On the Power of Reaction Time in Deterring Collective Actions*

Philippe Jehiel† Mathieu V. Leduc‡

March 31, 2025

Abstract

Multiple agents play a dynamic coordination game, each having an opportunity to act at a separate random time. A principal with a conflicting interest can, at random times, take costly punitive actions to dissuade the agents from acting adversely. When the Principal is too slow in her reaction time, the agents successfully coordinate, but when she is sufficiently quick relative to the agents, we show natural conditions under which she can deter the agents from coordinating irrespective of her budget. This brings new perspectives on policing and on the political economy of social movements.

Keywords: Dynamic games, coordination games, equilibrium selection, reaction time, deterrence
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, Leeat Yariv, Andrew Koh, Tom Palfrey and David Myatt.
†Email: philippe.jehiel@gmail.com
‡(Corresponding author) Email: mattvleduc@gmail.com.

1 Introduction

When attempting to build a political movement, civil society activists 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 activists have manifested themselves—or are expected to do so in the near future—does it become worth expressing support for the movement, leading to a dynamic coordination problem.

In addition to this dynamic coordination problem, the activists in the previous examples (henceforth, the *agents*) may also face opposition from a *Principal* with diverging interests. This Principal may represent a government, a police force, or an opposing lobby. The Principal may try to dissuade the agents from coordinating by taking retaliatory actions against those who express support for the movement.

It would seem natural that such a Principal would need considerable resources to contain the social movement, allowing it to take retaliatory actions against a potentially large number of people. Indeed, preventing civil society activists from organizing could require extensive policing resources, the ability to conduct a large number of lawsuits, etc.

In this article, we examine this claim and will argue that, in our environment in which actions are taken sequentially at different times, a Principal's ability to react quickly to the agents' actions is a key determinant of the effectiveness of the Principal's activity, more so than its ability to take retaliatory action against a large number of agents.

Specifically, we study a dynamic model where agents with shared interests are each given an opportunity, at a random time, to take a binary action (i.e. expressing support for a political movement 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 movement is increasing in the number of agents who have (and who will) express support for it, thus defining a dynamic coordination game among the agents. On the other side, a Principal with interests diverging from those of the agents is also given random opportunities to take retaliatory actions against the agents who engaged in the adversarial activity. 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, possibly small number of agents.

We first show that in the absence of the Principal (or when the Principal is slow enough), the ability of the agents to dynamically coordinate depends on the arrival rate of their opportunities to express their support for the movement. 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, agents all choose to express support for the movement. This effectively selects a unique equilibrium behavior on the part of the agents.

We next assume that the Principal can react quickly to the agents' actions. We provide an upper bound on how many agents can engage in the adversarial activity in any equilibrium. This upper bound depends on the budget of the Principal and the severity of the punishment felt by an agent, as we discuss later. Interestingly, we note that in any equilibrium and irrespective of the budget of the Principal, no agent engages in the adversarial activity when (i) the Principal can react quickly enough and (ii) sufficiently many agents need to coordinate on the adversarial action for it to become worthwhile. 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.

The intuition for our results is as follows. When the agents are much quicker than the Principal, the effective cost of being punished in case of contributing to the collective action is not very high. This is because, by the time the Principal is given an opportunity to punish, many agents would already have contributed to the collective action (and thus the punishment gets diluted among the many offenders). By contrast, when the Principal can react quickly to the agents' actions, the agents choosing to contribute to the collective action can be punished more effectively. In fact, in the limit, each agent could face the threat of a prompt individual punishment. The remaining key non-trivial consideration is whether the Principal is indeed willing to punish when given an opportunity to, and thus whether her punishment threat is credible. In fact, since she could deplete her limited budget, punishing an agent 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 punishment, then the Principal would find punishment optimal (and it would hence be credible in the eyes of the agents). This is how we derive our upper bound on how many agents can choose the offending action in any equilibrium. However, when this upper bound is smaller than the number of agents required to make this action worthwhile, no agent chooses the offending action in any equilibrium.

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 expressing support for the movement. This simultaneous action game can model a demonstration or riot (as opposed to asynchronous expressions of support, as is the case of social media) and this interestingly shows that deterring such public demonstrations can require more resources from the Principal than for deterring activism on social media.

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). More specifically, we derive conditions under which the Principal can dispense with explicit commitment devices to deter collective adversarial actions.

Our result that without the Principal, agents 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 have no counterpart in the literature as far as we know.[1] We believe our results are particularly well suited to the understanding of activism in the age of social media technologies. Indeed, the latter have had the effect of increasing the speed at which opportunities arrive for agents (i.e. activists) and for the Principal (i.e. an adversarial lobby, a government, etc.). To the extent that a Principal is better at handling social media technologies (and thus can react quickly when the agents express their views), our analysis suggests that a Principal does not need a large budget in order to be effective.

From a technical viewpoint, our modeling of stochastic decision times is somehow 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. However, to our knowledge, that literature has not considered the kind of games involving both a Principal and agents as in our setting.[2] 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

[1] Note that 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.

[2] 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.

(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 dynamic collective action games (see Battaglini and Palfrey (2024a), Battaglini and Palfrey (2024b), Koh et al. (2024) for recent references which, like the other papers in that literature, do not have the analog of a Principal as in our setting) and 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 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 conclude and discuss some extensions and robustness checks. Detailed proofs are relegated to Section 5.

2 Model

2.1 Setting

We study a setting where at random times governed by a Poisson process, different infinitely-lived 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. 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.[4] 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 police force, an opposing lobby, or a government) with interests diverging from those of the agents also gets the chance to police the population of agents, but at some random times. The Principal initially holds a budget $B_0 \in \mathbb{N}^+$. She can act at random times governed by a Poisson process. At each time, the Principal can act and 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, 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 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 k-th event of the Principal's Poisson point process. At each such time, as long as the budget permits it, i.e., as long as $B_{\tau_k} \geq 1$, the Principal can choose an action $a_{P,k} \in \{0,1\}$ where $a_{P,k} = 1$ means that the Principal uses one bullet to punish one of those agents who previously chose $a_i = 1$ and $a_{P,k} = 0$ means that the Principal does not take punitive action.[5] When $a_{P,k} = 1$, the budget is reduced by one unit, i.e. $B_{\tau_k^+} = B_{\tau_k} - 1$ and one of the agents who chose $a_i = 1$ is picked uniformly at random to be punished.

Let the *full 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 $\sigma_P : \mathcal{H} \to \Delta(\{0,1\})$. That is, at some action time τ, based on an observed history of past play h_τ, the Principal can choose to punish

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

[4]This implies that the number of agents who get a chance to play goes to infinity as $t \to \infty$.

[5]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.

(i.e. $a_P = 1$) if $B_\tau \geq 1$. She can alternatively decide not to take punitive action (i.e. $a_P = 0$). She can also choose a mixed strategy.

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 $\phi_{i,t}$ is the event that agent i has been punished (possibly multiple times) by time t, 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.[6] 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 (i.e., at least $N - 1$) 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, irrespective 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.

[6] We could alternatively make the punishment cost be proportional to the number of times the agent has been punished, but this would make no change in the analysis.

2.2.2 Principal's payoff

The Principal's flow payoff at time t is

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

where $\epsilon \in (0,1)$ and $\psi_{j,t} = \sum_{k:\tau_k \leq t} \tilde{\phi}_{j,\tau_k}$ is the number of times agent j has been punished by time t.

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 ϵ from having punished agents who had chosen action $a_j = 1$.[7] $\tilde{\phi}_{j,\tau_k}$ is then a Bernouilli random variable such that if the principal chooses to take punitive action at time τ_k (i.e. $a_{P,k} = 1$), and if agent j has chosen action $a_j = 1$ at some time $t_j \leq \tau_k$, then[8]

$$\tilde{\phi}_{j,\tau_k} = \begin{cases} 1 & \text{w.p. } \frac{1}{\sum_l a_{l,\tau_k}} \\ 0 & \text{w.p. } \frac{\sum_l a_{l,\tau_k} - 1}{\sum_l a_{l,\tau_k}}. \end{cases}$$

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.

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

We will analyze the Subgame Perfect Nash Equilibria of the game described in Section 2. The shape of the equilibria will depend crucially on the reaction time of the Principal to the agents' actions. This is the subject of the following subsections.

3.1 Successful agent coordination

The dynamic nature of the game allows us 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 3.3, the dynamics can allow us to select a unique equilibrium behavior on the part of the agents.

[7] 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.

[8] Our random punishment formulation can be viewed as formalizing a refinement of SPNE based on symmetry and the Markovian payoff-relevance restriction (to the extent that all past offenders affect the Principal's payoff in the same way). Alternatively, our random punishment formulation can be viewed as formalizing an imperfect information assumption that the Principal, when called to play, would only observe those who chose $a = 1$ (as opposed to $a = 0$) but not the identity of the corresponding agents. We will mention in the discussion section how the analysis would be affected if the Principal could observe the identity of the agents.

To gain some intuition into Proposition 1, note that if the Principal'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 by the Principal

We next consider the case in which the intensity (i.e., λ_P) of the Principal's activity is high enough *relative* to the agents' (i.e., λ_A) or, in other words, when the Principal can *react* quickly enough to the actions of the agents.

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

Intuitively, one might think that when the Principal reacts quickly enough, she would always be able to deter agents from choosing $a = 1$ by punishing an offender when given an opportunity to do so. Such a Principal's strategy would deter agents from choosing $a = 1$, because if a first agent were to choose $a = 1$, the Principal would then likely be given an opportunity to punish him before another agent can choose $a = 1$.[9]

The missing aspect in this intuition is whether the Principal would always find it optimal[10] to choose to punish an agent who chose $a = 1$. As it turns out, it cannot *always* be in the interest of the Principal to punish the offenders in equilibrium, and this invalidates the deterrence argument just suggested.

To see this concretely, assume that the Principal's punishment budget allows her to punish only one agent (i.e. $B_0 = 1$). 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 (as long as her budget permits). Then, in the subgame where one agent has chosen $a = 1$ and the Principal is given an opportunity to act, it is *not* optimal for the Principal to punish when C satisfies the condition $\lim_{n\to\infty} \Delta v(n) - \kappa < \frac{C}{2}$, which holds for large enough C.

Indeed, if the Principal chooses to punish, her punishment budget would be completely depleted (i.e. $B_1 = 0$) and all subsequent agents would then choose $a = 1$ (by Proposition 1).

By contrast, if the Principal chose not to punish at that time, 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 quickly punished with probability $\frac{1}{2}$.[11] However, the condition $\lim_{n\to\infty} \Delta v(n) - \kappa < \frac{C}{2}$ implies that this second agent (and any subsequent agent) would strictly prefer not to choose $a = 1$.

[9] The condition $\lim_{n\to\infty} \Delta v(n) - \kappa < C$ in Assumption 1 indeed guarantees that it is suboptimal for an agent to choose $a = 1$ when his punishment probability is high.

[10] If as considered here the Principal cannot ex ante to commit to such a punishment strategy. We consider later the case where the Principal has the ability to commit.

[11] Indeed, when the Principal reacts quickly enough, the punishment would almost surely take place before a third agent has an opportunity 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 a single agent has chosen $a = 1$.

Thus, even if the Principal can react very quickly to the agents' actions, the threat of punishment may not always be *credible* in the eyes of the agents, and thus it may not prevent them from choosing the adversarial action $a = 1$.

The next Proposition takes into account the credibility constraint of the punishment threat and it establishes an upperbound on the number of agents who can possibly choose $a = 1$ in any equilibrium.

Proposition 2 (Partial deterrence) *Let $B_0 \geq 1$ and assume that $\frac{C}{m} < \lim_{n\to\infty} \Delta v(n) - \kappa < \frac{C}{m-1}$, for some $m \in \mathbb{N}_+$. There exist $\underline{\lambda}_A > 0$ (or $\underline{\delta}_A \in (0,1)$) and $\underline{\eta}$ such that if $\lambda_A > \underline{\lambda}_A$ (or $\delta_A > \underline{\delta}_A$) and $\lambda_A - \lambda_P < \underline{\eta}$, then in any equilibrium there are at most $max(0, m - 1 - B_0)$ agents choosing $a_i^* = 1$.*

The intuition for Proposition 2 can be understood as follows. Let $B_0 = 1$ to start with and assume that m satisfies the conditions of the Proposition. When $m - 1$ agents have already chosen $a = 1$, the next agent to act will choose $a = 1$, since even assuming there will be a quick punishment, when $\frac{C}{m} < \lim_{n\to\infty} \Delta v(n) - \kappa$, $a = 1$ is better than $a = 0$ for the agent under the (rational) expectation that all subsequent agents will choose $a = 1$. Consider next the case in which $m-2$ agents have chosen $a = 1$. If the next agent chooses $a = 1$, he should know that the Principal will choose to punish next, since as just established, no matter what the Principal decides then, all subsequent agents will choose $a = 1$ and the Principal has a preference for punishing as early as possible those who picked $a = 1$. Punishment is hence credible. Since $\lim_{n\to\infty} \Delta v(n) - \kappa < \frac{C}{m-1}$ and the Principal reacts quickly (so that she will most likely act before another agent is given an opportunity to move), the agent will optimally choose $a = 0$. This establishes that when $B_0 = 1$, there can be at most $m - 2$ agents who choose $a = 1$ in any equilibrium. It is then not difficult to establish by backward induction that a larger budget $B_0 > 1$ allows the Principal to deter an additional $B_0 - 1$ agents from choosing $a = 1$, thereby establishing Proposition 2.

3.2.2 On the non-monotonic role of the punishment cost C felt by agents

An interesting aspect suggested by Proposition 2 is that a larger punishment cost C felt by the agent is not necessarily more dissuasive and does not necessarily lead to a smaller number of agents choosing $a = 1$ in equilibrium.[12]

While it might seem that a larger C is always beneficial to the Principal, a closer inspection suggests otherwise. As a stark illustration, consider the threat by a country to use a nuclear weapon and whether it would actually be used after another country has committed a small offense. If this other country anticipates that this weapon will not be used, then it would clearly commit the small offense in the first place. Likewise, the former country may not want to deplete its limited arsenal, as it knows this arsenal will be dissuasive in the future.[13]

As it turns out, the magnitude of the punishment cost C can be viewed as playing a dual role. On the one hand, the punishment cost C serves 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 C is more dissuasive and thus beneficial to the Principal. On the other hand, a small enough C 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 C is then

[12] To see this more explicitly, consider a case in which $N = 1$ (say because $\kappa = 0$) and $B_0 = 1$. When $\frac{C}{2} < \lim_{n\to\infty} \Delta v(n) - \kappa < C$, we know by Proposition 2 that no agent will choose $a = 1$ in equilibrium. However, when $\frac{C}{m} < \lim_{n\to\infty} \Delta v(n) - \kappa < \frac{C}{m-1}$ for some $m > 2$, it is easily verified that one can construct an equilibrium in which the first $m - 2$ agents choose $a = 1$. Thus, a larger C is not always good.

[13] Such deterrence considerations have been central in Schelling (1960) and have led practitioners to develop devices to make the use of nuclear weapons credible in some circumstances.

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 2 puts these two considerations together by relating the punishment cost C to the number of agents who can possibly choose $a = 1$ in equilibrium.

3.2.3 Successful deterrence of all agents, irrespective of the Principal's budget B_0

Proposition 2 suggests a potential role for the size of the Principal's budget B_0 in deterring agents from choosing $a = 1$. But as we shall now see, in some natural cases the role of the budget turns out to be insignificant.

More precisely, when the number of agents required for coordination to be worthwhile (absent the punishment cost) is large enough (compared to the m introduced in Proposition 2), we can strengthen Proposition 2 by establishing that, in any equilibrium, no agent chooses $a = 1$, irrespective of the size of the Principal's budget.

Proposition 3 (Threshold criterion for full deterrence) *Let* $B_0 \geq 1$ *and assume that* $\frac{C}{m} < \lim_{n\to\infty} \Delta v(n) - \kappa < \frac{C}{m-1}$ *for some* $m < N + 2$*. There exist* $\underline{\lambda}_A > 0$ *(or* $\underline{\delta}_A \in (0, 1)$*) and* $\underline{\eta}$ *such that if* $\lambda_A > \underline{\lambda}_A$ *(or* $\delta_A > \underline{\delta}_A$*) and* $\lambda_A - \lambda_P < \underline{\eta}$*, then in any equilibrium all agents choose* $a_i^* = 0$*.*

Under the conditions of Proposition 3, 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 (i) the minimum number of agents needed to make coordination worthwhile is large enough (i.e. $N > m - 2$), and (ii) the Principal can react quickly enough[14] to the actions of the agents (i.e. $\lambda_A - \lambda_P < \underline{\eta}$).

This result is easy to understand. By Proposition 2, we know that (as long as $B_0 \geq 1$) 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 $a = 1$. Clearly, keeping fixed the punishment cost C and the maximal marginal benefit of choosing $a = 1$ (as captured by $\lim_{n\to\infty} \Delta v(n) - \kappa$), these conditions are met when N, as defined in Assumption 1, is large enough. This thus establishes that when the threshold required to make the collective action worthwhile is large enough, all agents are deterred from choosing $a = 1$ in any equilibrium and irrespective of the Principal's budget size B_0.

An application to social media activism lends itself well to this setting with asynchronous actions. 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. A Principal (i.e. an opposing lobby or a government) can then succeed in preventing the movement from forming simply by reacting and punishing quickly enough, even if it does not have the ability to actually punish many activists. This conclusion is especially true if the movement would require a large number of people to join it so that it starts becoming attractive. This shows how brittle social movements can be.

3.3 Contrast 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. While the asynchronous actions game lent itself well to social media activism, a simultaneous-actions game can better model a 'classical' demonstration, where agents can choose to join in a mass protest or riot at some physical location and run the risk of arrest.

Consider M agents, each of whom can choose an action $a \in \{0, 1\}$ at time 0. The Principal then observes the action profile $\vec{a}$ and, at time 1, chooses whether to punish a *set* of agents, i.e. her action

[14] 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 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.

$a_P \subset \emptyset \bigcup \{1, ..., M\}$ 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).[15] 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

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

As in previous sections, we will restrict our attention to Principal strategies that target agents anonymously (i.e. irrespective of an agent's label i). We thus call $\sigma_P : \{0, 1\}^M \to \overline{\Delta}(2^{\{1,2,...,M\}})$ the Principal's punishment strategy, which is a mapping from a time-0 agents' actions profile $\vec{a}$, which she observes, to a probability measure over anonymously-chosen subsets of agents.[16]

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 $\min(\frac{B}{\sum_j a_j}, 1)$. 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 $B/M > \frac{\Delta v(M-1)-\kappa}{C}$.

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 an initial budget $B_0 \geq 1$ was enough to deter any $a_i = 1$ in equilibrium under the conditions of Proposition 3.

We see that public demonstrations still have a purpose in the age of social media, as they allow the expected punishment to be diluted among the protesters. In this sense, it has a somewhat similar effect as if the Principal did not react quickly in the dynamic version of the game. It also requires the Principal

[15] Note that here we let $B_0 = B$ and we dispense with the time subscript, as the principal acts only once.

[16] This means that any two sets a_P and a'_P of the same size, i.e. such that $|a_P| = |a'_P|$, have the same chance of being chosen, i.e. $\sigma_P(a_P|\vec{a}) = \sigma_P(a'_P|\vec{a})$, thus guaranteeing anonymity. Here, $\sigma_P(a_P|\vec{a})$ denotes the probability of choosing the set of agents a_P, given the agents' action profile $\vec{a}$ under strategy σ_P. Note that this construction also implies that a given agent cannot be punished more than once.

to have a large policing budget (i.e. a budget that scales with the size of the population of agents) in order to effectively deter the agents.

4 Discussion and extensions

Finally, we discuss some extensions.

4.1 On the effect of personalized punishments

4.1.1 Modified setting

In the above, 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.[17]

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.

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 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 the event[18] that agent j is punished at time τ is defined as $\tilde{\phi}_{j,\tau} = 1$ with probability 1 if $a_{P,\tau} = j$, and $\tilde{\phi}_{j,\tau} = 0$ with probability 1 if $a_{P,\tau} \neq j$. The payoffs are otherwise defined as in Section 2.2.

4.1.2 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 $B_0 = 1$ 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

[17]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.

[18]See Section 2.2 for how $\tilde{\phi}_{j,\tau}$ was defined in main setting of the paper.

agents (i.e., $N = 1$ in Assumption 1). Then punishing the m-th agent who chooses $a = 1$, with the first $m - 1$ 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. $\lambda_P > \overline{\lambda}_P^c$, for some $\overline{\lambda}_P^c > 0$), or that the agents do not discount the future too much (i.e. $\delta_A \in (\underline{\delta}_A^c, 1)$ for some $\underline{\delta}_A^c \in (0, 1)$). 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 magnitude of $\lambda_P - \lambda_A$) is now irrelevant. This is in sharp contrast with the analysis in the main model.

This is formalized in the following definition and proposition.

Definition 1 (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_{τ_k} be a history where i is such that $a_i = 1$ and $a_j = 0$ for all $t_j < t_i$, and $\tau_{k-1} < t_i < \tau_k$. Then $\sigma_P(i|h_{\tau_k}) = 1$, where $\sigma_P(i|h_{\tau_k})$ denotes the probability of selecting agent i for punishment after history h_{τ_k}.*

Proposition 5 (Equilibrium with commitment and personalized punishment) *Let the Principal commit to a first offender punishment strategy σ_P^*. There exists $\overline{\lambda}_P^c > 0$ such that for* any *budget size $B_0 > 0$, if $\lambda_P > \overline{\lambda}_P^c$, then the equilibrium involves agents choosing $a_i^* = 0$ for all i. Moreover, $\overline{\lambda}_P^c$ does not depend on λ_A.*[19]

In Proposition 5, 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.

Again, the threat of a single bullet (i.e. any $B_0 > 0$) can discipline an entire population, but the Principal does not even need to react quickly anymore.

4.2 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

$$\tilde{\pi}_{i,t}^F(a_{i,t}, \vec{a}_{-i,t}, \vec{a}_{P,t}) = v(a_{i,t}, \sum_j a_{j,t}) - \kappa \cdot a_{i,t} \tag{7}$$

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, if the Principal can announce and commit to a personalized punishment strategy, she can choose a first offender punishment strategy as in Section 4.1 (Definition 1). She will then have to punish the agent who chose $a = 1$ first.

[19] Equivalently, there exists $\underline{\delta}_A^c \in (0, 1)$ such that the same equilibrium outcome holds when $\delta_A \in (\underline{\delta}_A^c, 1)$.

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.[20] 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 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 3 are satisfied, and that the Principal can condition his 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 $\lambda_A - \lambda_P$ is sufficiently negative), then she could choose a strategy by which she punishes the first $B_0 - 1$ 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 $B_0 = 1$ 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_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

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

[20] Indeed, going forward, this second agent becomes the next 'first agent' to choose action $a = 1$. A slight variation on Definition 1 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.

i's expected marginal payoff from choosing $a_i = 1$ as opposed to $a_i = 0$ can be written as

$$\begin{aligned}\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i},\vec{a}_P)|\sigma_P,\sigma_A,h_{t_i}] &= \int_{s=t_i}^{\infty}\delta_A^{s-t_i}\Big(\mathbb{E}[\Delta v(\sum_j a_{j,s})-\kappa|\sigma_P,\sigma_A,h_{t_i}]-\mathbb{E}[C\mathbb{1}_{\phi_{i,s}}|a_i=1,\sigma_P,\sigma_A,h_{t_i}]\Big)ds\\ &= \int_{s=t_i}^{\infty}\delta_A^{s-t_i}\Big(\mathbb{E}[\Delta v(\sum_j a_{j,s})-\kappa|\sigma_P,\sigma_A,h_{t_i}]-C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}\Big)ds \quad (8)\end{aligned}$$

where the expectation on the right-hand side is taken over a_j and t_j.

Proof of Proposition 1.

Note that if the last term in Eq. (8), 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

$$\begin{aligned}\int_{s=t_i}^{\infty}\delta_A^{s-t_i}C\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds &= C\int_{s=t_i}^{T}\delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds+\\ &\qquad C\int_{s=T}^{\infty}\delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds. \quad (9)\end{aligned}$$

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 right-hand side of Eq. (9) can be made arbitrarily small.

The second term on the right-hand side of Eq. (9) can also be made arbitrarily small when T gets large. Indeed,

$$\begin{aligned}C\int_{s=T}^{\infty}\delta_A^{s-t_i}\mathbb{P}\{\phi_{i,s}|a_i=1,\sigma_P,\sigma_A,h_{t_i}\}ds &\leq C\int_{s=T}^{\infty}\delta_A^{s-t_i}ds\\ &= C\frac{\delta_A^{s-t_i}}{\ln\delta_A}\Big|_{s=T}^{\infty}\\ &= 0-C\frac{\delta_A^{T-t_i}}{\ln\delta_A}\\ &= K(T)\\ &> 0, \quad (10)\end{aligned}$$

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 guarantees 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.

Consider first the case where $B_0 = 1$.

The fact that λ_P is large relative to λ_A ensures that between two consecutive agents' actions, the Principal has an opportunity to move with high probability. Suppose then that, at time t_{i-1}, $m-2$ agents have already chosen $a = 1$ and the Principal has not yet used her budget, i.e $B_{t_{i-1}} = 1$. If an additional agent (i.e. agent $i-1$) chooses $a_{i-1} = 1$, then the Principal will choose $a_{P,k} = 1$ as soon as she has an opportunity to move (i.e. at the first $\tau_k > t_{i-1}$). This is so because otherwise, if the Principal does not punish, all subsequent agents will choose $a = 1$ as $\frac{C}{m} < \lim_{n\to\infty} \Delta v(n) - \kappa$ (and 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 t_i such that $t_{i-1} < \tau_k < t_i < \tau_{k+1}$. Using Eq. (8), and when λ_A (or δ_A) is large enough, this agent i's expected marginal payoff from choosing action $a_i = 1$ can be expressed as

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

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

The second inequality is established by the following argument: Since $\lim_{n\to\infty} \Delta v(n) - \kappa > \frac{C}{m}$, 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 $\frac{C}{m}$. Indeed, under the correct expectation that all future agents choose $a = 1$, then for all $s \geq t_i$ and $\epsilon > 0$, there exists λ'_A such that when $\lambda_A > \lambda'_A$, then $\left(\lim_{n\to\infty} \Delta v(n) - \kappa\right) - \mathbb{E}[\Delta v(\sum_j a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{t_i}] < \epsilon$. It thus follows that there exists $\underline{\lambda}_A > 0$ such that $\int_{s=t_i}^{\infty} \delta_A^{s-t_i}\left(\mathbb{E}[\Delta v(\sum_j a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{t_i}] - \frac{C}{m}\right)ds > 0$ when $\lambda_A > \underline{\lambda}_A$, under the correct expectation that all future agents will choose $a = 1$.[21]

Thus, since $\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{t_i}] > 0$, 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. $\epsilon > 0$ in Eq. (3)). She prefers doing it at the earliest opportunity (i.e. at time t_k) 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 punishment and $\lim_{n\to\infty} \Delta v(n) - \kappa < \frac{C}{m-1}$. 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 $t_{i-1} < \tau_k < t_i < \tau_{k+1}$. Using Eq. (8), and when λ_P is large enough relative to λ_A, this agent i's expected marginal payoff from choosing action $a_i = 1$ can be expressed as

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

where the first inequality follows from $\mathbb{E}[\Delta v(\sum_j a_{j,s})|\sigma_P, \sigma_A, h_{t_i}] \leq \lim_{n\to\infty} \Delta v(n)$. To establish the second inequality, first note that, by an argument analogous to the one stated earlier, $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} \leq \frac{1}{m-1}$. Moreover, $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} \to \frac{1}{m-1}$ as λ_P increases, for any given fixed λ_A. This implies that $\forall\epsilon > 0$, $\exists\underline{\eta}$ such that when $\lambda_A - \lambda_P < \underline{\eta}$, then $\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} > \frac{1}{m-1} - \epsilon$. Since, $\lim_{n\to\infty} \Delta v(n) - \kappa - \frac{C}{m-1} < 0$, this in turn implies that $\lim_{n\to\infty} \Delta v(n) - \kappa - C\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\} < 0$ when $\lambda_A - \lambda_P < \underline{\eta}$, thus establishing the second inequality above.

Therefore, the Principal reacting quickly enough to an agent's action (i.e. $\lambda_A - \lambda_P < \underline{\eta}$) ensures that $\mathbb{E}[\Delta\pi_{i,t_i}(\vec{a}_{-i}, \vec{a}_P)|\sigma_P, \sigma_A, h_{t_i}] < 0$, and thus that the $(m-1)$-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 $max(0, m - 1 - B_0)$ 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. ■

Proof of Proposition 3. By Proposition 2, 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 $N - 1$ 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 Proposition 4.

Part (I):

In the absence of the Principal, $a_i^* = 0, \forall i$, is a pure strategy equilibrium. Indeed, $\Delta\pi_i(\vec{a}_{-i}) =$

[21] A similar argument can be made to show that $\int_{s=t_i}^{\infty} \delta_A^{s-t_i}\mathbb{E}[\Delta v(\sum_j a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{t_i}]ds$ converges to $\int_{s=t_i}^{\infty} \delta_A^{s-t_i}\left(\lim_{n\to\infty} \Delta v(n) - \kappa\right)ds$ as δ_A increases, and thus that the same conclusion holds for δ_A greater than some $\underline{\delta}_A \in (0, 1)$.

$v(1,0) - \kappa - v(0,0) < 0$ 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 $\sigma_P : \{0,1\}^M \to \overline{\Delta}(2^\Omega)$, where 2^Ω denotes the power set (the set of all subsets a_P of agents), with the restriction that $\sigma_P(a_P|\vec{a}) = 0$ when $|a_P| > B$, and that $\sigma_P(a_P|\vec{a}) = \sigma_P(a'_P|\vec{a})$ whenever $|a_P| = |a'_P|$. That is, $\sigma_P(\cdot|\vec{a})$ 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, and with agents being anonymously chosen (irrespective of their labels).

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. ■

Proof of Proposition 5.

Let σ_P be such that the Principal uses a first offender punishment strategy.

Given some agent i's action time t_i, the Principal gets a first chance to punish agent i at some random time $\tau = t_i + w$, where $w \sim exp(1/\lambda_P)$. Given any history h_{t_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. (8) as

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

where $f(\tau_{k_i})$ denotes the probability density function of τ_{k_i}. 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 a first offender 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 $\lambda_P > \overline{\lambda}_P^c$,

$$\int_{s=t_i}^{\infty} \delta_A^{s-t_i} C\mathbb{P}\{\phi_{i,s}|a_i = 1, \sigma_P, \sigma_A, h_{t_i}\}ds > -\frac{1}{\ln\delta_A}C(1-\epsilon). \tag{11}$$

Moreover, since

$$\begin{aligned}\int_{s=t_i}^{\infty} \delta_A^{s-t_i}\mathbb{E}[\Delta v(\sum_j a_{j,s}) - \kappa|\sigma_P, \sigma_A, h_{t_i}]ds &< \int_{s=t_i}^{\infty} \delta_A^{s-t_i}\big(\lim_{n\to\infty}\Delta v(n) - \kappa\big)ds \\ &= -\frac{1}{\ln(\delta_A)}\cdot\big(\lim_{n\to\infty}\Delta v(n) - \kappa\big)\end{aligned} \tag{12}$$

and $\lim_{n\to\infty}\Delta v(n) - \kappa < C$ by Assumption 1, it then follows from Eqs. (11) and (12) that

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

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

Hence no agent i wants to be the first to choose $a_i = 1$. Denoting by $\sigma_A(0|h_{t_i})$ the probability of choosing action $a = 0$ after history h_{t_i} under strategy σ_A, we conclude that $\sigma_A(0|h_{t_i}) = 1$ 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 all i.

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. ■

References

Battaglini, M. and T. R. Palfrey (2024a): "Dynamic collective action and the power of large numbers," Tech. rep., National Bureau of Economic Research.

——— (2024b): "Organizing for Collective Action: Olson Revisited," *Journal of Political Economy*, 132, 2881–2936.

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.

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

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.

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.

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