to prevent them from coordinating. These random opportunities are driven by another Poisson process with a different arrival rate. The Principal has a limited budget and can thus only take retaliatory actions against a finite number of agents.

We first show that in the absence of the Principal (or when the Principal is slow enough), the ability of the policymakers to dynamically coordinate depends on the arrival rate of their opportunities to express their support for the policy. Indeed, if the sequential asynchronous opportunities that they are given arrive at a fast enough rate, we show by a subgame perfection argument that in any equilibrium, policymakers all choose to express support for the policy. This effectively selects a unique equilibrium behavior on the part of policymakers.

We then show that if the lobby can react quickly enough to the actions of the policymakers, in equilibrium, policymakers are fully dissuaded from expressing support for the policy, *irrespectively* of the size of its budget. More precisely, for this conclusion to hold, we only need the lobby to have a budget ensuring the effective punishment of a single policymaker, and we also need the lobby to react sufficiently

[1] See, for instance, Hasen (2012), Briffault (2008) or Johnson (2006) for a broad discussion of this topic.

quickly after the action of any policymaker. Under these conditions, after such a putative action by a policymaker and considering the time at which the lobby can implement its punishment, there would be little chance that another policymaker would also be given the opportunity to act. Thus, the lobby would find it unambiguously optimal to punish that one policymaker who expressed his support and, as a result, in equilibrium no policymaker dares express support for the policy.

We next note that if a lobby can perfectly observe the separate times at which the policymakers expressed their viewpoints, and if it can publicly announce and commit to a retaliation strategy, then it is able to fully dissuade the policymakers from expressing support for the policy, no matter how the reaction speeds of the lobby and the policymakers compare to one another. It can do so by threatening to punish the policymakers in the order in which they expressed their support, since then no policymaker will want to be the first to express his support. As already mentioned, for our deterrence results to hold, the lobby only needs to have a budget large enough to make the punishment of a single agent effective. In other words, with a "single bullet" the lobby can discipline an entire population of policymakers.

Our findings contrast with those obtained in the equivalent simultaneous-actions coordination game, in which the size of the lobby's budget (i.e. its ability to punish a large number of policymakers) is key to ensuring that, in all equilibria, policymakers are deterred from expressing support for the policy.

Our result that without the Principal, policmakers are able to coordinate on the outcome that is efficient for them in a dynamic version of the coordination game is reminiscent of the work of Gale (1995), who developed a similar insight in a different context of private provision of public goods with asynchronous (yet deterministic) decision times. Our results in the presence of the Principal, while simple, have no counterpart in the literature as far as we know. We believe our results are particularly well suited to the understanding of lobbying activities in the age of social media technologies. Indeed, the latter have had the effect of increasing the speed at which opportunities arrive for policymakers and lobbyists. To the extent that lobbies are better at handling social media technologies (and thus are quicker in their reaction times than policymakers are in expressing their views), our analysis suggests that lobbies do not need large budgets in order to be effective.[2]

From a technical viewpoint, our modeling of stochastic decision times is somehow similar to that adopted in the recent 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. However, to our knowledge, that literature has not considered the kind of games discussed in our application. 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), Morris and Shin (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.[3] 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.

Our paper is also related to the literature on the economics of policing (e.g. Owens and Ba (2021); Knowles, Persico, and Todd (2001)) and to the literature on community enforcement (e.g. Kandori (1992); Takahashi (2010); Kandori and Obayashi (2014)).

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 policymakers and of the lobby when the policymakers' action times are perfectly observable and when the lobby may or may not be able to

[2]To be effective, the actions of the policmakers would have to be decided jointly or in such a way that the lobbyist cannot observe them separately.

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

publicly announce and commit to a retaliation strategy. We state our main results and also compare our dynamic setting to a game where policymakers act simultaneously. In Section 4, we conclude and discuss some extensions and robustness checks. A Supplementary Online Appendix further discusses such extensions.

2 Model

2.1 Setting

We study a setting where at random times governed by a Poisson process, different infinitely-lived agents (e.g. policymakers) 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. The number of agents $N_A(0,t)$ getting the chance to play the game in an interval of time $[0,t]$ is thus $N_A(0,t) \sim Poiss(\lambda_A t)$ where λ_A is the agents' arrival rate. Let t_i be the event time of the ith event of this Poisson point process. At such a time, the ith agent will get the chance to choose an action $a_i \in \{0,1\}$.

A Principal (e.g. a lobby with interests diverging from those of the agents) also gets the chance to police the population of agents, but at some random times. The number of times $N_P(0,t) \sim Poiss(\lambda_P t)$ the Principal gets to act in an interval of time $[0,t]$ is thus also governed by a Poisson process where λ_P is the intensity of the Principal's policing activity. Call τ_k the event time of the kth event of this Poisson point process. At each such time, the Principal can choose an action $a_{P,k} \in \mathbb{N}_+$ (the identity of the agent who is to be punished) or $a_{P,k} = \emptyset$ (not take punitive action).

Let the history of play at time t be denoted by

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

Calling $\mathcal{H}$ the set of possible histories, the Principal's strategy is then $\sigma_P : \mathcal{H} \to \Delta(\{\emptyset, \mathbb{N}_+\})$. That is, at some action time τ, based on a history of past play h_τ, the Principal can choose to punish any 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. It is important to note that the Principal has limited resources and is thus restricted in the number of agents she can punish. Call B_t, with $B_0 = B \in \mathbb{N}$, the Principal's finite budget at time t. It represents the maximum number of agents she can choose to punish now and in the future. Thus, for $\tau_{k-1} < t \leq \tau_k$, where τ_{k-1} and τ_k are action times for the Principal, $B_t = B_{\tau_{k-1}} - \mathbb{1}_{\{a_{P,k-1} \neq \emptyset\}}$. That is, the Principal has the resources to punish at most B agents and her budget B_t decreases every time she punishes an agent.

Likewise, the agents' strategy is $\sigma_A : \mathcal{H} \to \Delta(\{0,1\})$. That is, at some action time t_i, based on a history of past play h_{t_i}, an agent i can choose to take action $a_i = 0$ or $a_i = 1$ and he can also randomize.

2.2 Payoffs

Given action times $\{t_j\}_{j=1}^{\infty}$ and $\{\tau_k\}_{k=1}^{\infty}$ for the agents and Principal, let us denote the agents' and Principal's actions profiles as $\vec{a} = (a_1, a_2, ...)$ and $\vec{a}_P = (a_{P,1}, a_{P,2}, ...)$.

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 $a_{j,t} = a_j \in \{0,1\}$ if $t_j \leq t$ and thus agent j has already acted. By default, $a_{j,t} = 0$ if $t_j > t$ and the agent has not yet acted. Likewise $\vec{a}_{P,t} = \{a_{P,1,t}, a_{P,2,t}, ...\}$ with $a_{P,k,t} = a_{P,k} \in \{0,1\}$ if $\tau_k \leq t$ and, by default, $a_{P,k} = 0$ if $\tau_k > t$.

2.2.1 Agents' payoffs

At any time t, agent i receives a flow payoff

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

In the above equation, $v : \mathbb{N}^2 \to \mathbb{R}$ 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 $\sum_j a_{j,t}$. Thus, at time t an agent benefits from the actions of all the agents who chose $a_j = 1$ up to time t. $\kappa > 0$ is the intrinsic cost to agent i of taking action $a_i = 1$. $C > 1$ is the punishment cost felt by agent i if he is punished by the Principal and $\phi_{i,t} = \{\exists \tau_s \leq t : a_{P,s} = i\}$ is the event that agent i is punished no later than time t. Without loss of generality, we let $v(0, 0) = 0$. Moreover, we let $v(1, 0) < \kappa$ and $v(1, n-1) - v(0, n-1) > \kappa$ for all $n \geq N$ and some $N \in \mathbb{N}_+$, capturing the fact that agents play a coordination game among themselves. Thus, it is not worth taking action $a_i = 1$ if no other agent takes it, while it becomes worth taking action $a_i = 1$ when sufficiently many other agents also take it. We also assume that $\lim_{n\to\infty} v(1, n) - v(0, n) - \kappa < C$ so that an agent always suffers from being punished, irrespectively of how many other agents have chosen action $a_i = 1$. These properties of v are summarized in the following assumption.

Assumption 1 (Properties of benefit function) *(i)* $v(0, 0) = 0$. *(ii) Let* $\Delta v(n) = v(1, n) - v(0, n)$. $\Delta v(n)$ *is increasing in* n*, with* $\Delta v(0) < \kappa$, $\Delta v(N-1) > \kappa$ *for some* $N \in \mathbb{N}_+$ *and* $\lim_{n\to\infty} \Delta v(n) - \kappa < C$.

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

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

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 $\sigma_A^*(h_{t_i})$ to maximize his expected payoff $\mathbb{E}[\pi_{i,t_i}(a_i, \vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{t_i}]$, given the Principal's strategy, the other agents' strategy, and a history of play at time t_i.

2.2.2 Principal's payoff

The Principal's flow payoff at time t is

$$\tilde{\pi}_{P,t}(\vec{a}_{P,t}, \vec{a}_t) = -\sum_j a_{j,t} + \sum_j \epsilon \mathbb{1}_{\psi_{j,t}} \cdot a_{j,t}, \tag{3}$$

where $\epsilon \in (0, 1)$ and $\psi_{j,t} = \{\exists \tau_k \leq t : \{a_{P,k} = j\} \bigcap \{a_{P,s} \neq j, \forall \tau_s \leq t \text{ such that } \tau_s \neq \tau_k\}\}$.

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[4] a permanent benefit ϵ from having punished (i.e. $\{a_{P,k} = j\}$) agents who had chosen action $a_j = 1$ and who had not yet been punished (i.e. $\{a_{P,s} \neq j, \forall \tau_s \leq t \text{ such that } \tau_s \neq \tau_k\}$).

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

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

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

[4]The Supplementary Online Appendix offers a microfoundation for ϵ, in which the punishment results in the action of the targeted agent to be possibly cancelled.

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

3 Equilibrium analysis

Let the full history of play h_τ be observable to the Principal at a time τ when she takes her own action. Since $h_\tau = (\{t_j\}_{t_j \leq \tau}, \{a_j\}_{t_j \leq \tau}, \{\tau_k\}_{\tau_k \leq \tau}, \{a_{P,k}\}_{\tau_k < \tau})$, this includes both the identity j of the agents who took action $a_j = 1$ and the times t_j at which they took it. This fits well with an application where policymakers (the agents) are given opportunities, at random times, to express their support in the media for a given policy, while a lobby (the Principal) can then attack them if they took a certain position, when it is given an opportunity to react.

The equilibrium behavior will depend on the timing of actions (namely, on how quickly the Principal can react to the agents' actions), on the Principal's ability to commit to a strategy, as well as on her budget. We will analyze these in the following subsections.

3.1 Successful agent coordination

The dynamic nature of the game allows to select a unique equilibrium behavior for the agents. Namely, if the intensity of the agents' activity is high enough and the intensity of the Principal's policing activity is low enough, the agents always succeed in coordinating on action $a = 1$.

Proposition 1 (Successful agent coordination) *There exist $\underline{\lambda}_P > 0$ and $\overline{\lambda}_A(\delta_A, N) > 0$, such that when $\lambda_P < \underline{\lambda}_P$ and $\lambda_A > \overline{\lambda}_A(\delta_A, N)$, then 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 Proposition 1, note that if the lobby's policing activity is slow enough, then the chance of being punished in the not-too-distant future — by which we mean in a period of time that is not too severely discounted by the discount factor δ_A^{s-t} — can be low enough that agents always have an interest in choosing $a = 1$. Indeed, if the agents' arrival rate λ_A is high enough, 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 $a_N = 1$. As this happens in the not-too-distant future with high probability when λ_A is high enough, it is then 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. This allows agents to coordinate dynamically.

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 (yet deterministic) decision times.

3.2 Successful deterrence without commitment from the Principal

On the contrary, if the intensity of the agents' activity is low enough *relative* to the Principal's, then she can successfully deter the agents from choosing action $a = 1$, *irrespectively of the size of her budget.* This is formalized in the following proposition.

Proposition 2 (Equilibrium without commitment from the Principal) *There exist $\underline{\eta}$ such that if $\lambda_A - \lambda_P < \underline{\eta}$, then any equilibrium involves $a_i^* = 0$ for all i and for any $B > 0$.*

Thus, if the Principal acts sufficiently quickly relative to the agents, then with a single bullet she can discipline an entire population. Indeed, as long as $B > 0$ (which holds for $B = 1$), the actual size of the budget B is not important to obtain that, in any equilibrium, all agents choose the action $a = 0$. This illustrates how the reaction time of a lobby is more important than its ability to take punitive actions against a large number of policymakers (e.g. its financial ressources).

If $\lambda_A - \lambda_P$ is too large, then an equilibrium where agents coordinate on $a = 1$ can be sustained. This can be done requiring that the Principal randomizes her punishment uniformly among agents having chosen $a = 1$. Then, agents would be confident enough that when the Principal moves, there would be a large number of agents having chosen $a = 1$, making the punishment ineffective at deterring this action.

In an application to crime and policing, it is interesting to note a parallel with the famous[5] "broken window theory", 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 λ_P. In this famous theory, the actions of the police must also be visible, which is the case in our model as $a_{P,k}$ is observed by all agents.

3.3 Successful deterrence with commitment from the Principal

If the Principal can credibly commit to a strategy σ_P at time 0, we will see that she is allowed to react much more *slowly* than in the case without commitment and still succeed in deterring agents from coordinating on action $a = 1$. Indeed, since we suppose the full history of play h_τ is observable to the Principal at a time τ when she takes her own action, she has access to both the identity of the agents who took action $a = 1$ and the times t_j at which they took it, from which she is able to deduce the order in which agents took action $a = 1$. Interestingly in this case, if the Principal announces and commits to a strategy that consists in punishing the agents *in the order* in which they have chosen action $a_i = 1$, then (unless λ_P is so low that an agent has little chance of being punished in the not-too-distant future) the equilibrium involves $a_i^* = 0$ for all i. In other words, not only is the size of her budget $B > 0$ once again irrelevant to successfully dissuading the agents from choosing action $a = 1$, but the required intensity λ_P of the Principal's policing activity is also independent of the intensity λ_A of the agents' activity. λ_P only needs to be large enough so that an agent has a reasonable chance of being policed in the not-too-distant future.

This is formalized in the following definition and proposition.

Definition 1 (Ordered punishment strategy) *σ_P is called an* ordered punishment strategy *if the Principal punishes agents in the order in which they took action $a = 1$. That is $\sigma_P(h_{\tau_1}) = i$, where i is such that $a_i = 1$ for $t_i < \tau_1$ and $a_j = 0$ for all $t_j < t_i$. Likewise, $\sigma_P(h_{\tau_2}) = i'$, where i' is such that $a_{i'} = 1$ for $t_{i'} < \tau_2$ and $a_j = 0$ for all $t_j < t_{i'}$ except for $t_j = t_i$ where $a_i = 1$, and so on.*

Proposition 3 (Equilibrium with commitment from the Principal) *Let the Principal commit to an ordered punishment strategy σ_P^*. (I) There exists $\overline{\lambda}_P^c > 0$ such that for* any *budget size $B > 0$, if $\lambda_P > \overline{\lambda}_P^c$, then the equilibrium involves agents choosing $a_i^* = 0$ for all i. (II) Moreover, $\overline{\lambda}_P^c$ is decreasing in an agent's discount rate δ_A and it does not depend on λ_A.*

Thus, when a lobby can announce and commit to a punishment strategy, the information it possesses about the order in which policymakers took action $a = 1$ (e.g. made statements in the media in support of a policy that goes against the lobby's interests) is not only more important than the size of its budget (and thus its ability to take retaliatory action against multiple policymakers), it is also more important than its reaction time. Again, the threat of a single bullet can discipline an entire population, but under even less restrictive conditions than in the case without commitment.

[5]See, for instance, Corman and Mocan (2005).

In Proposition 3, λ_P must only be high enough so that an agent has a large enough chance of being punished *in the not-too-distant future*. $\overline{\lambda}_P^c$ is thus completely independent of the agents' activity rate λ_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.

Specifically, the Principal can achieve the same outcome as in Proposition 3 if she only recalls the identity of the first[6] offender!

3.4 Contrasts with a game where agents act simultaneously

The previous analysis contrasts sharply with the equivalent simultaneous-actions coordination game, in which agents must act at the same time. Consider M agents, each of whom can choose an action $a \in \{0,1\}$ at time 0. The set of agents $\{1,2,...,M\}$ is *unordered*, in the sense that an agent's identity i is just a label or a name. The Principal then observes the action profile $\vec{a}$ and, at time 1, chooses which agents to punish, i.e. her action $a_P \subset \emptyset \bigcup \{1,...,M\}$ with $|a_P| \leq B$ is the set of agents she punishes (noting that she cannot punish more than B agents, the size of her budget). Payoffs are realized at time 1.

Agent i has payoff

$$\pi_i(a_i, \vec{a}_{-i}, a_P) = v(a_i, \sum_j a_j) - \kappa \cdot a_i - C \cdot \mathbb{1}_{\phi_i} \tag{5}$$

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

The Principal has payoff

$$\pi_P(a_P, \vec{a}) = -\sum_j a_j + \sum_j \epsilon \mathbb{1}_{\phi_j} \cdot a_j. \tag{6}$$

Call $\sigma_P : \{0,1\}^M \rightarrow \Delta(2^{\{1,2,...,M\}})$ the Principal's punishment strategy. It is a mapping from a time-0 agents' actions profile $\vec{a}$, which she observes, to the set of probability measures over all subsets of agents.

As in Assumption 1, we let $v(1,0) - v(0,0) < \kappa < v(1, M-1) - v(0, M-1)$ 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 $1 \leq B < M$, 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 $\min(\frac{B}{\sum_j a_j}, 1)$. Indeed, here the Principal cannot condition her punishment strategy $\sigma_P(\vec{a})$ on the (unordered) identities of the agents, but only on their actions. 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 investing would be too low while the probability of being punished is too large. This is summarized in the following proposition.

Proposition 4 *Consider the game where M 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 $B/M < \frac{\Delta v(M-1) - \kappa}{C}$, 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 $B/M > \frac{\Delta v(M-1) - \kappa}{C}$.

[6]The Online Supplementary Appendix provides such an elaboration.

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, $\frac{\Delta v(M-1)-\kappa}{C}$ is bounded away from 0, we conclude that the Principal would need a budget B that also grows very large as M gets large to be sure to deter any $a_i = 1$ in equilibrium. This is to be contrasted with our finding in the dynamic version of the game, for which we obtained that $B \geq 1$ was enough to deter any $a_i = 1$ in equilibrium under the conditions of Propositions 2 or 3.

4 Conclusions and extensions

Finally we discuss certain extensions, a more detailed version of which is included in the Supplementary Online Appendix.

4.1 Effect of the information structure available to the Principal

Suppose now that the actual times t_j at which agents acted are not publicly observed, but their actions are. Thus, the Principal cannot perfectly deduce the order in which agents took action $a = 1$. This fits well with situations like political demonstrations or riots, where the Principal (e.g. the government) can observe which agents (e.g. demonstrators) participated in a demonstration, but not the exact timing at which they joined the protest.

As the Principal cannot do better than punishing a randomly chosen agent, the probability that an agent gets punished will depend on the intensities λ_A and λ_P governing the action opportunities of the agents and of the Principal. It is important to note that this is true both without *and* with commitment. Thus, the benefits of being able to announce and commit to a punishment strategy—which, as stated in Proposition 3, allowed the Principal to deter collective action irrespectively of the intensity of the agents' activity λ_A—disappear. Commitment is useful when information about the timing of the agents' actions allows the Principal to announce and design an ordered punishment strategy (cf. Definition 1). Without such information, commitment loses its advantage and the Principal must rely on his reaction speed to deter agents.

4.2 Presence of fearless agents

Suppose there are two types of agents: rational and fearless, i.e. $\theta_i \in \{R, F\}$. The fearless agents do not fear punishment and thus cannot be deterred by the Principal.

If θ_i is publicly observable, 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. However, if θ_i is private (not publicly observable), the Principal will have to punish any agent taking action $a = 1$ to maintain credibility and will thus unavoidably deplete her budget. The presence of fearless agents can thus ultimately allow rational agents to coordinate on action $a = 1$. The only way for the Principal to deter rational agents from ultimately choosing action $a = 1$ in this setting would be to have an infinite budget B, irrespectively of her reaction speed.

Interestingly, the Principal would benefit more from an improvement in the type detection technology–which allows her to differentiate fearless from rational agents–than from an increase in her budget, again illustrating the importance of factors such as information in allowing the Principal to deter collective actions.

5 Proofs

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

At time t_i, given some history h_{t_i}, a Principal's strategy σ_P and a strategy σ_A for the agents, agent i's expected marginal payoff from choosing $a_i = 1$ as opposed to $a_i = 0$ can be written as

$$\begin{aligned}\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i},\vec{a}_P)|\sigma_P,\sigma_A,h_{t_i}] &= \int_{s=t_i}^{\infty}\delta_A^{s-t_i}\Big(\mathbb{E}[\Delta v(\sum_j a_{j,s}) - \kappa|\sigma_P,\sigma_A,h_{t_i}] - \mathbb{E}[C\mathbb{1}_{\phi_{i,s}}|a_i=1,\sigma_P,\sigma_A,h_{t_i}]\Big)ds \\ &= \int_{s=t_i}^{\infty}\delta_A^{s-t_i}\Big(\mathbb{E}[\Delta v(\sum_j a_{j,s}) - \kappa|\sigma_P,\sigma_A,h_{t_i}] - C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}\Big)ds \quad (7)\end{aligned}$$

where the expectation on the righthand side is taken over a_j and t_j.

Proof of Proposition 1.

Note that if the last term in Eq. (7), that is $\int_{s=t_i}^{\infty}\delta_A^{s-t_i}C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds$, is small enough for all i, then by continuity the equilibrium will be the same as in a game without the Principal. This occurs when $\lambda_P < \underline{\lambda}_P$. Indeed, we can rewrite it as

$$\begin{aligned}\int_{s=t_i}^{\infty}\delta_A^{s-t_i}C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds &= C\int_{s=t_i}^{T}\delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds + \\ &\quad C\int_{s=T}^{\infty}\delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds. \quad (8)\end{aligned}$$

Moreover, $\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\} \leq \mathbb{P}\{t_i \leq \tau_{k_i} \leq s\}$, where τ_{k_i} is the first time the Principal gets a chance to act after t_i. Since for any $\epsilon' > 0$ and $T > 0$, there exists $\underline{\lambda}_P > 0$ such that $\mathbb{P}\{t_i \leq \tau_{k_i} \leq s\} < \epsilon'$ when $\lambda_P < \underline{\lambda}_P$ and $s < T$, then the first term on the righthand side of Eq. (8) can be made arbitrarily small.

The second term on the righthand side of Eq. (8) can also be made arbitrarily small when T gets large. Indeed,

$$\begin{aligned}C\int_{s=T}^{\infty}\delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds &\leq C\int_{s=T}^{\infty}\delta_A^{s-t_i}ds \\ &= C\frac{\delta_A^{s-t_i}}{\ln\delta_A}\Big|_{s=T}^{\infty} \\ &= 0 - C\frac{\delta_A^{T-t_i}}{\ln\delta_A} \\ &= K(T) \\ &> 0, \quad (9)\end{aligned}$$

where $K(T) \downarrow 0$ as $T \to \infty$.

Thus, we conclude that $\forall\epsilon > 0$, there exists $\underline{\lambda}_P > 0$ such that

$$\int_{s=t_i}^{\infty}\delta_A^{s-t_i}C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds < \epsilon$$

when $\lambda_P < \underline{\lambda}_P$.

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 when $\lambda_P < \underline{\lambda}_P$, 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 $\Delta v(N-1) - \kappa > 0$:

$$\mathbb{E}[\Delta\pi_{N,t_N}(\vec{a}_{-N}, \vec{a}_P)|\sigma_P, \sigma_A, h_{t_N}] = \int_{s=t_N}^{\infty} \delta_A^{s-t_N} \Big(\mathbb{E}[\Delta v(\sum_j a_{j,s})|\sigma_P, \sigma_A, h_{t_N}] - \kappa - C\mathbb{P}\{\phi_{N,s}|a_N = 1, \sigma_P, \sigma_A, h_{t_N}\}\Big) ds > 0,$$

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

Finally, let t_1 be the first time at which an agent acts and call this agent $i = 1$. Note that if λ_A is high enough, then agent 1 will have positive expected marginal benefit of choosing action $a = 1$, since then, with high probability, he precipitates a subgame in which all agents will choose $a = 1$.

Consider agent $i = 1$. Let $\overline{\lambda}_A(\delta_A, N)$ be such that, given a fixed profile of actions $\vec{a}_{-1} = \vec{1}$ for the other agents, then

$$\int_{s=t_1}^{\infty} \delta_A^{s-t_1} \Big(\mathbb{E}[\Delta v(\sum_j a_{j,s})|\sigma_P, \sigma_A, h_{t_1}] - \kappa\Big) ds > 0, \quad \forall \lambda_A > \overline{\lambda}_A(\delta_A, N).$$

A high enough λ_A indeed garantees that, in expectation, sufficiently many other agents (i.e. more than $N-1$) will get the chance to act (and take action $a_j = 1$) in the not-too-distant future (which depends on the discount factor δ_A) in order to make it worthwhile for agent $i = 1$ to take action $a_i = 1$.

Specifically, if $N = 2$ and $\lambda_A > \overline{\lambda}_A(\delta_A, 2)$, 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 $a_2 = 1$. Thus, agent 1 will never choose $a_1 = 0$ and thus $a_i = 1$ for all i is part of any subgame perfect Nash equilibrium.

Likewise, if $N = 3$ and $\lambda_A > \overline{\lambda}_A(\delta_A, 3)$, 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 $a_2 = 1$ and it follows that agent 1 will never choose $a_1 = 0$. 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 $\lambda_A > \overline{\lambda}_A(\delta_A, N)$, then $a_i = 1$ for all i is part of any subgame perfect Nash equilibrium. It is trivial to show that $\overline{\lambda}_A(\delta_A, N)$ is increasing in N and decreasing in δ_A. ■

Proof of Proposition 2. Let τ_I be the first event time at which the Principal observes at least one agent who chose action $a = 1$. Call this earliest cohort of offending agents $\mathcal{A}_I$. From Eq. (3), the Principal gets the same benefit from punishing *any* offending agent. Thus let σ_P be some strategy by which the Principal punishes some member of $\mathcal{A}_I$.

Let us now examine an agent i's marginal benefit from choosing action $a_i = 1$.

If $\lambda_A - \lambda_P$ is sufficiently negative, then the first agent will get punished with very high probability after he chooses $a_1 = 1$, since the Principal tends to react very quickly (before a next agent has the chance to choose $a_2 = 1$). Then as long as $B > 0$, no agent will dare take the action $a = 1$ and thus $a_i = 0$, for all i, will be part of any equilibrium.

To formalize this, call τ_{k_i} the first time the Principal has the opportunity to act after t_i. Then, $\tau_{k_i} = t_i + s$ where $s \sim exp(1/\lambda_P)$ is exponentially distributed with mean $1/\lambda_P$. Thus, $\forall \epsilon > 0$, $\exists\, \underline{\eta}$ such that $\mathbb{P}\{\tau_{k_i} < t_{i+1}\} > 1 - \epsilon$ when $\lambda_A - \lambda_P < \underline{\eta}$ and for any $B > 0$.

Now consider any history h_{t_i}, where $a_j = 0$ for all $t_j < t_i$ (i.e. no agent before agent i has chosen action $a = 1$). In agent i's expected marginal payoff, note that the first term of the integral in Eq. (7) can be expressed as

$$\begin{aligned} \int_{s=t_i}^{\infty} \delta_A^{s-t_i} \mathbb{E}[\Delta v(\sum_j a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{t_i}] ds &< \int_{s=t_i}^{\infty} \delta_A^{s-t_i} \cdot \big(\lim_{n\to\infty} \Delta v(n) - \kappa\big) ds \\ &= -\frac{1}{\ln(\delta_A)} \cdot \big(\lim_{n\to\infty} \Delta v(n) - \kappa\big), \end{aligned}$$

where the inequality follows from Assumption 1.

Second, note that as said earlier, $\forall \epsilon > 0$, $\exists\ \underline{\eta}$ such that $\mathbb{P}\{\tau_{k_i} < t_{i+1}\} > 1-\epsilon$ when $\lambda_A - \lambda_P < \underline{\eta}$ and for any $B > 0$. Thus, the second term of the integral in Eq. (7) can be expressed as

$$\begin{aligned}
\int_{s=t_i}^{\infty} \delta_A^{s-t_i} C\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\}ds &= \mathbb{E}[C\int_{s=\tau_{k_i}}^{\infty} \delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\}ds] \\
&= \int_{\tau_{k_i}=t_i}^{\infty} \Big(C\int_{s=\tau_{k_i}}^{\infty} \delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\}ds\Big)f(\tau_{k_i})d\tau_{k_i} \\
&\geq \int_{\tau_{k_i}=t_i}^{\infty} \Big(C\int_{s=\tau_{k_i}}^{\infty} \delta_A^{s-t_i}\mathbb{P}\{\tau_{k_i} < t_{i+1}\}ds\Big)f(\tau_{k_i})d\tau_{k_i} \\
&> \int_{\tau_{k_i}=t_i}^{\infty} \Big(C\int_{s=\tau_{k_i}}^{\infty} \delta_A^{s-t_i}(1-\epsilon)ds\Big)f(\tau_{k_i})d\tau_{k_i} \\
&= \int_{\tau_{k_i}=t_i}^{\infty} \Big(-C(1-\epsilon)\frac{\delta_A^{\tau_{k_i}-t_i}}{\ln(\delta_A)}\Big)f(\tau_{k_i})d\tau_{k_i}.
\end{aligned}$$

The first equalities follow from the fact that $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} = 0$ for $s < \tau_{k_i}$, since agent i cannot be punished before the Principal has had a chance to act at time τ_{k_i}. The first (weak) inequality follow from the fact that, with some Principal's strategy σ_P as described above, $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} \geq \mathbb{P}\{\tau_{k_i} < t_{i+1}\}$. Indeed, $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} = 1$ for $\tau_{k_i} < s < t_{i+1}$, as then agent i is punished with certainty if he chose $a_i = 1$, while $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} \geq 0$ for $t_{i+1} < \tau_{k_i} < s$ as then agent i may still be punished with some probability. The second (strict) inequality follows when $\lambda_A - \lambda_P < \underline{\eta}$, as previously mentioned. Finally, noting that $\int_{\tau_{k_i}=t_i}^{\infty} \frac{\delta_A^{\tau_{k_i}-t_i}}{\ln(\delta_A)} f(\tau_{k_i})d\tau_{k_i} \uparrow \frac{-1}{\ln(\delta_A)}$ as $(\lambda_A - \lambda_P) \to -\infty$, then it follows that $\forall \epsilon > 0$, there exists $\underline{\eta}$ such that when $\lambda_A - \lambda_P < \underline{\eta}$, then

$$\begin{aligned}
\int_{s=t_i}^{\infty} \delta_A^{s-t_i} C\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\}ds &> \int_{\tau_{k_i}=t_i}^{\infty} \Big(-C(1-\epsilon)\frac{\delta_A^{\tau_{k_i}-t_i}}{\ln(\delta_A)}\Big)f(\tau_{k_i})d\tau_{k_i} \\
&> -C(1-\epsilon)\frac{1-\epsilon}{\ln(\delta_A)}.
\end{aligned}$$

From these observations and from Assumption 1, it then follows that when $\lambda_A - \lambda_P < \underline{\eta}$,

$$\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{t_i}] < 0.$$

Thus, agent i never takes action $a_i = 1$ and thus no agent ever takes action $a = 1$. It follows that in any equilibrium, $a_i^* = 0$ for all i. ∎

Proof of Proposition 3.

Let σ_P be such that the Principal uses an ordered punishment strategy.

Part (I):

Given some agent i's action time t_i, the Principal gets a first chance to punish agent i at some random time $\tau = t_i + s$, where $s \sim exp(1/\lambda_P)$. Given any history h_{t_i} such that all previous agents $j < i$ choose $a_j = 0$, then similarly as in the proof of Proposition 2, we can write the last term on the righthand side of Eq. (7) as

$$\begin{aligned}\int_{s=t_i}^{\infty} \delta_A^{s-t_i} C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds &= \int_{\tau_{k_i}=t_i}^{\infty}\Big(C\int_{s=\tau_{k_i}}^{\infty}\delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds\Big)f(\tau_{k_i})d\tau_{k_i} \\ &= \int_{\tau_{k_i}=t_i}^{\infty}\Big(C\int_{s=\tau_{k_i}}^{\infty}\delta_A^{s-t_i}ds\Big)f(\tau_{k_i})d\tau_{k_i} \\ &= \int_{\tau_{k_i}=t_i}^{\infty}\Big(-C\frac{\delta_A^{\tau_{k_i}-t_i}}{\ln\delta_A}\Big)f(\tau_{k_i})d\tau_{k_i}. \qquad (10)\end{aligned}$$

The first equality follows from the fact that $\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}=0$ for $s<\tau_{k_i}$, while the second equality follows from the fact that $\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}=1$ for $s\geq\tau_{k_i}$ under an ordered punishment strategy, as i will surely be punished at time τ_{k_i}.

As $\int_{\tau_{k_i}=t_i}^{\infty}\frac{\delta_A^{\tau_{k_i}-t_i}}{\ln(\delta_A)}f(\tau_{k_i})d\tau_{k_i}\uparrow\frac{-1}{\ln(\delta_A)}$ when $\lambda_P\to\infty$, then $\forall\delta_A\in(0,1)$ and $\epsilon>0$, there exists $\overline{\lambda}_P^c$ such that when $\lambda_P>\overline{\lambda}_P^c$,

$$\int_{s=t_i}^{\infty}\delta_A^{s-t_i}C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds>-\frac{1}{\ln\delta_A}C(1-\epsilon).$$

Moreover, as in the proof of Proposition 2, since

$$\int_{s=t_i}^{\infty}\delta_A^{s-t_i}\mathbb{E}[\Delta v(\sum_j a_{j,s})-\kappa|\sigma_P,\sigma_A,h_{t_i}]ds<-\frac{1}{\ln(\delta_A)}\cdot\big(\lim_{n\to\infty}\Delta v(n)-\kappa\big),$$

it then follows from Assumption 1 that

$$\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i},\vec{a}_P)|\sigma_P,\sigma_A,h_{t_i}]<0$$

when $\lambda_P>\overline{\lambda}_P^c$.

Hence no agent i wants to be the first to choose $a_i=1$ and $\sigma_A(h_{t_i})=0$ with probability 1 is an optimal strategy after any such history h_{t_i}. Applying this reasoning by induction to all $i'>i$ yields that the unique equilibrium involves $a_i^*=0$ for al i. We have thus established Part (I).

Part (II):

It is immediate from the above that here, and in contrast to the proof of Proposition 2, the bound $\overline{\lambda}_P^c$ is independent of λ_A.

To show that $\overline{\lambda}_P^c$ is decreasing in δ_A, first note that as

$$\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i},\vec{a}_P)|\sigma_P,\sigma_A,h_{t_i}]<-\Big(\frac{1}{\ln(\delta_A)}\cdot\big(\lim_{n\to\infty}\Delta v(n)-\kappa\big)-C\int_{\tau_{k_i}=t_i}^{\infty}\frac{\delta_A^{\tau_{k_i}-t_i}}{\ln(\delta_A)}f(\tau_{k_i})d\tau_{k_i}\Big),$$

a sufficient condition for $\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i},\vec{a}_P)|\sigma_P,\sigma_A,h_{t_i}]<0$ is that

$$\big(\lim_{n\to\infty}\Delta v(n)-\kappa\big)-C\int_{\tau_{k_i}=t_i}^{\infty}\delta_A^{\tau_{k_i}-t_i}f(\tau_{k_i})d\tau_{k_i}<0. \qquad (11)$$

Eq. (11) still holds for $\tilde{\delta}_A>\delta_A$, as $\int_{\tau_{k_i}=t_i}^{\infty}\tilde{\delta}_A^{\tau_{k_i}-t_i}f(\tau_{k_i})d\tau_{k_i}>\int_{\tau_{k_i}=t_i}^{\infty}\delta_A^{\tau_{k_i}-t_i}f(\tau_{k_i})d\tau_{k_i}$. Moreover, by continuity in λ_P, we can say that $\exists\xi>0$ such that for all $\lambda_P'\in(\lambda_P-\xi,\lambda_P]$, we have that $\int_{\tau_{k_i}=t_i}^{\infty}\tilde{\delta}_A^{\tau_{k_i}-t_i}f'(\tau_{k_i})d\tau_{k_i}>\int_{\tau_{k_i}=t_i}^{\infty}\delta_A^{\tau_{k_i}-t_i}f'(\tau_{k_i})d\tau_{k_i}$, where $f'(\tau_{k_i})$ denote the pdf of τ_{k_i} under the intensity parameter λ_P'. It therefore follows that Eq. (11) holds for $\tilde{\delta}_A>\delta_A$ and $\lambda_P'\in(\lambda_P-\xi,\lambda_P]$ and thus that $\overline{\lambda}_P^c$ is decreasing in δ_A. ■

Proof of Proposition 4.

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 $\Delta\pi_i(\vec{a}_{-i}) = v(1, M-1) - \kappa - v(0, M-1) > 0$ 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):

Call $\Omega = \{1, 2, ..., M\}$. Recall that a Principal strategy is defined as $\sigma_P : \{0,1\}^M \to \Delta(2^\Omega)$, where 2^Ω denotes the power set (the set of all subsets a_P of agents), with the restriction that $\sigma_P(a_P) = 0$ when $|a_P| > B$. That is, σ_P is a probability measure that assigns a probability to each (possibly empty) subset a_P of Ω, with subsets of size $|a_P| > B$ necessarily having probability 0 since the principal cannot punish more than B agents.

The Principal will direct punishment only at agents who have chosen action $a = 1$, since she gets no utility from punishing agents who chose $a = 0$. Punishing any selection of agents who have chosen $a = 1$ will give her the same utility. As she cannot condition punishment on an agent's label, she will choose a uniformly random punishment strategy σ_P, by which $\mathbb{P}\{\phi_i|\sigma_P\} = \min(\frac{B}{\sum_j a_j}, 1)$ 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/M (her budget relative to the number of agents) is low enough.

Let $1 \leq B < M$. 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 $\mathbb{P}\{\phi_i|\sigma_P\} = B/M$. There exists $\beta \in (0,1)$ such that, when $B/M < \beta$, then $\Delta v(M-1) > \kappa$ (by Assumption 1) and $\mathbb{P}\{\phi_i|\sigma_P\} = B/M < \beta$ for all i. Thus with the action profile $\vec{a} = \vec{1}$,

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

for all i when $B/M < \beta$ 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/M 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 $\frac{\Delta v(M-1) - \kappa}{C}$.

By a similar argument, there exists $\gamma \in (0,1)$, with $\gamma \geq \beta$, such that when $B/M > \gamma$, then the only equilibrium involves $a_i^* = 0$ for all agents, since then each agent has a large enough probability $\mathbb{P}\{\phi_i|\sigma_P\}$ of being punished. Here, such a γ also corresponds to $\frac{\Delta v(M-1) - \kappa}{C}$. Indeed, in such a case

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

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, $\frac{\Delta v(M-1)-\kappa}{C} < \frac{B}{M}$ implies that $\frac{\Delta v(n-1)-\kappa}{C} < \frac{B}{n}$ for any $n < M$ and thus $\mathbb{E}[\Delta\pi_j(\vec{a}_{-i})|\sigma_P] < 0$ under any agents' actions profile. ■

References

Briffault, R. (2008): "Lobbying and campaign finance: Separate and together," *Stanford Law and Policy Review*, 19, 105.

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

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

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.

Hasen, R. L. (2012): "Lobbying, rent-seeking, and the constitution," *Stanford Law Review*, 191–253.

Johnson, V. R. (2006): "Regulating lobbyists: Law, ethics, and public policy," *Cornell Journal of Law and Public Policy*, 16, 1.

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.

Knowles, J., N. Persico, and P. Todd (2001): "Racial bias in motor vehicle searches: Theory and evidence," *Journal of political economy*, 109, 203–229.

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.

Owens, E. and B. Ba (2021): "The economics of policing and public safety," *Journal of Economic Perspectives*, 35, 3–28.

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

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

6 Supplementary Online Appendix

6.1 Extensions

6.1.1 Effect of the information structure available to the Principal

Suppose now that the actual times t_j at which agents acted are not publicly observed, but their actions are. Thus, the Principal cannot perfectly deduce the order in which agents took action $a = 1$. This fits well with situations like political demonstrations or riots, where the Principal (e.g. the government) can observed which agents (e.g. demonstrators) participated in a demonstration, but not the exact timing at which they joined the protest (and thus the order in which they joined it).

For that purpose, let $\mathcal{A}_k = \{j | a_j = 1 \text{ and } t_j \leq \tau_k\}$ be the set of agents who chose $a_j = 1$ before time τ_k. Here the identities j, j' of such agents reveal no information about the relative timing of their actions (j, j' are just unordered labels or names). Let the Principal observe the information $(\mathcal{A}_k \setminus \mathcal{A}_{k-1}, \tau_k, a_{P,k})$ at each time τ_k when she gets to act. With τ_1 being the first time she gets to act, denote the observed history of play by

$$\tilde{h}_{\tau_k} = \tilde{h}_{\tau_{k-1}} \bigcup (\mathcal{A}_k \setminus \mathcal{A}_{k-1}, \tau_k, a_{P,k}), \tag{12}$$

with $\tilde{h}_{\tau_0} = \emptyset$ and $\mathcal{A}_0 = \emptyset$.

With this information structure, the best the Principal can do is to split the agents who took action $a = 1$ into ordered cohorts of offenders $\mathcal{A}_1$, $\mathcal{A}_2 \setminus \mathcal{A}_1$, ..., reflecting the imperfect information that she has about the timing of the agents' actions.

As the Principal cannot do better than punishing a randomly chosen agent in a given cohort each time τ_k she gets to act, the probability that an agent gets punished will depend not only on the actions chosen by other agents, but also on the intensities λ_A and λ_P governing the action opportunities of the agents and of the Principal. It is important to note that this is true both without *and* with commitment. Thus, the benefits of being able to announce and commit to a punishment strategy—which, as stated in Proposition 3, allowed the Principal to deter collective action irrespectively of the intensity of the agents' activity λ_A—disappear.

If the Principal's policing is fast enough relative to the agents' activity, then it has the same effect as knowing the order in which agents took action $a = 1$, since the ordered cohorts of offending agents contain at most a single agent with arbitrarily high probability. Note again that here the size of the budget B is not important to deter the agents (just like in Proposition 2).

Thus, commitment is useful when information about the timing of the agents' actions allows the Principal to announce and design an ordered punishment strategy (cf. Definition 1). Without such information, commitment loses its advantage and the Principal must rely on his reaction speed to deter agents.

6.1.2 Presence of fearless agents

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 $q \in (0, 1)$. The fearless agents have flow payoff

$$\tilde{\pi}^F_{i,t}(a_{i,t}, \vec{a}_{-i,t}, \vec{a}_{P,t}) = v(a_{i,t}, \sum_j a_{j,t}) - \kappa \cdot a_{i,t} \tag{13}$$

and thus do not fear punishment, just like in a pure coordination game. The rational players have the payoff function as in Eq. (1), as before.

If the intensity λ_A of the agents' activity is sufficiently high and if there is a sufficiently high fraction $q > \underline{q}(\lambda_A)$ of all agents who are fearless (with $\underline{q}(\lambda_A)$ decreasing in λ_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 agent always choose action $a = 1$.

Suppose an agent's type θ_i is private and thus not publicly observable. Then, when agents' action times are perfectly observable and if the Principal can announce and commit to a punishment strategy, she can choose an ordered punishment strategy as in Definition 1. 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. 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, irrespectively 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. $B > N_F$) 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 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 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 $B > 0$.

The same insights apply when the Principal cannot publicly announce and commit to a punishment strategy. Indeed, then if θ_i is private, she would choose a strategy that punishes all offenders, irrespectively 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 $\lambda_A - \lambda_P$ is sufficiently negative), then she could choose a strategy by which she punishes the first $B - 1$ agents, since she enjoys punishing offenders (recall that $\epsilon \in (0, 1)$ is positive in her payoff function (cf. Eq. (3))). Perpetually keeping a budget of size $B = 1$ thereafter would then be optimal, since she would credibly dissuade the rational agents from taking action $a = 1$—as she enjoys punishing offenders—and she maximizes her expected payoff by minimizing 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 all rational agents 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. 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.

6.2 Microfoundation for parameter ϵ

To give a microfoundation for the parameter ϵ in Eq.(3), which makes the principal enjoy punishing a faulty agent, we could assume that, with some probability ξ, the agent's action is "cancelled" or overturned upon punishment. For instance, if the lobby (or interest group) is not only able to make the policymaker suffer a reputational cost (i.e. C) for posting something on social media, but also to have the post removed from the social media platform. For that purpose, we could slightly rewrite the flow payoff in Eq. (3) as $\tilde{\pi}_{P,t}(\vec{a}_{P,t}, \vec{a}_t) = -\sum_j a_{j,t} + \sum_j \mathbb{1}_{\overline{\psi}_{j,t}} \cdot a_{j,t}$, where $\overline{\psi}_{j,t} = \{\exists \tau_k \leq t : \{a_{P,k} = j\} \bigcap \{a_{P,s} \neq j, \forall \tau_s \leq t \text{ such that } \tau_s \neq \tau_k\} \bigcap \{a_j \text{ cancelled}\}\}$ is now the event ψ_j *together with* the agent's action being cancelled. We could then slightly rewrite an agent's flow payoff (Eq. (1)) with the first term replaced by $v(\tilde{a}_{i,t}, \sum_j \tilde{a}_{j,t})$, where $\tilde{a}_{k,t} = a_{k,t} \cdot (1 - \mathbb{1}_{\omega_{k,t}})$ with $\omega_{k,t} = \{\phi_{k,t} \bigcap \{a_k \text{ cancelled}\}\}$ being the event that agent k is punished *and* his action is cancelled. The results presented throughout the paper would be unaffected by such a modification. We thus kept the payoffs as defined in equations Eq. (1) and Eq. (3), as they are notationally simpler.

6.3 Further comments on Proposition 3 and the case with commitment

The Principal can achieve the same outcome as in Proposition 3 if she only recalls the identity of the first offender. To see this, let $\underline{h}_t = h_{\min(t_l, t)}$, where t_l is the first time an agent chooses $a = 1$. We call $\underline{h}_t$ the truncated history of play at time t, by this we mean the history that contains only the information of h_t up to and including the first time an agent chooses $a = 1$. This is stated in the following corollary.

Corollary 1 (Equilibrium when only the identity of the first offender is observable) *Let the Principal observe only the truncated history of play $\underline{h}_t$ at any time t. Moreover, let her commit to a strategy σ_P^* under which she punishes only the first agent who chose action $a = 1$. For* any *budget size $B > 0$, if $\lambda_P > \overline{\lambda}_P^c$, then the equilibrium involves $a_i^* = 0$ for all i.*

Proof of Corollary 1. Let σ_P^* be such that the Principal punishes only the first agent who chose action $a = 1$. That is $\sigma_P(h_{\tau_1}) = i$, where i is such that $a_i = 1$ for $t_i < \tau_1$ and $a_j = 0$ for all $t_j < t_i$. Then, by the same argument as in the proof of Proposition 3, no agent ever takes action $a = 1$ since no agent wants to be the first to take it. It follows that the equilibrium involves $a_i = 0$ for all i. ■

It is interesting to note that in Proposition 2 (the non-commitment case), if the lobby reacts quickly enough (i.e. if $\lambda_A - \lambda_P$ is low enough), then it has the same effect as being able to commit to a strategy in which it promises to punish offenders in the order in which they took action $a = 1$. Indeed, with arbitrarily high probability, *at most one* agent will have had the chance to act between any two times τ_{k-1} and τ_k at which the Principal gets to take policing actions. Then any strategy by which she punishes, at time τ_k, *some* offending agent who acted in the previous time interval will correspond, with high probability, to punishing agents in the order in which they took action $a = 1$.

Finally, also note that when the Principal has commitment ability, we are free to set $\epsilon = 0$ in Eq. (3). In other words, she no longer even needs to enjoy punishing a faulty agent: the term $\sum_j \epsilon \mathbb{1}_{\psi_{j,t}} \cdot a_{j,t}$ is no longer needed and her flow payoff at time t can be simply written as $\tilde{\pi}_{P,t}(\vec{a}_{P,t}, \vec{a}_t) = -\sum_j a_{j,t}$. While in the case without commitment, the only way for punishment to be credible was for the Principal to enjoy some utility from it, here her ability to credibly commit to a strategy removes any such need.