

Voting with your Pocketbook - A Stochastic Model of Consumer Boycotts*

Daniel Diermeier and Jan A. Van Mieghem

This version: March 17, 2008

Abstract

This paper presents a model of consumer boycotts where the discrete choices of concerned consumers are represented as a stochastic processes. Boycotts are interpreted as a form of voting where consumers are trying to shape the behavior of firms. We solve for the limiting distribution of the process and analyze its properties. We then discuss how the model relates to standard game-theoretic approaches to the same phenomenon and show that our model selects one of the many solutions of the corresponding game-theoretic treatment. The type of solution selected depends on the costs and benefits of boycotts to consumers. Specifically, boycotts will occur if and only if they are efficient for consumers.

*Professors Diermeier and Van Mieghem are at the Department of Managerial Economics and Decision Sciences (MEDS), Kellogg School of Management, Northwestern University, Evanston, IL 60208-2009; d-diermeier@kellogg.northwestern.edu and VanMieghem@kellogg.northwestern.edu

1 Introduction

The regulation of economic activity is one of the main arenas of political competition. The impetus for changes to regulatory regimes frequently originates with concerned citizens, often motivated by social or ethical concerns. Examples include areas such as product safety, fair lending standards, or environmentally responsible lending practices. Traditionally, concerned citizens, which we will call "activists" henceforth, have used public institutions such as legislatures, executive agencies, and courts to advance their agenda. Winning majorities in elections for public office then becomes a critical step to success. In recent years, however, many activists have concluded that public processes respond too slowly and can be blocked too easily by special interest. In response they have turned to "private politics" instead. Private politics refers to actions by private interests such as activists that target private agents, often in the institution of public sentiment (Baron and Diermeier 2007; Diermeier 2007). Michael Brune, executive director of the Rainforest Action Network (RAN), a leading global activist group, commented that "Companies were more responsive to public opinion than certain legislatures were. We felt we could create more democracy in the marketplace than in the government." (Baron and Yurday 2004) Democracy in the marketplace means that citizen consumers express in markets their assessment of firm and industry practices.

Consumer boycotts are the most widely used strategic tool of political activists to change corporate practice (Baron 2003, Friedman 1999, Shaw 1996). Yet, they represent a puzzling phenomenon to the modern theory of collective action (e.g. Granovetter 1978, Oliver and Marwell 1988, Oliver 1993, Olson 1965). Consumers are not part of an existing identifiable group or social network, nor do they share a common identity or social activities, all factors that the existing literature has identified as facilitating collective action. In addition, there are no readily available selective incentives, the social benefits of a boycott are not exclusionary, and there usually is no repeated interaction among participants. Thus, since these factors are absent, consumer boycotts should not occur.

Still boycotts do occur, and in many cases they are highly successful. In one of the few quantitative studies of boycott success, Wolman (1914) reports that 72% of the concluded labor-sponsored boycotts at the turn of the century were successful in attaining their stated objective. Indeed they

were so successful that businesses began to devise political and legal strategies to effectively make them illegal. Their actions bore fruit in both Supreme Court decisions and Federal legislation that effectively outlawed “coercive” secondary boycotts in labor disputes (Friedman 1999).

Today, boycotts are the weapon of choice used by political activists with various agendas ranging from environmental concerns, global labor standards, to animal welfare or opposition to genetically modified food products. Boycotts critically rely on the participation of *concerned consumers* who are consumers that also care about the social dimension of a product such as its environmental impact or the way the product is manufactured or marketed. Concerned consumers are an increasingly important segment of the market. They may be willing to pay a higher price for a socially responsible product, or will switch to alternative products if their preferred products are considered socially unacceptable. In the oil industry, concerned consumers are estimated to represent up to 70% of all consumers.¹

To fix ideas consider the famous example of the confrontation between Shell and Greenpeace over the decommissioning of the Brent Spar oil storage facility (e.g. Diermeier 1996, Jordan 2001). In 1991 Shell UK, the British operating company of multinational Royal Dutch/Shell Group, was facing the necessary disposal of the Brent Spar, an aging North Sea oil storage facility and tanker loading buoy. Regulatory guidelines (in this case by the UK. Ministry of Energy and Environmental Affairs) govern petroleum companies in the process of offshore facilities disposal; companies are required to rigorously evaluate disposal options and submit their preference, the Best Practical Environment Option (BPEO), for government approval. Two options survived Shell’s screening process: on-shore dismantling and deep-water disposal. The former requires the transport of the buoy to shore for dismantling while the latter involves towing the structure to a deep-water disposal sight for sinking. Shell UK submitted deep-water disposal as their BPEO concluding that it was both less costly and less likely to result in mishaps that could be dangerous to the environment and the workers. In February of 1995, the British government accepted Shell’s BPEO, deep-water disposal.

Meanwhile, one of the world’s largest environmental groups, Greenpeace International, had become aware of Shell’s plan and had commissioned their own study concluding that removal

¹Presentation by Steve Percy, former President of BP America. Kellogg School of Management. October 28, 2002.

to shore was a better option than deep-water disposal. Greenpeace subsequently acquired satellite communications and video equipment, and on April 30, 2005, 14 activists and 9 journalists boarded the Brent Spar rig. After a three week occupation, the activists were expelled by Shell and local authorities using water cannons, an act that one Greenpeace official, Harold Zindler, characterized as having “portrayed Shell as unresponsive and inconsiderate big business.” In response German motorists engaged in an informal boycott of Shell stations which led to a drop in sales of up to 40%.

On June 20, Shell announced that they would abandon the sinking of the Brent Spar rig. The Chairman stressed that while Shell still believed deep-water disposal to be the best environmental option, Shell UK was in an “untenable position” because of its failure to convince stakeholders in the North Sea. Shell also started an advertising campaign admitting mistakes and promising change, despite a University of London study arguing that deep-sea disposal would likely have been less dangerous to the environment than on-shore dismantling. Shell’s additional costs were estimated at around \$60 million.

Despite their practical importance, boycotts have not attracted much interest among political economists. This is particularly true of formal analyses.² In this paper we focus on the decision problem faced by concerned consumers. That is, we are interested in the dynamics of boycott participation where a large number of consumers needs to take coordinated and costly action.

2 The Basic Model

As our base-line model, we consider the interaction between (potentially) concerned consumers. These consumers have the usual consumption preferences but they also care about the social characteristics of a product. Interaction is modeled as a complete information non-cooperative game with simultaneous moves.

The decision of consumers whether to participate in a boycott can easily be modeled. First

²Recently, however, David Baron has proposed a series of formal models of private politics, i.e. actions by interest groups against private parties such as firms with the goal of changing a firm’s behavior or industry standards (Baron 2002, 2003b, 2003c). His models focus predominantly on the interaction between activists and a firm, and the media, not the dynamics of boycott participation. For a recent model of strategic activism see Baron and Diermeier (2007).

consider only concerned consumers. Assume that concerned Shell costumers need to decide whether to switch their consumption decision to BP in order to force Shell to abandon deep-water disposal of the Brent Spar. On the dimension of the private qualities of the product (quality, price, location of nearest gas station, etc.) these consumers have a preference for buying Shell. That is, if they switch to BP they will pay a private cost c . If the alternative product (here “BP”) is a cheap substitute, c will be low. We also assume that on the social dimension, all concerned customers believe that on-shore disposal is better for the environment than deep-water disposal. This social benefit is denoted b , which we normalize at $b = 1$. This benefits has the features of a public good. If Shell decided to change its decommissioning strategy all concerned consumers would benefit from the decision whether they bore the cost of participating in the boycott or not. A boycott thus results in a drop of sales for Shell. We assume that if the drop is substantial enough, Shell will yield to pressure and choose on-shore disposal.

In the complete information case, the focus on Shell’s concerned customers is without loss of generality. Concerned consumers with a strict private preference for BP, e.g. because of better gas station location, have a dominant strategy to buy from BP. On the other hand, Shell customers that do not care about the Brent Spar or believe deep-water disposal is the preferred environmental option, have a dominant strategy to buy Shell. That is, the only customers who face a strategic dilemma are *concerned Shell* customers. For them it is only worthwhile to participate in the boycott if enough other consumers participate as well. That is, they face a collective action problem (Olson 1965).

The model thus corresponds to an N -player ($N \geq 2$) discrete public goods game as defined e.g. by Palfrey and Rosenthal’s (1984) where c stands for the (net) opportunity cost of participating (e.g. the extra distance a driver has to drive to buy his gasoline from BP rather than Shell) while b ($0 < c < b$) stands for the (collective) benefit of stopping the Brent Spar from being sunk. This benefit constitutes a pure public good. If (and only if) a sufficient number of consumers k (with $1 < k \leq N$) boycott the bad product (Shell), the management of Royal Dutch/Shell will decide to dismantle to Brent Spar on-shore.

Formally, agents have two choices: they can either boycott (labeled choice 1) or decide not to

participate in a boycott (choice 0). 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 $u(z; X_{-d})$, where $z \in \{0, 1\}$ represents the agent's choice. Agent d 's payoffs 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	b
$z = 1$	$-c$	$b - c$	$b - c$

As Palfrey and Rosenthal have shown, the game has many equilibria. Specifically, there are $\binom{N}{k}$ pure strategy equilibria (each with exactly k boycotting consumers), and one pure strategy equilibrium where no boycott takes place. 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.

The importance of the Palfrey-Rosenthal model for our application lies in the fact that it demonstrates how boycotts can occur as equilibrium phenomena even if there is only a single interaction. Boycotts are thus consistent with rational action taken by concerned consumers. However, the game theoretic approach also faces some limitations. First, the Palfrey-Rosenthal game has many equilibria, some with a protest level of zero. Game-theoretic analysis, however, only specifies which outcomes are consistent with the incentives specified in the game. It does not indicate which one is more likely. Specifically, for large populations, the Palfrey-Rosenthal model implies that either boycotts will not occur with probability one, or (also with probability one) they will occur at exactly the efficient level. In the game-theoretic context we are thus left with an equilibrium multiplicity problem. Second, note that the two types of equilibria exist for all $k > 1$ and $0 < c < b$. Thus, the model cannot explain any of the following empirical phenomena: calls for boycotts are more likely to be successful if cheap substitute products are available (i.e. c is low), if the issue has high importance of salience (i.e. b is high), or if the company can ill afford to loose a large number of

customers (i.e. k is high) (Friedman 1999). Third, for protests to occur, agents must be able to solve a complex coordination problem (especially in large populations) with no apparent coordination device because all equilibria where the collective good is provided are asymmetric if $k < N$. That is, although the game is symmetric in payoffs and actions, the predicted behavior is not: some agents participate while others free-ride. This leaves us with a puzzle: how do large populations manage to overcome a stark coordination problem, especially if there is no apparent coordination device like previous experience or existing social structures?

A common solution to the problem of equilibria is to invoke the theory of “focal points” (Schelling 1960) based on the observation that agents use salient features of a particular equilibrium to coordinate. However, many focal mechanisms such as prior experience or related conventions (e.g. Schelling’s famous example of meeting in a foreign city at the train station at noon) are not available in the case of boycotts. Extensive media coverage may be interpreted as providing a focal point. However, the mechanism of how coordination is achieved through the media remains unclear. Below we will suggest such a mechanism.

Theoretical sociologists have developed an alternative formal methodology to study collective action: so-called “threshold” or “critical mass” models (Granovetter 1978, Oliver and Marwell 1988, Schelling 1978).³ Individuals in a population are assumed to vary in their willingness to participate in a collective action such as a boycott. These variations may stem from differences in costs and benefits (Oliver and Marwell 1988), or may be directly specified 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. In contrast to game-theoretic approaches critical mass models explicitly model the dynamic nature of collective action. However, while there have been some informal attempts to explicitly model the implicit adjustment processes (e.g. Schelling 1978), a rigorous treatment of the their underlying dynamics is still lacking.⁴ We

³These models have experienced a recent renaissance as “tipping point” models (Gladwell 2000).

⁴Most of the theoretical development of tipping models relies exclusively on numerical examples and simulations (e.g. Granovetter 1978, Oliver 1993).

propose a dynamic model to bridge this gap.

3 A Probabilistic Model

To explicitly analyze coordination in large populations we present a stochastic, dynamic model of collective action.⁵ This approach differs from standard game-theory in two respects: (a) the behavioral assumptions, and (b) the predictive concept. 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. Moreover, the model's predictions are not given by an equilibrium, but by a probability distribution. Specifically, we use the game's normal form to define a Markov process and then use the process' limiting distribution as our solution concept.

The Markov process consists of an action rule and a selection rule. In classical game theory agents are assumed to use best-response correspondences as their action rule. That is, behavior is completely determined by the incentives specified in the game (unless the agent is exactly indifferent between two actions). We generalize this assumption to allow for random choice behavior.⁶ Specifically, we use a random utility model (McFadden 1978). So, while each agent's mean utility is fixed, individual realizations may vary. This approach seems especially appropriate in models of boycotts where the perceived costs and benefits may well vary over time as a consequence of media coverage and other idiosyncratic sources of information.

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 . Under the standard extreme-value assumptions⁷ for the error term each individual's choice for all $d \in N$ will be characterized by the probability distribution:

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

⁵There is a large related literature on the use of stochastic models in economics. See Blume (1997), Fudenberg and Levine (1998) or Young (1998) for detailed overviews,

⁶The case of (pure) best-response is discussed in detail in section 7.

⁷See McFadden (1978) for details.

which is equivalent to the familiar log-linear choice rule. It captures the assumption that the pair-wise probability ratios of choosing actions are proportional to the respective pay-off differences. The log-linear choice model is closely connected to the best-response correspondence. The parameter β formally captures the degree to which the deterministic component of utility (given by the payoff matrix) determines choice. 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 .

In addition to an action rule we need to define a selection rule that specifies when agents act. In the Palfrey-Rosenthal game agents are assumed to act simultaneously. In our model they act sequentially: In each period t one specific agent out of N is randomly chosen with probability $1/N$. The agent then looks at the current configuration X^t of actions in the population and chooses an action according to $p^\beta(z|X_{-d}^t)$. The next period, again a player is chosen at random, and so on. Given the current configuration, an actor will then probabilistically adjust her participation behavior to improve her pay-off.

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 (with replacement) and so forth. The key idea of our model is to “decompose” the simultaneous choice of classical game-theory (where agents form conjectures about each others beliefs) into a dynamic adjustment process. As in game-theoretic models, some features of the model are mainly technical, while others are of substantive importance.

One of the technical assumptions pertains to selecting exactly one agent in each period. This does not imply that agents cannot change their behavior “fast.” After all, periods between revisions can be arbitrarily small.⁸ The *informational* implication of this assumption, however, is critical. That is, when revising their actions, agents have full information about the state of the dynamic

⁸While we adopt a discrete framewrok for simplicity, our analysis continues to hold for agents that adjust their actions in continuous time provided that the time between revisions is exponentially distributed.

system. This assumption is a natural base-line, but it also models an informational environment where boycott activity is reported in the mass media ⁹

Among the substantive assumptions perhaps the most important pertains to bounded rationality. 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 payoff function. Indeed, they do not expect that their action may influence the future decisions of other participants. Agents simply adopt the action that maximizes their current pay-off given information about the global state of the system.

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

⁹Alternative informational structures that could be investigated include agents observing a random sample of population behavior (e.g. Diermeier and Van Mieghem 2008) or only the actions in some local neighborhood (e.g. Blume 1993). See also Young (1998).

limiting distribution, denoted by the column vector π , where

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

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 $\pi = \pi P$.¹⁰ These equations are called the *global balance equations* because, rearranging $\pi_i = \sum_j \pi_j P_{ji}$, yields

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

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 $P_{i,i+1}$ by the “birth” probability λ_i (i.e., the probability that the number of participants increases by one) and $P_{i,i-1}$ 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:

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

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.

¹⁰To 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 $\pi = \pi P$.

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 $\lambda_n = P_{n,n+1} = \Pr\{X^t = n+1 | X^t = n\}$ and $\mu_n = P_{n,n-1}$. 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 subtype $(d, 0)$ 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:

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

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

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

$$\forall n \neq k-1 : \frac{\pi_{n+1}}{\pi_n} = \frac{p_0(n)}{p_1(n+1)}e^{-\beta c} = \frac{N-n}{n+1}e^{-\beta c}. \quad (2)$$

$$\text{for } n = k-1 : \frac{\pi_k}{\pi_{k-1}} = \frac{e^{\beta(1-c)}p_0(k-1)}{p_1(k)} = \frac{N-k+1}{k}e^{\beta(1-c)}. \quad (3)$$

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:*

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

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 (2), 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 (3) and (1) 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 (2) 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 (2), 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 combinatorial $\binom{N}{n}$, and the results of the action process, represented by $e^{-\beta nc}$ or $e^\beta e^{-\beta nc}$. 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 z with $z \leq x$ and $\lceil x \rceil$ as the smallest integer z with $z \geq x$ 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^*

$$n^* = \max \left\{ 0, \frac{Ne^{-\beta c} - 1}{1 + e^{-\beta c}} \right\} \quad \text{and} \quad k^* = \frac{N + 1}{1 + e^{-\beta(1-c)}}, \quad (4)$$

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, one at $\lfloor n^* \rfloor$ and another at k , of which k is the most-likely long-run state if

$$\pi_{\lfloor n^* \rfloor} < \pi_k \Leftrightarrow g(k) := (1 - (k - \lfloor n^* \rfloor)c)\beta + \sum_{i=\lfloor n^* \rfloor}^{k-1} \ln \frac{N-i}{i+1} > 0. \quad (5)$$

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 $f(0) = Ne^{-\beta c}$ and $f(N) = 0$. From (2), 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 $Ne^{-\beta c} > 1$, 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 $x^* = \frac{Ne^{-\beta c}-1}{1+e^{-\beta c}}$. 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 $\pi_1/\pi_0 = Ne^{\beta(1-c)} > 1$, 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 $f(0) = Ne^{-\beta c} > 1$ so that $n^* := x^*$ and $[n^*]$ constitutes the unique maximum for π_n .

Now consider $k > 1$. If $Ne^{-\beta c} \leq 1$, then $\pi_1/\pi_0 \leq 1$ so that π reaches a maximum at $[n^*] = 0$. If $Ne^{-\beta c} > 1$, then as before $[n^*]$ 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 $k < [n^*]$. But since π_{n+1}/π_n is increasing below n^* , we have $\pi_{[n^*]}/\pi_k > 1$ so that $n = k$ cannot be a maximum. For $n = k - 1$ to be a maximum, we need

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

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

To characterize the most likely long-run state note that (2) and (3) imply

$$\pi_{[n^*]} < \pi_k \Leftrightarrow \frac{k!(N-k)!}{[n^*]!(N-[n^*])!} < e^{-\beta((k-[n^*])c-1)}. \quad (6)$$

Condition (5) then follows immediately. ■

The proposition states that the most likely state is either $[n^*]$ or k . Notice that k is the state where an efficient number of people participates, while state $[n^*]$, on the other hand, represents random participation. That is, $[n^*]$ is entirely driven by the error component in the log-logistic choice rule; it is independent of the threshold k and depends only on N , c , and β . Indeed, as we

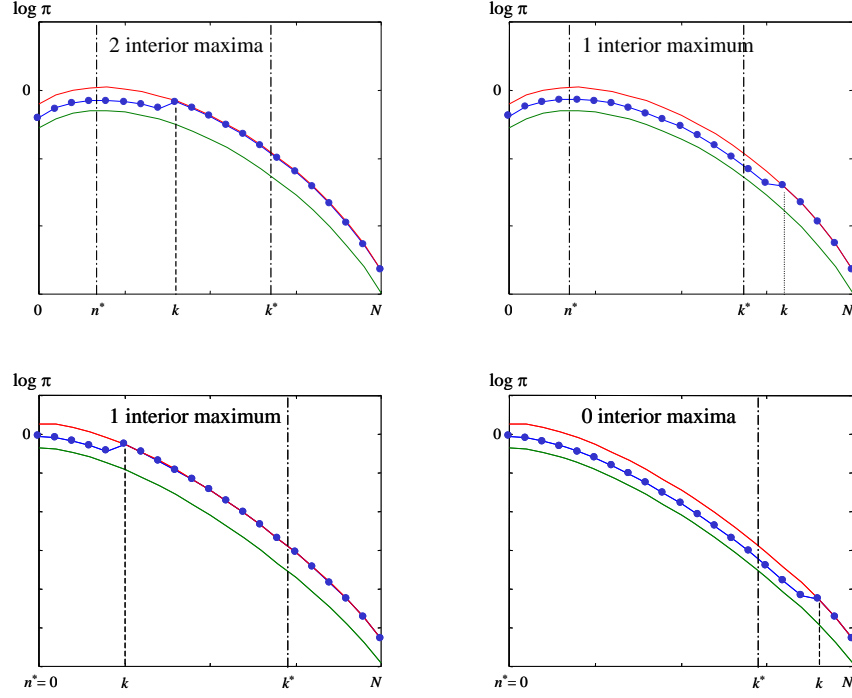


Figure 1 The limiting distribution π assumes one of four possible cases, depending on the parameters n^* , k^* and k .

reduce randomness at the individual level so that $\beta \rightarrow \infty$ (and approach best-response in the limit), $[n^*]$ approaches 0.

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 (2) 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 (5). 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 , 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.




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 $c = 0.5$.

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.

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

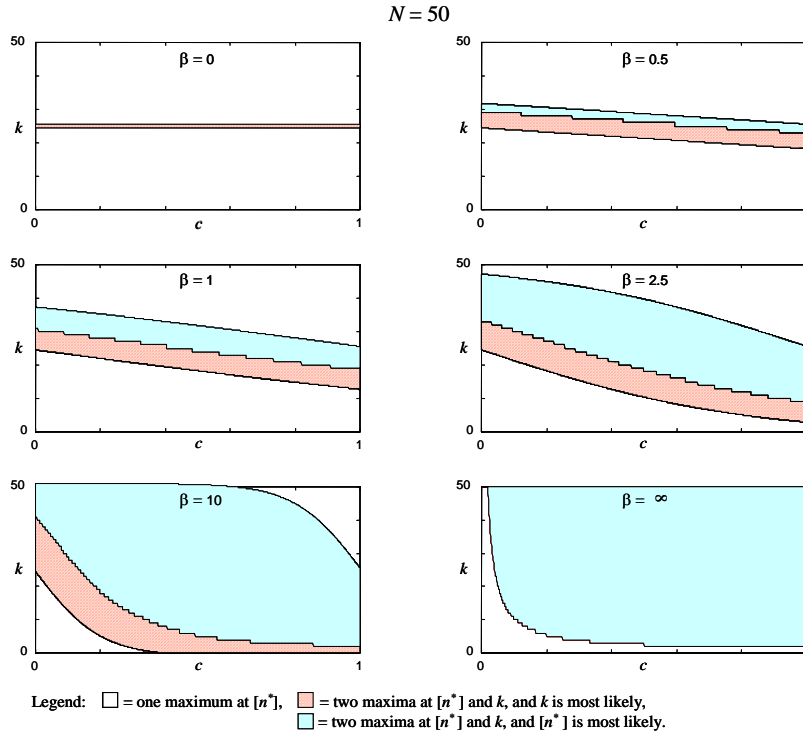


Figure 3 Strategy regions in (k, c) -space for different values of β for $N = 50$.

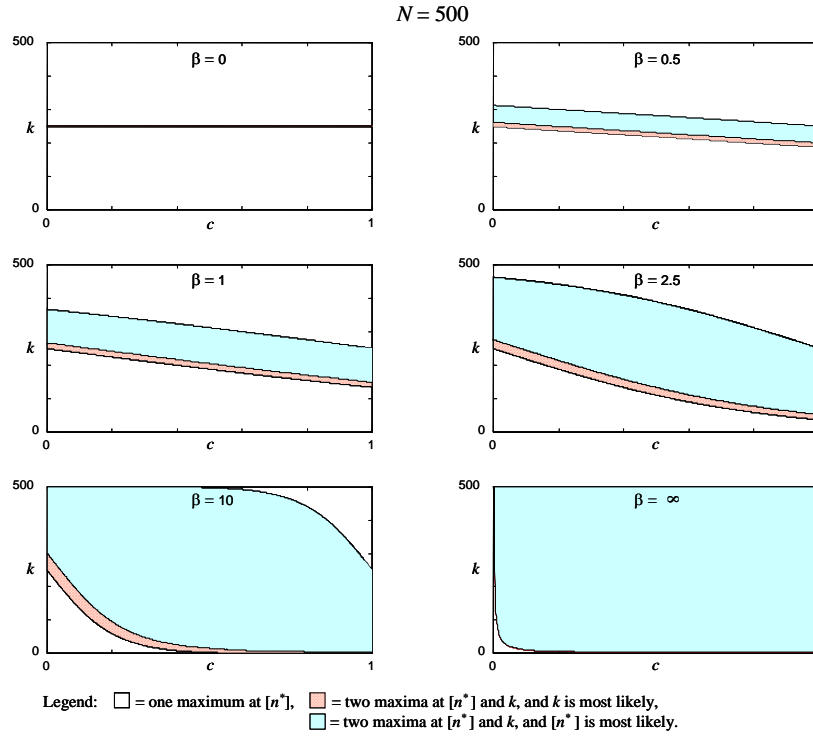


Figure 4 Strategy regions in (k, c) -space for different values of β for $N = 500$.

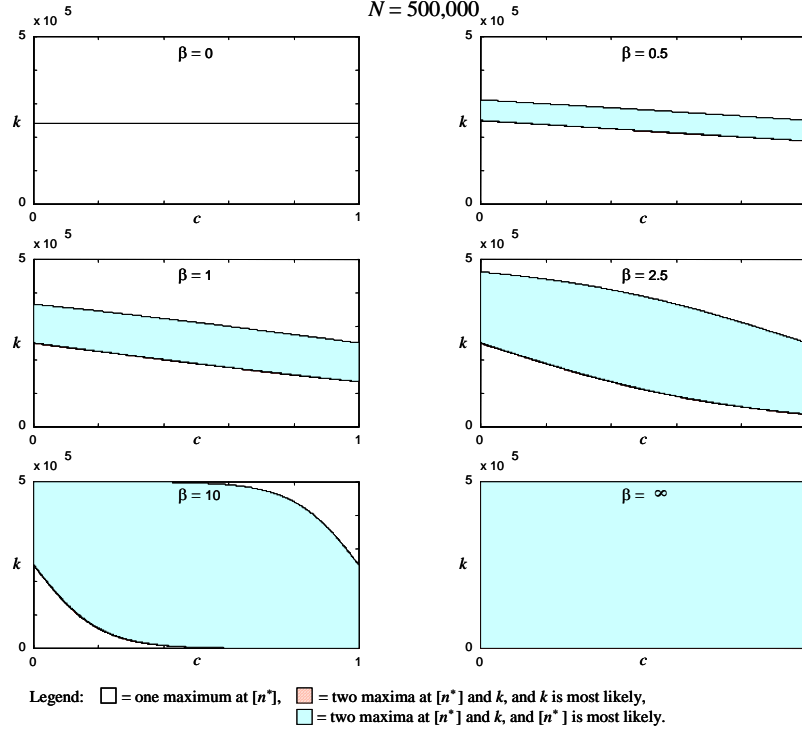


Figure 5 Strategy regions in (k, c) -space for different values of β for $N = 500,000$.

of the population N .¹¹ 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 .¹²

The existence of a local maximum at k even for very large N is one 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

¹¹Note 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.

¹²This 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."

home). Nevertheless, mass collective action may occur in the absence of any apparent coordination device.

6

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} . We can use therefore 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. 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 (5), 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 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 . That is, from the point of view of concerned consumers the selected equilibrium satisfies an efficiency property. However, the analysis in Figures 3-5, of course, presupposes that each parameter configuration is "equally likely." But it follows from the model that strategic activists will try to lower costs, increase collective benefits, or decrease the threshold k . According to the model once the threshold of $kc < 1$ is crossed, we will switch to a regime

where high participation in a boycott is very likely. Such a switch, technically a phase transition in kc -space, formally captures the fact that this particular phenomenon “has legs.”

[[-]]

[[The result implies that activists should design their campaigns carefully such that participation costs and thresholds k are low, while collective benefits are high. For example, activists should select industries where consumers have cheap substitutes, and, within the targeted industry, should target companies (or company units) with lowest switching costs. The Shell-Greenpeace controversy illustrates both points. First, vertically integrated oil companies are good targets since consumers have low costs of switching; filling up one’s car at a BP instead is enough. Second, activists that seek to change industry-practice should target a *single* firm in the same industry. In the case of the Brent Spar, Shell was targeted because of its strong global brand recognition. Third, activists may target unrelated business units of the same company if this lowers switching costs for consumers or increases perceived benefits. In the Brent Spar case Greenpeace targeted Shell Germany (not Shell UK, the truly responsible party) even though Shell Germany had nothing to do with the initial decision to seek approval for deep-water disposal (Diermeier 1996). The reason? Greenpeace expected a better strategic environment in Germany where global environmentalism has wide appeal and recycling is a national passion.

Companies, on the other side, should anticipate these incentives and then could assess their risk of being the target for an activist campaign. Possible counter-strategies include industry-wide standards or self-regulation, which may lower the benefits of targeting a specific company. Such strategic interactions between companies and activists are discussed in more detail in Baron and Diermeier (2007).]]

7 Conclusion

This paper provides a formal model of consumer boycotts as a collective action problem between concerned consumers. We show that in this model a unique equilibrium is selected. The type of equilibrium depends on the switching costs, the threshold for success, and the importance of the social dimension of the boycott to concerned consumers. If switching costs are sufficiently low, an

optimal number of agents will join the boycott, leading to mass participation.

We then discuss the model's consequences for activists' strategies. The following empirical phenomena are consistent with the model:

1. Activists should frequently rely on secondary boycotts, i.e. boycotts where the target is not the business entity engaged in the offensive practice. Secondary targeting should also occur in cases where the primary target is a well-known consumer brand. Targeting is predominantly driven by switching costs and multiplier effects. This can lead to complicated targeting chains.
2. In cases where activists try to change industry practice, they will not target the firm that caused the most egregious offense, but the most vulnerable. Activists should also limit their actions to a single target.
3. Union-sponsored boycotts should occur predominantly in cases of rights violations or exploitative working conditions, not in wage disputes.

The model provides a general, flexible model, that can be incorporated into more comprehensive models of strategic activism and counter-strategies by firms and industries (e.g. Baron and Diermeier 2007). However, the formal and empirical analysis of such interactions is still in its infancy. We hope that our approach can serve as a “work-horse” model to facilitate such analyses.

References

- [1] Baron, David.P. 2002. "Private Politics and Private Policies: A Theory of Boycotts." Working paper, Stanford University.
- [2] Baron, David P. 2003a. *Business and Its Environment*, 4th ed. Upper Saddle River, NJ: Prentice Hall.
- [3] Baron, David P. 2003b. "Private Politics." *Journal of Economics and Management Strategy*. 12 (Spring): 31-66.
- [4] Baron, David P. 2003c. "Competing for the Public through the News Media." Working paper, Stanford University.
Blume, Lawrence E. 1993. "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 4:387-424.
- [5] Baron, David P. and Daniel Diermeier. 2007. "Strategic Activism and NonMarket Strategy." *Journal of Economics and Management Strategy*, 16 (3): 599-634. .
- [6] Baron, David P. and Erin Yurday. 2004. "Anatomy of a Corporate Campaign: Rainforest Action Network and Citigroup." Case P-42A,B,C, Graduate School of Business, Stanford University, Stanford, CA.
- [7] 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.
- [8] Diermeier, D. 1996. "Shell and Greenpeace" (A), (B), and (C). Harvard Business School Case P19. Reprinted in David Baron. 2006. *Management and its Environment* 2nd-6th Edition. Prentice Hall 2008.
- [9] Diermeier, D. 2003. "Huntingdon Life Sciences" (A), (B), (C). Teaching Case. Kellogg School of Management.
- [10] Diermeier, D. 2007. "Private Politics- A Research Agenda." 2007. *The Political Economist*, XIV, 2: pp 1-2.

- [11] Diermeier, D. and J. A. Van Mieghem. 2008. "Coordination and Turnout in Large Elections." *Mathematical and Computer Modeling*. (Forthcoming).
- [12] Friedman, Monroe. 1999. *Consumer Boycotts*. New York: Routledge.
- [13] Fusfeld, D.R. 1980. *The Rise and Repression of Radical Labor in the United States, 1877-1918*. Chicago, Charles H. Kerr Publishing Company.
- [14] Fudenberg, Drew, and David Levine. 1998. *The Theory of Learning in Games*. Cambridge: MIT Press.
- [15] Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83:1420-43.
- [16] Jordan, Grant. 2001. *Shell, Greenpeace. and the Brent Spar*. New York: Palgrave.
- [17] McFadden, David. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." in P. Zarembka, ed., *Frontiers in Econometrics*. New York: Academic Press.
- [18] Oliver, Pamela E. 1993. "Formal Models of Collective Action." *Annual Reviews of Sociology* 19:271-300.
- [19] 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.
- [20] Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge: Harvard University 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] Schelling, Thomas C. 1960. *The Strategy of Conflict*.
- [23] Schelling, Thomas C. 1978. *Micro-Motives and Macro-Behavior*. New York: W. W. Norton
- [24] Shaw, R. 1996. *The Activist's Handbook*. Berkeley: University of California Press.
- [25] Taylor, Howard M., and Samuel Karlin. 1994. *An Introduction to Stochastic Modeling*. Second Edition. Boston et al.:Academic Press

- [26] Young, H. Peyton. 1993. "The Evolution of Conventions." *Econometrica* 61:57-84.
- [27] Young, H. Peyton. 1998. *Individual Strategy and Social Structure*. Princeton: Princeton University Press.
- [28] Wolman, L. 1914. *The Boycott in American Trade Unions*. Baltimore: John's Hopkins Press.