

Spontaneous Collective Action*

Daniel Diermeier and Jan A. Van Mieghem

August 14, 2000

Abstract

We propose a dynamic, probabilistic model of collective action. While agents' incentives are approximately captured by the normal form suggested in Palfrey and Rosenthal (1984), agents' behavior is only boundedly rational. The model defines a discrete time, discrete state Markov process. We identify the process' unique limiting distribution and show that in the long-run even in large populations mass protests may occur as rare, sudden events of comparatively short duration. Moreover, mass collective action is possible in the absence of any coordination device. We then show that our model can be used to give a formal "as-if" interpretation of game-theoretic analyses of collective action.

1 Introduction

The political transformation of Eastern Europe in the Fall of 1989 was initiated by a series of mass demonstrations. A well-studied example are the 1989 Leipzig Monday Demonstrations that triggered the collapse of the German Democratic Republic and led to a reunified Germany (Lohmann 1994, Opp, Voss, and Gern 1995). Not only was there no apparent leadership, political organization, public event, or even previous experience that would lead to coordinated mass behavior,¹ but the protests also occurred to the utmost surprise of both political rulers and observers. These were *spontaneous* revolutions (Opp, Voss, and Gern 1995). Nevertheless, the demonstrators were able not only to sustain waves of mass demonstrations, but also to coordinate on a strict time-table:

Monday, 5 PM, prayer for peace, followed at about 6 PM by gathering at the Karl-Marx-Platz, and finally a march around the city on the Ring to the new city hall, break-up, and return home. (Opp, Voss, and Gern 1995; p.22)

The goal of this paper is to provide a formal model of collective action that can capture these and related instances of collective action. It is intended to account for the following empirical regularities²:

1. Mass political protest usually occurs in waves. For example, participation in political demonstrations is not evenly spread across time, but tightly clustered. (e.g. Koopmans 1993).
2. In many countries mass demonstrations, especially violent protests, are rare events. Long periods of political calm may suddenly be interrupted by outbreaks of political protest of comparatively short duration (e.g. Koopmans 1993, Opp, Voss, and Gern 1995). That is, any model that explains why e.g. revolutions are possible should also explain why they occur so rarely.
3. Political protests are frequently not foreseen. Perhaps the most famous example are the Eastern European revolutions of 1989 (e.g. Francisco, 1993, Kuran 1991, Lohmann 1994, Opp, Voss, and Gern 1995). But similar instances can be found in Western democracies such as the protests in the UK over the poll tax that led to the fall of then Prime Minister Margaret Thatcher.

*Department of Managerial Economics and Decision Sciences (MEDS), Kellogg Graduate School of Management, Northwestern University, Evanston, IL 60208-2009; d-diermeier@kellogg.nwu.edu and VanMieghem@kellogg.nwu.edu

¹The last political demonstrations opposed to the SED government occurred in 1953.

²In our discussion we mainly focus on demonstrations and mass protest. Our model, however, is also applicable to other examples of collective action like rallies, boycotts, wildcat strikes, and lynch mobs etc. For an overview see e.g. Chong (1991).

4. Many instances of mass collective action occur (at least initially) in the absence of political leadership or explicit coordination. Examples include the 1989 Leipzig Monday Demonstrations (Lohmann 1994, Opp, Voss, and Gern 1995) and the origins of the U.S. Civil Rights Movement (e.g. Chong 1991).

To account for these phenomena we build on two influential formal approaches in the study of collective action.³ The first is Palfrey and Rosenthal's (1984) game-theoretic analysis of collective action. The key idea of this approach is to interpret collective action as a discrete public goods problem.⁴ For example, demonstrators may call for free elections, the resignation of a political leader, a change in policy, or the termination of a controversial construction project like an airport. If (and only if) a sufficient number of citizens demonstrate, will these demands be granted. More radical demands are associated with higher thresholds.⁵ Protesters all collectively benefit from a change, but incur individual costs. These costs capture not only spent time and resources, but also the risks of protest such as physical injury, loss of life, arrest or prosecution. Palfrey and Rosenthal demonstrate that mass collective action can occur even if the protesters' costs are substantial.

The strength of game-theoretic models is to capture the subtle strategic incentives in mass collective action. However, the game theoretic approach also faces some challenges. First, the Palfrey-Rosenthal game has many equilibria, some with a protest level of zero. Game-theoretic analysis, however, only specifies which outcomes are possible. It does not indicate which one is more likely. Some researchers have argued that the "non-participation" equilibrium is focal, because it constitutes a natural base-line (Parikh and Cameron 1999). But such an argument not only seems circular (it uses focality to explain predominant occurrence, and predominate occurrence to explain focality), but it also fails to explain why demonstrations ever stop once they have occurred. Second, for protests to occur, agents must be able to solve a complex coordination problem (especially in large populations) with no apparent coordination device.⁶ An analysis of solutions to these coordination problems is usually beyond the realm of traditional game-theoretic analysis. Nevertheless it plays an important role in informal accounts of revolutions and other forms of mass political behavior (e.g. Kuran 1991, Opp, Voss, and Gern 1995). Third, game-theoretic models cannot explain the occurrence of short-lived protest waves. In the game-theoretic formalism, these phenomena correspond to an unexplained switch from one equilibrium to another and back. Fourth, many authors have rejected game-theoretic models of mass political behavior like collective action or elections on the grounds that these approaches make unrealistic assumptions about the rationality of actors and the amount of common knowledge that must be present among the agents (e.g. Green and Shapiro 1996). A discussion of these issues is beyond the scope of this paper, but an account of the strategic incentives of mass collective action that relies on less stringent assumptions about the rationality and knowledge of actors is certainly desirable.

A related, but distinct tradition in the study of collective action is usually referred to as "threshold" or "critical mass" models (Schelling 1978, Granovetter 1978, Kuran 1991, Oliver and Marwell 1988). Individuals in a population are assumed to vary in their willingness to participate in collective action. These variations may stem from differences in costs and benefits (Oliver and Marwell 1988), or directly as propensities to act as a function of the number of others who are already acting (Granovetter 1978). Collective action will occur only if there is a sufficiently large critical mass of agents who are willing to take the first step and thus trigger mass participation. Whether collective action occurs thus depends on the distribution of individual participation thresholds in the populations. For example, contrary to Olson's (1965) seminal analysis of collective action large groups should *not* be expected to be less successful in generating collective action than small ones, since large groups may be more likely to contain a critical mass of participants (Oliver and Marwell 1988).

In contrast to game-theoretic approaches critical mass and threshold models nicely capture the dynamic nature of collective action. However, as pointed out by Lohmann (1994), they imply a dynamic path of mass protests that is either stagnant or monotonically increasing. They do not account for protest waves or sudden outbreaks of short-lived mass demonstrations. Moreover, threshold models are not well-suited to

³For a comprehensive overview of formal models of collective action see Oliver (1993).

⁴See Lohmann (1994) for an informational model of collective action. In Lohmann's approach costly political actions transmit information. So, information plays the role of the public good.

⁵For an example of political thresholds in practice consider Francisco's (1993; table 1) summary of the 1989 events in Czechoslovakia. At 1% cumulative protest level the political arrests stopped, at 11% minor leaders began to resign, at 61% the party surrendered power.

⁶Political leaders may try to overcome this problem, but then followers face a second-order coordination problem on whether to follow the leader. Also, a political leadership frequently develops *after* initial protests have occurred.

capture the strategic complexities of collective action indicated by game-theory. Finally, while there have been some informal attempts to explicitly model the implicit adjustment processes (e.g. Schelling 1978), a rigorous treatment of the dynamics of collective actions is still lacking.

In this paper we develop a dynamic, behavioral model of mass political protests to integrate these approaches. In particular, while agents' incentives are captured by normal forms, agents are not assumed to have common knowledge about the game-form, to calculate optimal long-run strategies, or have perfect foresight. Rather, their behavior exhibits myopia, lack of information, and inertia. That does not mean that they are entirely devoid of rationality. In fact, as in the game-theoretic model, their behavior does respond to incentives, just not in the perfect, friction-less, and time-independent manner envisioned by non-cooperative game-theory. This approach not only allows us to formally capture the informal dynamics suggested by threshold models. It also explains why mass protests occur rarely; why they occur suddenly if they occur; and why they are unlikely to last long.⁷

The model is able to capture both dynamic effects such as "critical masses" and the strategic richness of game-theoretic models. In our model agents' incentives are given by the same normal form as in the Palfrey-Rosenthal model. Moreover, our approach provides a rigorous formulation of the informal claims (often uttered by defenders of rational choice models) that the assumptions in rational choice models of politics should not be interpreted literally, since they are best understood as approximations or interpreted "as-if".⁸ Once the methodological link to game-theoretic models is established, we can show that our predictions correspond to a strict subset of the Palfrey-Rosenthal equilibria as our model approaches the rational choice assumption. Moreover, in the limit, our model selects a *unique* Nash equilibrium as the predicted outcome. We further characterize the conditions under which this equilibrium is a protest or no-protest equilibrium.

2 The Participation Game

We consider a fixed population of size $N \geq 2$ with d denoting a generic agent (a potential demonstrator). Agents can choose one of two actions $z \in Z = \{0, 1\}$, where $z = 0$ means "not participating." We assume the same normal form as Palfrey and Rosenthal (1984)⁹: If at least k agents participate (with $0 < k \leq N$), each player receives a (collective) benefit of size 1; otherwise the collective benefit is not provided. In addition there is a private cost c of participating independent of whether the collective benefit is provided. Throughout the analysis we assume $0 < c < 1$. Let X denote the number of agents participating; similarly, let X_{-d} denote that number excluding agent d . Since an agent's payoff depends only on his action and on the number of other players participating we can write an agent d 's pay-off as $u(z; X_{-d})$. Agent d 's pay-offs can be summarized in the following matrix:

Payoffs $u(z; X_{-d})$	$X_{-d} < k - 1$	$X_{-d} = k - 1$	$X_{-d} \geq k$
$z = 0$	0	0	1
$z = 1$	$-c$	$1 - c$	$1 - c$

The game has an abundance of equilibria. If $k = 1$ there are N pure strategy equilibria, each with exactly one participant.¹⁰ If $k > 1$, there are $\binom{N}{k}$ pure strategy equilibria (each with exactly k contributors), and one pure strategy equilibrium where nobody participates.

In addition, there are equilibria where some agents use mixed strategies. These agents must be indifferent between c and their pivot probability, i.e., the probability that their participation will lead to the provision of the collective good. Palfrey and Rosenthal show that as $N \rightarrow \infty$ mixed strategy equilibria disappear. That is, in large populations either the collective good is provided for sure or not at all. We are thus left with a severe equilibrium multiplicity problem. The game-theoretic analysis predicts either universal participation or no

⁷The focus of our model is on mass behavior, not decision making by elites. Action by political rulers (such as modifying punishments, granting small concessions etc.) are only incorporated to the extent that they can be captured in the parameters of the behavioral model. See DeNardo (1985) for a model that also incorporates elite decision making.

⁸For a recent example of this debate see the critique of rational choice models Green and Shapiro (1994) and the responses by rational choice theorists collected in Friedman (1996).

⁹Palfrey and Rosenthal also consider other specifications of public good problems. See their paper for precise definitions of the normal form game in terms of strategies, pay-offs etc.

¹⁰The case of $k = 1$ corresponds to the well-known "chicken game".

participation at all. Note that these two types of equilibria exist for all $k > 1$ and $0 < c < 1$. Thus, in large populations, we cannot formally capture the intuition that higher costs or higher thresholds may reduce participation. Moreover, if $k < N$, all equilibria where the collective good is provided are asymmetric¹¹. That is, although the game is symmetric in pay-offs and actions, the predicted behavior is not: some agents participate while others free-ride. This leaves us with a puzzle: how do large populations like the citizens of Leipzig manage to overcome a stark coordination problem, especially if there is no apparent coordination device, like previous experience or political leadership? How is spontaneous collective action possible?

3 A Probabilistic Model

To explicitly analyze coordination in large populations we present a stochastic, dynamic model of the participation game.¹² This approach differs from standard game-theory in two respects: (a) the behavioral assumptions, and (b) the equilibrium concept. However, it shares with game-theory the use of the normal form to capture the relevant incentives.

In contrast to standard game-theoretic models, the model does not assume common knowledge of the game form or perfect foresight by voters. Rather, agents adjust their actions according to some behavioral rule. We may then think about this process as follows: In each period t one specific agent out of N is randomly chosen with probability $1/N$.¹³ The agent then looks at the current configuration X^t of actions in the population and adjusts his action according to a given behavioral rule.¹⁴ The next period, again a player is chosen at random, and so on. Given the current configuration, an actor will then probabilistically adjust his participation behavior to improve his pay-off.

Let $p^\beta(z|X_{-d}^t)$ denote the conditional probability that in period $t+1$ agent d will play action z given that the current configuration of play is X^t . Specifically, we assume the following behavioral rule for all $d \in N$:

$$p^\beta(z|X_{-d}^t) = \frac{\exp[\beta u(z; X_{-d}^t)]}{\sum_{z' \in Z} \exp[\beta u(z'; X_{-d}^t)]}, \quad (1)$$

which is equivalent to the familiar log-linear choice rule¹⁵. It captures the assumption that the pair-wise probability ratios of choosing actions are proportional to the respective pay-off differences. This rule can either be interpreted as “perturbed” decision making (e.g. Blume 1993) or as a random utility model (e.g. McFadden 1973). In the latter interpretation, rather than specifying that agents have fixed incentives, utilities are assumed to vary randomly according to a given probability distribution with a fixed mean.¹⁶ Given these incentives agents choose optimal actions. This interpretation is particularly suitable for a model of collective action, since the (perceived) benefits and costs of participating, e.g. in a demonstration, may well vary substantially over time. Note, however, that nothing in our model presupposes a particular interpretation of the log-linear rule. All we assume is that the agents’ behavioral regularities can be captured by it.

In both interpretations agents are not assumed to be as rational and knowledgeable as in game-theoretic models. Rather, individuals exhibit elements of boundedly rational behavior. Agents do respond to incentives, but not perfectly. For example, they optimize conditional on the current behavior in the population without anticipating the future strategic consequences of their actions. Agents need not believe that other actors reason in the same way as they do, or that they have the same pay-off function. Indeed, they need

¹¹See Myerson (1997) for a critique of asymmetric equilibria in the context of large turnout games. Myerson points out that for asymmetric equilibria to occur, the identity (“the name”) of the agents must be common knowledge.

¹²There is a large literature on the use of stochastic models in economics (see Blume (1997), Fudenberg and Levine (1998) or Young (1998) for detailed overviews) and sociology (e.g. Oliver 1993), but little work in political science (see, however, Bendor, Diermeier and Ting 2000). Our approach is most closely related to the models proposed by Young (1993) and especially Blume (1993).

¹³For simplicity, we assume that revisions are made each period. All results, however, continue to hold in continuous time when the time between revisions is exponentially distributed.

¹⁴Alternatively, we may assume that the actor sees an unbiased random sample of the population’s current configuration.

¹⁵This rule has a long history in psychology and economics (Block and Marschak 1960) and has recently also been used in adaptive models (Blume 1993). It is also used in many empirical studies of mass political behavior (the “logit model”).

¹⁶See McFadden (1973) for details.

not expect that their action may influence the future decisions of other participants. Agents simply adopt the action that maximizes their current pay-off.

Despite these apparent differences, the log-linear choice model is closely connected to the familiar best-response correspondence. The parameter β formally captures the degree to which participation decisions are captured by the assumed pay-off matrix. A low β corresponds to the case where a participation decision is not much influenced by the incentives specified in the model. For $\beta = 0$ choice is completely random. That is, for all possible configurations, d will play each action with probability $1/2$. For $\beta \rightarrow \infty$, log-linear choice converges to a distribution that puts positive probability only on best-responses to X_{-d}^t .

The model can now be summarized as follows. In each period one agent is randomly selected to change his behavior. That agent's action then is drawn from a log-linear behavioral rule given the current configuration of play. The realization of that action then determines the next period's configuration of play; again an agent is chosen and so forth.

Our stochastic model defines a discrete time, discrete state Markov process (or Markov "chain"). Formally, we have a family of random variables $\{X^t : t \in \mathbb{N}\}$ where X^t assumes values on the state space $S = \{0, 1, 2, \dots, N\}$. The value of X^t is updated at the beginning of each period t , such that, given the value of X^t , the values of X^s for $s > t$ do not depend on the values of X^u for $u < t$. The probability of X^{t+1} being in state j (that is, $X^{t+1} = j$) given that X^t is in state i is called the transition probability P_{ij}^t . In our model, these transition probabilities are fully specified by the log-linear choice rule and the selection process. Since both stochastic components are independent of the time variable t , we have a Markov chain with stationary transition probabilities, denoted by the transition matrix P . A Markov process is completely defined once its transition matrix P and initial state X^0 (or, more generally, the initial probability distribution over X^0) are specified.

A Markov chain with transition matrix P is said to be *regular* if for some m the matrix P^m has only strictly positive elements. The following two conditions are jointly sufficient for regularity (Taylor and Karlin 1994; p.171):

1. For every pair of states i and j there is a path l_1, \dots, l_r for which $P_{il_1}P_{l_1l_2} \cdots P_{l_rj} > 0$.
2. There is at least one state i for which $P_{ii} > 0$.

The most important fact concerning a finite, regular Markov chain is the existence of a unique limiting distribution, denoted by the column vector π , where

$$\pi_j = \lim_{t \rightarrow \infty} \Pr\{X^t = j | X^0 = i\}, \quad (2)$$

and $\pi_j > 0$ for all $j \in S$ (Taylor and Karlin 1994). Thus, π_j is the long-run ($t \rightarrow \infty$) probability of finding the process in state j , irrespective of the initial state. A second interpretation of the limiting distribution is that π_j also gives the long-run mean fraction of time that the process is in state j .

It can easily be shown that π is the unique distribution that solves $\pi = \pi P$.¹⁷ These equations are called the *global balance equations* because, rearranging $\pi_i = \sum_j \pi_j P_{ji}$, yields

$$(1 - P_{ii}) \pi_i = \sum_{j \neq i} \pi_j P_{ji}, \quad (3)$$

which can be interpreted as saying that the probability "flow" out of state i must equal the probability flow into state i .

Because at most one individual can change his behavior in any period, X^t can change by at most 1 at a time. That is, we have $P_{ij} = 0$ if $|i - j| > 1$. Such Markov process is called a *birth-death process*. To simplify notation, denote P_{ii+1} by the "birth" probability λ_i (i.e., the probability that the number of participants increases by one) and P_{ii-1} by the "death" probability μ_i (i.e., the probability that the number of participants decreases by one). Hence, $P_{ii} = 1 - \lambda_i - \mu_i$. For a birth-death process, the balance of probability flow satisfies a stronger property:

$$\lambda_{i-1} \pi_{i-1} = \mu_i \pi_i \Leftrightarrow \frac{\pi_i}{\pi_{i-1}} = \frac{\lambda_{i-1}}{\mu_i}. \quad (4)$$

¹⁷ To see this, let $P_{ij}^{(t)} = \Pr\{X^t = j | X^0 = i\}$ denote the " t -step" transition probabilities. We have that $P^{(t+1)} = P^{(t)}P$. Now letting $t \rightarrow \infty$ and using the definition that $\pi_j = \lim_{t \rightarrow \infty} P_{ij}^{(t)}$, yields $\pi = \pi P$.

These equations are called *detailed balance equations*. It is easy to verify that they indeed also solve the global balance equations, which now read

$$(\lambda_n + \mu_n)\pi_n = \lambda_{n-1}\pi_{n-1} + \mu_{n+1}\pi_{n+1}.$$

Since in a birth-death process the limiting probability ratio equals the transition probability ratio, we easily can derive a closed form solution of the limiting distribution in our probabilistic model.

4 Results

To analyze the limiting behavior of the participation model, we must first specify the transition matrix P . Given that only direct-neighbor transitions are possible, we only need to specify the birth and death parameters $\lambda_n = P_{nn+1} = \Pr\{X^t = n+1 | X^t = n\}$ and $\mu_n = P_{nn-1}$. The transition probabilities have two components. First, we have the probability that any one agent is selected to make a decision, which we call the “selection probability.” Second, there is the probability that a given action is chosen, which we call the “action probability.” The probability that any action is taken depends on the current configuration, i.e., the configuration X^t just before the revision time. If actor d did not participate, we characterize him as being of sub-type $(d, 0)$; otherwise he is of sub-type $(d, 1)$. Given that $X^t = n$, the probability that the randomly picked actor d is of a sub-type $(d, 0)$ or $(d, 1)$ is, respectively,

$$p_0(n) = \frac{N-n}{N} \quad \text{and} \quad p_1(n) = \frac{n}{N}.$$

This characterizes the selection probabilities.

Action probabilities are determined by the individual choice rule. It is useful to rewrite our pay-off matrix by sub-type. For example, the second row captures the next period pay-off of an agent that switches from non-participation to participation, conditional on the configuration of play (expressed by the columns).

Payoffs $u(z X)$	$X < k-1$	$X = k-1$	$X = k$	$X > k$
Type $(d, 0)$: $z = 0$	0	0	1	1
Type $(d, 0)$: $z = 1$	$-c$	$1-c$	$1-c$	$1-c$
Type $(d, 1)$: $z = 0$	0	0	0	1
Type $(d, 1)$: $z = 1$	$-c$	$-c$	$1-c$	$1-c$

Given log-logistic choice, actor d selects payoff action z with probability $p^\beta(z|X_{-d}^t)$. This allows us to specify the action probability matrix as:

Action Probabilities	$X < k-1$	$X = k-1$	$X = k$	$X > k$
Type $(d, 0)$: $z = 0$	$\frac{1}{1+e^{-\beta c}}$	$\frac{1}{1+e^{\beta(1-c)}}$	$\frac{e^\beta}{e^\beta + e^{\beta(1-c)}}$	$\frac{e^\beta}{e^\beta + e^{\beta(1-c)}}$
Type $(d, 0)$: $z = 1$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$	$\frac{e^{\beta(1-c)}}{1+e^{\beta(1-c)}}$	$\frac{e^{\beta(1-c)}}{e^\beta + e^{\beta(1-c)}}$	$\frac{e^{\beta(1-c)}}{e^\beta + e^{\beta(1-c)}}$
Type $(d, 1)$: $z = 0$	$\frac{1}{1+e^{-\beta c}}$	$\frac{1}{1+e^{-\beta c}}$	$\frac{1}{1+e^{\beta(1-c)}}$	$\frac{e^\beta}{e^\beta + e^{\beta(1-c)}}$
Type $(d, 1)$: $z = 1$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$	$\frac{e^{\beta(1-c)}}{1+e^{\beta(1-c)}}$	$\frac{e^{\beta(1-c)}}{e^\beta + e^{\beta(1-c)}}$

which simplifies to:

Action Probabilities	$X < k-1$	$X = k-1$	$X = k$	$X > k$
Type $(d, 0)$: $z = 0$	$\frac{1}{1+e^{-\beta c}}$	$\frac{1}{1+e^{\beta(1-c)}}$	$\frac{1}{1+e^{-\beta c}}$	$\frac{1}{1+e^{-\beta c}}$
Type $(d, 0)$: $z = 1$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$	$\frac{e^{\beta(1-c)}}{1+e^{\beta(1-c)}}$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$
Type $(d, 1)$: $z = 0$	$\frac{1}{1+e^{-\beta c}}$	$\frac{1}{1+e^{-\beta c}}$	$\frac{1}{1+e^{\beta(1-c)}}$	$\frac{1}{1+e^{-\beta c}}$
Type $(d, 1)$: $z = 1$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$	$\frac{e^{\beta(1-c)}}{1+e^{\beta(1-c)}}$	$\frac{e^{-\beta c}}{1+e^{-\beta c}}$

The birth probabilities stem from a sub-type $(d, 0)$ changing his action to “participate” ($z = 1$), while death probabilities derive from a demonstrating sub-type $(d, 1)$ changing his action to “not demonstrate”

($z = 0$). We can then calculate the total transition probability by de-conditioning on subtype as:

$$\begin{aligned}\lambda_n &= \begin{cases} \frac{e^{-\beta c}}{1+e^{-\beta c}}p_0(n) & \text{if } n \neq k-1, \\ \frac{e^{\beta(1-c)}}{1+e^{\beta(1-c)}}p_0(n) & \text{if } n = k-1. \end{cases} \\ \mu_n &= \begin{cases} \frac{1}{1+e^{-\beta c}}p_1(n) & \text{if } n \neq k, \\ \frac{1}{1+e^{\beta(1-c)}}p_1(n) & \text{if } n = k. \end{cases}\end{aligned}$$

Notice that our Markov chain is regular. Hence, it has a limiting distribution π that solves the detailed balance equations:

$$\forall n \neq k-1 : \frac{\pi_{n+1}}{\pi_n} = \frac{p_0(n)}{p_1(n+1)}e^{-\beta c} = \frac{N-n}{n+1}e^{-\beta c}. \quad (5)$$

$$\text{for } n = k-1 : \frac{\pi_k}{\pi_{k-1}} = \frac{e^{\beta(1-c)}p_0(k-1)}{p_1(k)} = \frac{N-k+1}{k}e^{\beta(1-c)}. \quad (6)$$

We can solve this recursive system of equations to characterize the limiting distribution. Intuitively, to calculate any π_n we will define an arbitrary reference state, in our case 0, and then “chain” the detailed balance conditions together along a path from 0 to n . This allows us to derive each π_n as a function of π_0 . The probability of the reference state (and thus the probability of every state) can then be derived using the normalization condition $\sum_{n=0}^N \pi_n = 1$.

Proposition 1 *The limiting distribution for the participation model is:*

$$\pi_n = \begin{cases} \binom{N}{n}e^{-\beta n c}\pi_0 & \text{if } n < k, \\ \binom{N}{n}e^{-\beta n c}\pi_0e^\beta & \text{if } n \geq k, \end{cases} \quad (7)$$

where π_0 is a normalization factor with $(1 + e^{-\beta c})^N \leq \pi_0^{-1} \leq e^\beta(1 + e^{-\beta c})^N$ such that $\sum_{n=0}^N \pi_n = 1$.

Proof: From (5), we have that $\forall n \leq k-1$:

$$\pi_n = \left(\prod_{i=0}^{n-1} \frac{p_0(i)}{p_1(i+1)} \right) e^{-n\beta c} \pi_0 = \frac{N(N-1)\dots(N-(n-1))}{1 \cdot 2 \cdot \dots \cdot n} e^{-n\beta c} \pi_0 = \frac{N!}{n!(N-n)!} e^{-n\beta c} \pi_0.$$

From (6) and (4) it follows that:

$$\pi_k = \frac{\lambda_{k-1}}{\mu_k} \pi_{k-1} = \frac{e^{\beta(1-c)}(N-(k-1))}{k} \frac{N!}{(k-1)!(N-(k-1))!} e^{-(k-1)\beta c} \pi_0 = \frac{N!}{k!(N-k)!} e^{-k\beta c} e^\beta \pi_0.$$

Finally, reapplying (5) yields that $\forall n > k$:

$$\begin{aligned}\pi_n &= \left(\prod_{i=k}^{n-1} \frac{p_0(i)}{p_1(i+1)} \right) e^{-(n-k)\beta c} \pi_k = \frac{(N-(n-1))(N-(n-2))\dots(N-k)}{(k+1) \cdot (k+2) \cdot \dots \cdot n} e^{-(n-k)\beta c} \pi_k \\ &= \frac{(N-(n-1))(N-(n-2))\dots(N-k)}{(k+1) \cdot (k+2) \cdot \dots \cdot n} e^{-(n-k)\beta c} \frac{N!}{k!(N-k)!} e^{-k\beta c} e^\beta \pi_0 \\ &= \frac{N!}{n!(N-n)!} e^{-n\beta c} e^\beta \pi_0.\end{aligned}$$

Applying the binomial theorem $\sum_n \binom{N}{n} x^n = (1+x)^N$ directly yields the bounds for π_0 . That is given (5), we have

$$1 = \sum_{n=0}^N \pi_n = \pi_0 \left[\sum_{n=0}^{k-1} \binom{N}{n} e^{-\beta n c} + \sum_{n=k}^N \binom{N}{n} e^{-\beta n c} e^\beta \right].$$

Hence

$$\sum_{n=0}^N \binom{N}{n} e^{-\beta n c} e^{\beta} \geq \pi_0^{-1} \geq \sum_{n=0}^N \binom{N}{n} e^{-\beta n c}.$$

■

Notice that the limiting distribution π_n combines the results of the selection process, as represented by the combinatorials $\binom{N}{n}$, and the results of the action process, represented by $e^{-\beta n c}$ or $e^{\beta} e^{-\beta n c}$. To characterize the long-run behavior of the probabilistic model we now need to identify the maxima of π_n . These are characterized in the next proposition. First, we need a definition:

Definition For any $x \in \mathbb{R}$ define $\lfloor x \rfloor$ as the largest integer with $\lfloor x \rfloor \leq x$ and $\lceil x \rceil$ as the smallest integer with $\lceil x \rceil \geq x$ and let

$$\lfloor x \rfloor := \begin{cases} \lfloor x \rfloor & \text{if } \pi_{\lfloor x \rfloor} \geq \pi_{\lceil x \rceil}, \\ \lceil x \rceil & \text{if } \pi_{\lfloor x \rfloor} \leq \pi_{\lceil x \rceil}. \end{cases}$$

Proposition 2 There exist two critical numbers n^* and k^*

$$n^* = \max \left\{ 0, \frac{N e^{-\beta c} - 1}{1 + e^{-\beta c}} \right\} \quad \text{and} \quad k^* = \frac{N + 1}{1 + e^{-\beta(1-c)}}, \quad (8)$$

with $n^* < \frac{N}{2} < k^*$ such that the following holds:

(i) If $k = 1$, then π_n has a unique maximum at

$$\begin{cases} n = k = 1 & \text{if } \frac{N-1}{2} e^{-\beta c} \leq 1, \\ \lfloor n^* \rfloor > 1 & \text{if } \frac{N-1}{2} e^{-\beta c} > 1. \end{cases}$$

(ii) If $k > 1$ and $k \notin (n^*, k^*)$, then π has a unique maximum at $\lfloor n^* \rfloor$.

(iii) If $k > 1$ and $k \in (n^*, k^*)$, then π has two maxima at $\lfloor n^* \rfloor$ and k , of which k is the most-likely long-run state if

$$\pi_{\lfloor n^* \rfloor} < \pi_k \Leftrightarrow g(k) := (1 - (k - \lfloor n^* \rfloor)c)\beta + \sum_{i=\lfloor n^* \rfloor}^{k-1} \ln \frac{N-i}{i+1} > 0. \quad (9)$$

Otherwise the most likely long-run state is $\lfloor n^* \rfloor$.

Proof: Define $f : [0, N] \rightarrow \mathbb{R} : x \rightarrow f(x) = \frac{N-x}{1+x} e^{-\beta c}$. Note that f is continuous and strictly decreasing over its domain $[0, N]$ with $f(0) = N e^{-\beta c}$ and $f(N) = 0$. From (5), it follows that the odds ratio $\pi_{n+1}/\pi_n = f(n)$ is strictly decreasing in n (with a possible jump at $n = k - 1$). Notice that if n were extended to a continuous variable x , π_x would reach an interior maximum at $x^* \in (0, N)$ where $f(x^*) = 1$ or at $x = 0$ otherwise. If $N e^{-\beta c} > 1$, then f is continuous and monotone decreasing with $f(0) > 1$ and $f(N) = 0$, so that there exists a unique x^* and solving $f(x^*) = 1$ for x^* yields $x^* = \frac{N e^{-\beta c} - 1}{1 + e^{-\beta c}}$. We now must consider the implications of the integer constraints on n and the possible jump at $n = k - 1$.

First consider the case where $k = 1$. For π_n to have a maximum at $n = k = 1$, we need $\pi_1/\pi_0 = N e^{\beta(1-c)} > 1$, which always holds because $N \geq 2$ and $\beta(1-c) \geq 0$, and $\pi_2/\pi_1 = \frac{N-1}{2} e^{-\beta c} \leq 1$, which is also sufficient for a unique maximum at $n = k = 1$ because π_{n+1}/π_n is strictly decreasing in $n \geq 1$. If $\frac{N-1}{2} e^{-\beta c} > 1$, then also $f(0) = N e^{-\beta c} > 1$ so that $n^* := x^*$ and $\lfloor n^* \rfloor$ constitutes the unique maximum for π_n .

Now consider $k > 1$. If $N e^{-\beta c} \leq 1$, then $\pi_1/\pi_0 \leq 1$ so that π reaches a maximum at $\lfloor n^* \rfloor = 0$. If $N e^{-\beta c} > 1$, then as before $\lfloor n^* \rfloor$ constitutes a maximum for π_n . We now need to check for other (possible) maxima, which can only occur around the ‘‘jump’’ at $n = k - 1$, namely at $n = k - 1$ or at $n = k$.

Suppose $k < n^*$. For two maxima we need $k < \lfloor n^* \rfloor$. But since π_{n+1}/π_n is increasing below n^* , we have $\pi_{\lfloor n^* \rfloor}/\pi_k > 1$ so that $n = k$ cannot be a maximum. For $n = k - 1$ to be a maximum, we need

$$\pi_{k-1} > \pi_k \Leftrightarrow \frac{N-k+1}{k} e^{\beta(1-c)} < 1 \Leftrightarrow k > k^*.$$

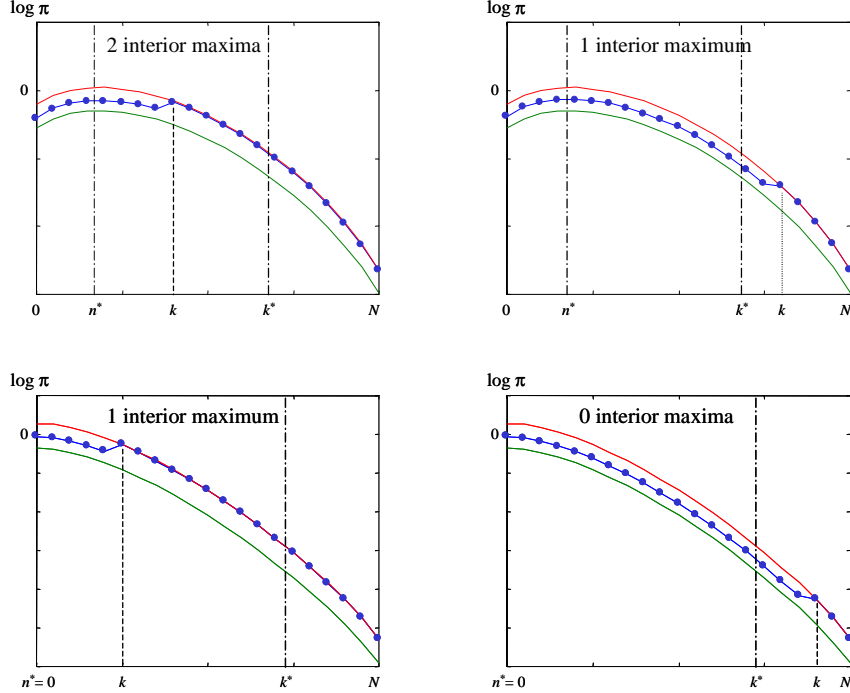


Figure 1: The limiting distribution π assumes one of four possible cases, depending on the parameters n^* , k^* and k .

Notice, however, that $n^* < \frac{N}{2} < k^*$ (because $0 \leq e^{-\beta(1-c)} \leq 1$ and $0 \leq e^{-\beta c} \leq 1$, given that $\beta \geq 0$ and $0 < c < 1$). Therefore, there cannot be a second maximum if $k < n^*$.

Suppose $k > n^*$. If $k - 1 = \lfloor n^* \rfloor = \lfloor n^* \rfloor$ then, since π_{n+1}/π_n is decreasing above n^* , k cannot be a maximum, and since $k - 1 = \lfloor n^* \rfloor$ there cannot be a second maximum. If $k - 1 > \lfloor n^* \rfloor$, then, since π_{n+1}/π_n is decreasing above n^* , there can only be a maximum at k . For a maximum at k we need $\pi_{k-1} < \pi_k \Leftrightarrow k < k^*$.

To characterize the most likely long-run state note that (5) and (6) imply

$$\pi_{\lfloor n^* \rfloor} < \pi_k \Leftrightarrow \frac{k!(N-k)!}{\lfloor n^* \rfloor!(N-\lfloor n^* \rfloor)!} < e^{-\beta((k-\lfloor n^* \rfloor)c-1)}. \quad (10)$$

Condition (9) then follows immediately. ■

The proposition states that the most likely state is either $\lfloor n^* \rfloor$, where any participation is entirely driven by noise, or k , which is the minimal revolt state.

While the integer restriction on n complicates Proposition 1, the basic intuition can be conveyed informally. From Proposition 1, it follows that the limiting distribution π has two components. At $n = k - 1$ the probability distribution π_n “jumps” from one component to the other. It thus suffices to characterize the maxima of the components and then identify possible maxima at the “jump” from $n = k - 1$ to $n = k$. The detailed balance equations (5) immediately imply that the probability ratio π_{n+1}/π_n is strictly decreasing in n . So either, there is a corner solution at $n = 0$ or one interior maximum where the probability ratios are approximately equal to one. Hence, for k smaller than the interior maximum, a maximum would have to be at $k - 1$. But, as we show, in the proof of Proposition 2, in this case the jump is too small. So, there can only be a second maximum at k larger than the interior maximum. The conditions for such a maximum are given by (9). Ignoring the knife-edge case of $k = 1$ we thus have four possible cases displayed in Figure 1.

5 Discussion

Proposition 2 now allows us to derive our model’s predictions concerning mass collective behavior. Note that the qualitative features of the limiting distribution change as a function of the cost c , the threshold k ,




Figure 2: Four cases for the distribution π_n depending on the threshold level k . Other parameters are fixed at $N = 50$, $\beta = 2.5$, and $c = 0.5$.

the responsiveness β and the size of the population N . We need to distinguish three cases:

1. There is one maximum at $[n^*]$, perhaps at 0.
2. There are two (local) maxima, one at $[n^*]$, the other at k , with k the most likely long-run state (global maximum).
3. There are two (local) maxima, one at $[n^*]$, the other at k , with $[n^*]$ the most likely long-run state (global maximum).

To see the effect of changes in k consider an example at $N = 50$, $c = 0.5$, and $\beta = 2.5$, for which $n^* = 10.4$, $k^* = 39.6$ and $[n^*] = 11$. Figure 2 illustrates how the qualitative features of the limiting distribution change in response to changes in k .

At low $k < n^*$ (here $k < 10.4$) there is a unique maximum at $[n^*]$, which thus must be the most likely long-run state. This corresponds to the case with permanent (very) low participation. Any participation is solely driven by randomness at the individual level. For example, using the random utility interpretation, on average there are some individuals that have an incentive to participate on their own. Note that as individual choice approaches best response behavior ($\beta \rightarrow \infty$) n^* converges to 0.

For higher k (here $k = 17$) there are two maxima with k the most likely long-run state. This captures the case of an unstable polity with frequent demonstrations and sustained levels of political protest.

At even higher k ($k = 20$), $[n^*]$ becomes the most likely long-run state, but k is still a local maximum. This case most closely corresponds to the empirical regularities outlined in the introduction. Political protest is possible, but it will be rare and comparatively short-lived.

For very high $k > k^*$ (here $k = 45 > 39.6$), we are back at the case where $[n^*]$ is the most likely long-run state without a local maximum at k .

A similar pattern can be observed for c . For general k and c we characterize maxima and long run states in Figures 3 to 5.




Figure 3: Strategy regions in (k, c) -space for different values of β , with fixed population size $N = 50$.

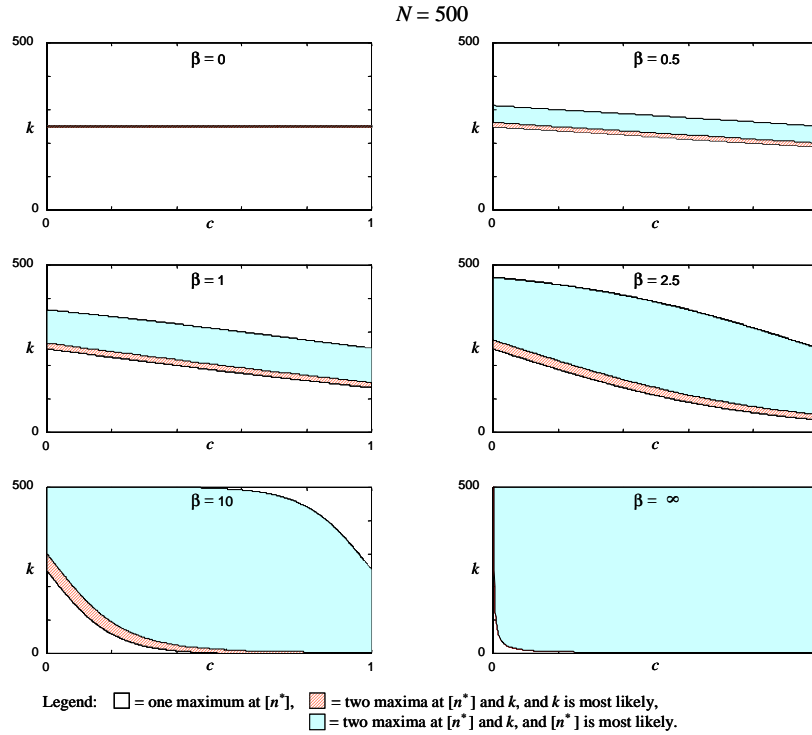


Figure 4: Strategy regions in (k, c) -space for different values of β , with fixed population size $N = 500$.

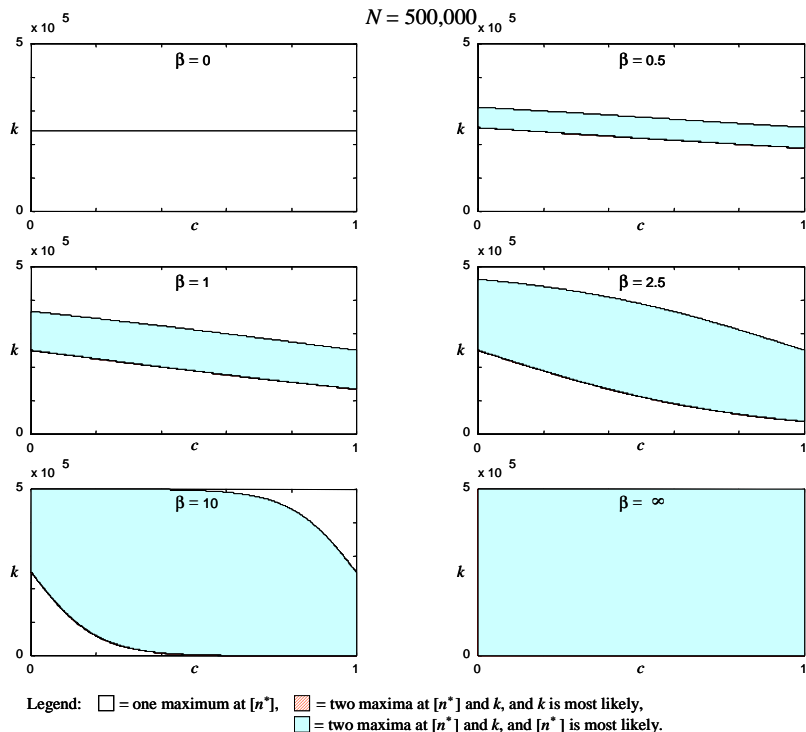


Figure 5: Strategy regions in (k, c) -space for different values of β , with fixed population size $N = 500,000$.

Note that for $\beta \rightarrow 0$, the critical numbers $n^* \rightarrow (N - 1)/2$ and $k^* \rightarrow (N + 1)/2$. Hence, π has a single maximum at $N/2$. In this case individual behavior is not at all governed by the incentives given in the model, it is purely random. This randomness at the individual level corresponds to a collective process with a binomial distribution. As β increases, however, the white areas (unique maximum at $[n^*]$) are shrinking. Even for moderately high β ($\beta = 10$) the largest region is the grey area (global maximum at $[n^*]$, local maximum at k). This effect is present independent of the size of the population N .¹⁸ It becomes, however, more pronounced as N increases. For very large N we virtually only have two regions: If individual randomness is high (low β), we have larger regions with $[n^*]$ as the most likely long-run state, but as individual behavior is better characterized by our normal form, we also have a local maximum at k .¹⁹

The existence of a local maximum at k even for very large N is perhaps the key insight from our model. It implies that at least some times agents are able to spontaneously coordinate on collective action. Note that these states are efficient and asymmetric (i.e. k agents participate, while $n - k$ agents stay home). Nevertheless, mass collective action may occur in the absence of any apparent coordination device.

We can now use our results to account for the empirical phenomena discussed in the introduction. For this purpose it is important to distinguish two alternative interpretations of the limiting distribution. So far our discussion has focussed on the primary interpretation of π_n as the limiting distribution. According to this “cross-sectional” interpretation after the process has been in operation for a long duration, the probability of finding the process in state j is π_j irrespective of the starting state. However, now our goal is to account for time-series phenomena such as protest waves, sudden outbreaks of mass demonstrations etc. This is where the second interpretation is useful. It can easily be shown (e.g. Taylor and Karlin 1994; p.176) that π_j

¹⁸Note that even in the case of $N = 500,000$ there exists a small region where k is the most likely long-run state (case 2), but this region is too small to be picked up by the figure.

¹⁹This result may surprise readers familiar with Olson’s (1965) seminal work on collective action. Olson’s central thesis was that large groups are much less likely than small groups to solve the free-rider problem. Subsequent work, however, has challenged Olson’s thesis (e.g. Marwell and Oliver 1988, Oliver 1993). In her comprehensive survey of the literature Oliver (1993; p.275) concludes: “Put simply, in some situations the group size effect will be negative, in others positive. You have to know the details of a particular situation before you can know how group size will affect the prospects for collective action.”

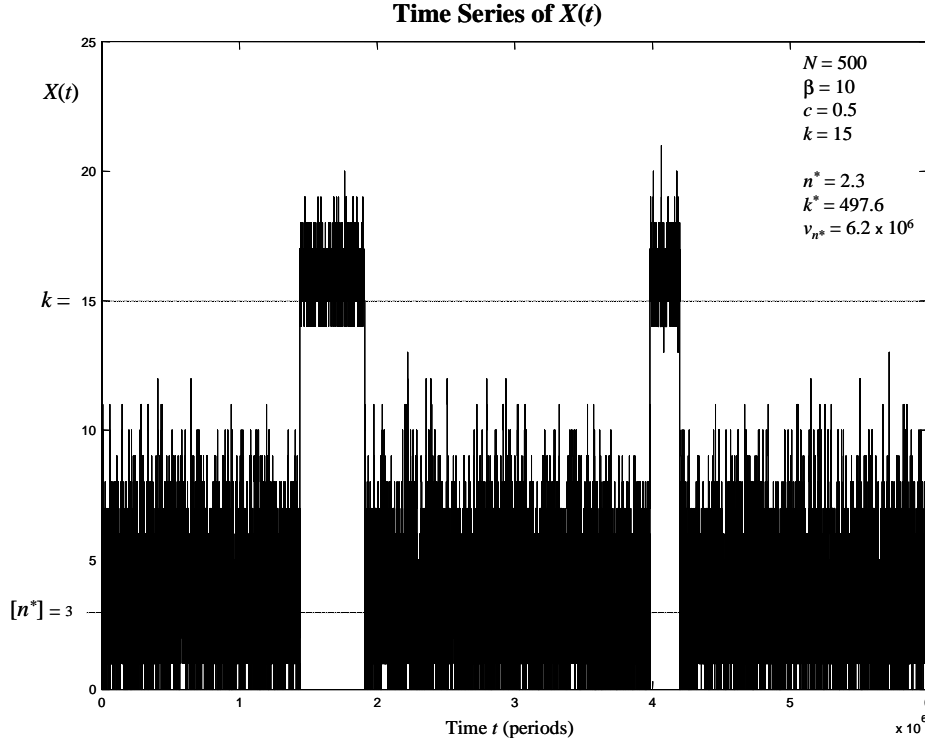


Figure 6: Time series of $X(t)$ for a simulated 6 million periods: X hovers around $[n^*]$ except for two uprisings ($X \geq k$). Parameters: $N = 500$, $\beta = 10$, $c = 0.5$ and $k = 15$.

also gives the long-run mean fraction of time that the process occupies state j . Formally,

$$\pi_j = \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{\tau=0}^{m-1} \Pr\{X^\tau = j | X^0 = i\}$$

Using this interpretation we can show that our model can replicate the typical protest “waves” or “clusters” identified in the empirical literature. These effects are most clearly present in the case of two maxima with n^* as the most likely long-run state (case 3). But as we showed above, even for moderate β this is the predominant case in the (k, c) -space.

Figure 6 shows a typical time series (sample path) of $X(t)$ during 6 million periods.²⁰

To see why these patterns occur it is useful to visualize $-\pi$ as a landscape or a potential with two minima at $[n^*]$ and at k that are separated by a hill. Random noise may push the process over the hill and back. Initially, the population hovers around n^* , but the probabilistic nature of the model may trigger a chain reaction: an initial (random) protester makes it more likely that a second agent switches to $z = 1$; but this may trigger at third etc. Once a critical threshold is reached the process suddenly jumps to the local maximum at k . These chain reactions are reminiscent of the dynamics identified in the critical mass literature. In contrast to these models, however, we can identify not only the likelihood that critical

²⁰Notice that the figure is consistent with theoretical predictions of the mean time between incidents of political protest. That is, the mean time v_i to reach state $n = k$ starting from state $i < k$, solves the linear system:

$$v_i = 1 + \sum_{j=0}^{k-1} P_{ij} v_j \quad \forall i = 0, \dots, k-1.$$

For our example, v_i is on the order of 6.2 million periods for all $i < k$. We simulated several times 10 million periods and indeed never saw more than two uprisings. A reader that is concerned with these large numbers should recall that periods can be arbitrarily small.

masses form, but also the probability that they disintegrate. Hence, our approach can account for both the emergence of collective action and its collapse as an (uniquely determined) *equilibrium phenomenon*. Moreover, since collapse is more likely than emergence, we can account for mass protests as rare, sudden events of comparatively short duration.

It is important to note that collective action, such as mass demonstrations, may occur although the underlying individual incentives *did not change*. This does not mean that external events (such as the events in the Soviet Union under Gorbachev) are not important to facilitate mass protests, only that they are not necessary to do so. Moreover, even if such changes decrease the costs to participate, agents still need to coordinate on protest states. Finally, models that exclusively rely on changes in external parameters also need to explain why protests ever stop again.

From a methodological point of view our approach suggests that it may be preferable not to focus on the prediction or explanation of specific instances of mass collective action, but to choose time-series (preferably of long duration) and cross-sectional samples as the unit of analysis.²¹ According to this approach, any successful model of collective actions should not only be measured in whether it can account for instances of political protests or demonstrations, but also whether it can account for long periods of no such activity. A sole focus on instances of mass collective action would result in a selection bias.

6 An Application - Equilibrium Selection in the Palfrey-Rosenthal Model

As discussed in section 3, the parameter β indicates how closely individual choice behavior approaches best response correspondences. For example, as $\beta \rightarrow \infty$, log-linear choice converges to a distribution that puts positive probability only on best-responses to X_{-d} . This suggests to use our model as a behavioral foundation for the game-theoretic analysis due to Palfrey and Rosenthal. That is, the Nash equilibria can be re-interpreted as approximations of limiting distributions generated by a behavioral model.

This approach²² has been the dominant one in economic applications of Markov learning models (e.g. Blume 1997, Young 1998). Its key purposes are (a) to provide behavioral foundations for game-theoretical solution concepts and (b) to provide a selection criterion among multiple Nash-equilibria in a game. The critical notion here is that of stochastically stable state (Young 1993). A state is *stochastically stable* if its assigned limiting probability is strictly positive as $\beta \rightarrow \infty$. Intuitively, these are states that are most likely to be observed over the long run if perturbations from best response behavior are arbitrarily small. The set of stochastically stable states in many games (but not in general finite games) is a subset of the set of Nash equilibria. If it is a proper subset, stochastic stability can be used as a selection criterion. For example, in 2×2 coordination games the stochastically stable states correspond to the risk-dominant equilibrium (Blume 1993, Young 1993). What makes these selection results remarkable is that they are able to distinguish between *strict* Nash equilibria.²³

As in the economic approaches we can use our analysis to select among the strict Nash-equilibria in Palfrey and Rosenthal's participation game. If $\beta \rightarrow \infty$, then $n^* \rightarrow 0$ and $k^* \rightarrow N$, so that there exist two maxima for large, but finite β , corresponding to either zero turnout or minimal critical turnout k . These maxima thus are analogues to the pure Nash equilibria in the Palfrey and Rosenthal model.²⁴ Note that in the limit of $\beta \rightarrow \infty$, the probabilistic model approaches the best-response model with the noted exception that at most one of the maxima corresponds to a stochastically stable state. This can be interpreted as the selection of one of the pure Nash equilibria in an environment with arbitrarily small (but persistent!) perturbations.

From (9), it follows that the selection depends on the sign of $g(k)$, which, for $\beta \rightarrow \infty$, is positive if $kc < 1$ and negative if $kc > 1$. Hence, the key factor that drives the selection is the sign of $1 - kc$. If $kc < 1$, then the unique long-run prediction is collective action at $n = k$ (almost surely); otherwise, the unique long-run

²¹Such an approach may help resolve the lively debate in political science on whether revolutions can be forecast in principle (e.g. Eckstein 1990, Kuran 1991, Francisco 1993).

²²It was already suggested in Nash's unpublished dissertation! See e.g. Blume (1997).

²³It is possible to identify stochastically stable states without explicitly calculating the limiting distribution π (see Young 1998 for details). However, in this case the model reduces to an equilibrium selection argument and cannot reproduce the dynamic patterns discussed in the introduction.

²⁴See also the discussion of best response dynamics below.

prediction is $n = 0$ (almost surely). Note that the selection does not depend on N . That is, once we control for k the absolute group size plays no explanatory role.

As we demonstrated in Figures 3-5, the case where $kc < 1$ is rare, especially if N is large. Intuitively it captures the case where even if the benefit of unit 1 was private (not public as assumed in our model), it could be redistributed among the minimum k participants needed for a revolt to cover their show-up cost c . Note that in either case the game-theoretic model could not capture the modal empirical case of rare, sudden protests of short duration.

7 Best-Response Dynamics

The reader may ask why we consider the case of $\beta \rightarrow \infty$ rather than analyzing best response dynamics directly. Under best response dynamics each randomly chosen agent plays a best response given the current configuration of play. Doesn't this amount to the same thing as $\beta \rightarrow \infty$? The answer is an emphatic no! A distinguishing feature of probabilistic game models is that their long-run behavior can differ radically from the corresponding best response dynamic even if the stochastic decision rule is arbitrarily close to best response decisions.

To see how this works in our model we derive the best response matrix for the participation game using our pay-off matrix from above:

Best response probabilities	$n < k - 1$	$n = k - 1$	$n = k$	$n > k$
Type $(d, 0)$: $z = 0$	1	0	1	1
Type $(d, 0)$: $z = 1$	0	1	0	0
Type $(d, 1)$: $z = 0$	1	1	0	1
Type $(d, 1)$: $z = 1$	0	0	1	0

Note that all action probabilities are either 0 or 1, so that with respect to actions the dynamic process is deterministic. The stochastic component is driven only by the selection probabilities. So, the Markov chain transition rates are:

$$\lambda_n = \begin{cases} 0 & \text{if } n \neq k - 1, \\ p_0(n) = \frac{N-n}{N} & \text{if } n = k - 1. \end{cases}$$

$$\mu_n = \begin{cases} p_1(n) = \frac{n}{N} & \text{if } n \neq k, \\ 0 & \text{if } n = k. \end{cases}$$

It is easy to see that under best response (provided $k > 1$), the chain is no longer regular, because all states no longer communicate. This follows, because for any n with $0 < n < k - 1$ only lower states are accessible from n . Similarly, if $n > k$, only lower states are accessible. Once such state is left, it is never revisited. To see why note that if $0 < n < k - 1$, the public benefit is not provided and no agent is pivotal. So, participating and non-participating agents have a strict incentive to choose $z = 0$. Hence the number of participants either stays the same (if a non-participating agent is selected) or decreases by one (if a participant is selected). Similarly, for $n > k$ the collective benefit is provided and no agent is pivotal. So, again every selected agent will choose to stay home. It follows that all states n with $0 < n < k - 1$ or $n > k$ are transient.

The states $n = 0$ and $n = k$, on the other hand, are absorbing: once such a state is reached, the system remains in that state forever. To see why note that for $n = 0$, no player is pivotal in providing the collective benefit. Hence, staying home remains optimal for each player. On the other hand consider $n = k$. If a participant is selected, she is pivotal in providing the collective benefit. Since $c < 1$, this makes it worthwhile to participate in the next period. On the other hand, if a non-participant is selected, since the collective benefit is provided no matter whether she participates, free-riding on the participation of the k participants is strictly optimal. So, again no actor has an incentive to change her behavior. The remaining state to consider is $n = k - 1$. Here the collective benefit is not provided and the behavior of the systems now depends on which agent is selected. Since a participant cannot bring about the collective benefit, but still would pay cost c , switching to non-participation is optimal. But then the next state reached is $n = k - 2$ and the system will end up at $n = 0$. On the other hand, any non-participant could now tip the balance if she participated in the next period. Given $c < 1$ this is always strictly optimal. Hence the system will reach the absorbing state

$n = k$. So, starting from $n = k - 1$, the system ends up in $n = k$ with probability $\frac{N-k+1}{N}$ or in $n = 0$ with probability $\frac{k-1}{N}$. Thus, the limiting probability is degenerate: $\lim_{t \rightarrow \infty} X^t$ is either 0 or k , and the selection of 0 or k depends solely on the initial conditions.

Now consider $k = 1$. In that case, $k = n = 1$ again is absorbing, and all $n > k$ are transient. However, now $n = 0$ is not longer absorbing but also transient, since each selected agent can bring about the collective benefit by herself. Hence, the system ends up in state $k = 1$ with probability one. We can summarize these results in the following proposition

Proposition 3 *Under best response dynamics, the limiting distribution (as $t \rightarrow \infty$) for the participation model is:*

$$\begin{aligned}\pi_0 &= \Pr(X^0 \leq k-2) + \frac{k-1}{N} \Pr(X^0 = k-1), \\ \pi_k &= \Pr(X^0 \geq k) + \frac{N-(k-1)}{N} \Pr(X^0 = k-1), \\ \pi_{i \notin \{0,k\}} &= 0,\end{aligned}$$

where $\Pr(X^0 = i)$ is the initial density of X^t at time 0.

Remark 1 *Note that for the special case where $k = 1$ we have $\pi_0 = 0$ and $\pi_1 = 1$.*

The reader will have noticed that Proposition 3 is closely connected to the Palfrey and Rosenthal results. The two absorbing states correspond to the pure strategy equilibria.²⁵ In contrast to the game theoretic model, however, there is no analog to the mixed strategy equilibria, even for small n . Methodologically, the best-response dynamic provides a micro-foundation for how even large groups of political actors may manage to coordinate their behavior even if they lack the strategic foresight required by game-theoretic models. However, in contrast to the probabilistic model presented above the predictions of the model are not unique. Which long-run state is reached depends on the initial state. So, the long-run behavior of the log-logistic model (even after we let $\beta \rightarrow \infty$) differs dramatically from the best response model. Technically, the limits $\lim_{\beta \rightarrow \infty}$ and $\lim_{t \rightarrow \infty}$ do not interchange. Intuitively, the use of some, though arbitrarily small, randomness in individuals' actions is critical. Somewhat surprisingly, assuming randomness in the model improves its predictive power rather than diminish it. Both best response dynamic and game theoretic analysis yield multiple qualitatively different predictions, while the probabilistic model yields a unique prediction, either as a unique limiting distribution or as a selection from the set of pure Nash equilibria. Moreover, once the best response dynamic settles down it will stay at the reached state indefinitely. Best-response models can no longer explain, for instance, the sudden outbreaks of mass demonstrations after long periods of clam, or their comparatively short duration. For these reasons, the log-logistic model seems to be the more promising approach to the study of mass political behavior.

8 Conclusion

We have presented a stochastic model of mass collective action. As in game-theoretic models the agents' incentives are given by a pay-off matrix. However, in contrast to game-theoretic models, actors are boundedly rational and do not share common knowledge of the strategic aspects of the collective decision situation. We argue that such a model is particularly well-suited to the study of mass political behavior such as demonstrations or voting in mass elections. Not only does the model make less demanding assumptions about the actors rationality and knowledge, it also provides a micro-foundation for mass coordination in political environments and yields unique predictions where the game-theoretic analysis predicts multiple, qualitatively different outcomes.

The approach is then applied to Palfrey and Rosenthal's (1984) model of collective action and is solved for the unique limiting distribution. Using its properties we are able to replicate some of the key regularities identified by the empirical literature on political protest: mass protests usually are rare, sudden events of comparatively short duration; they occur in waves or clusters and are frequently unexpected. Moreover, mass collective action is possible in the absence of any coordination device. We finally show that our model can

²⁵This is no coincidence. See Young (1998).

be used to give a formal “as-if” interpretation of game-theoretic collective action models. Once the linkage to game-theoretic models is completed, the model can be used as to select a unique Nash equilibrium.

Our model’s focus is on mass behavior. A natural next step would be to incorporate decision making by (competing) political elites. Other extensions include the study of riots and mass violence, turnout in elections, or voter coordination in multi-candidate elections.

References

- [1] Bendor, Jonathan, Daniel Diermeier, and Michael M. Ting. 2000. “A Behavioral Model of Turnout”. Mimeo. Graduate School of Business. Stanford University.
- [2] Blume, Lawrence E. 1993. “The Statistical Mechanics of Strategic Interaction”. *Games and Economic Behavior* 4:387-424.
- [3] Blume, Lawrence E. 1997. “Population Games”. In W. Brian Arthur, Steven N. Durlauf, and David A. Lane, eds. *The Economy as an Evolving Complex System II*. Reading: Addison-Wesley.
- [4] Chong, Dennis. 1991. *Collective Action and the Civil Rights Movement*. Chicago: University of Chicago Press.
- [5] DeNardo, James. 1985. *Power in Numbers*. Princeton: Princeton University Press.
- [6] Eckstein, Harry. 1990. “More About Applied Political Science.” *PS: Political Science and Politics* 23:54-56.
- [7] Francisco, Ronald A. 1993. “Theories of Protest and the Revolutions of 1989”. *American Journal of Political Science* 37 (August):663-680.
- [8] Friedman, Jeffrey. 1996. *The Rational Choice Controversy*. New Haven: Yale University Press.
- [9] Fudenberg, Drew, and David Levine. 1998. *The Theory of Learning in Games*. Cambridge: MIT Press.
- [10] Granovetter, Mark. 1978. “Threshold Models of Collective Behavior”. *American Journal of Sociology* 83:1420-43.
- [11] Green, Donald, and Ian Shapiro. 1994. *Pathologies of Rational Choice Theory*. New Haven: Yale University Press.
- [12] Koopmans, Ruud. 1993. “The Dynamics of Protest Waves: West Germany, 1965 to 1989.” *American Sociological Review* 58(October):637-658.
- [13] Kuran, Timur 1991. “Now or Never: The element of surprise in the Eastern European revolution of 1989.” *World Politics* 44 (October):7-48.
- [14] Lohmann, Suzanne. 1994. “The Dynamics of Information Cascades: The Monday demonstrations in Leipzig, East Germany, 1989-91.” *World Politics* 47 (October):42-101.
- [15] McFadden, David. 1973. “Conditional Logit Analysis of Qualitative Choice Behavior.” in P. Zarembka, ed., *Frontiers in Econometrics*. New York: Academic Press.
- [16] Myerson, Roger B. 1997. “Population Uncertainty and Poisson Games.” CMSEMS Discussion paper No. 1102R. Northwestern University
- [17] Oliver, Pamela E. 1993. “Formal Models of Collective Action.” *Annual Reviews of Sociology* 19:271-300.
- [18] Oliver, Pamela E., and Gerald Marwell. 1988. “The Paradox of Group Size in Collective Action: A Theory of the Critical Mass II.” *American Sociological Review* 53 (February):1-8.
- [19] Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge: Harvard University Press.

- [20] Opp, Karl-Dieter, Peter Voss, and Christiane Gern. 1995. *Origins of a Spontaneous Revolution: East Germany, 1989*. Ann Arbor: University of Michigan Press..
- [21] Palfrey, Thomas R., and Howard Rosenthal. 1984. "Participation and the Provision of Discrete Public Goods: A Strategic Analysis." *Journal of Public Economics* 24:171-193.
- [22] Parikh, Sunita, and Charles Cameron. 1999. "A Theory of Riots and Mass Political Violence." Mimeo. Department of Political Science. Washington University.
- [23] Schelling, Thomas C. 1978. *Micro-Motives and Macro-Behavior*. New York: W. W. Norton
- [24] Taylor, Howard M., and Samuel Karlin. 1994. *An Introduction to Stochastic Modeling*. Second Edition. Boston et al.:Academic Press
- [25] Young, H. Peyton. 1993. "The Evolution of Conventions". *Econometrica* 61:57-84.
- [26] Young, H. Peyton. 1998. *Individual Strategy and Social Structure*. Princeton: Princeton University Press.