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*

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

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

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

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

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

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

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

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

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

Otherwise the most likely long-run state is $[n^*]$.

Proof: Define $f : [0, N] \to \mathbb{R} : x \to f(x) = \frac{N-x}{1+x}e^{-\beta c}$. 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 $k - 1 = \lfloor n^* \rfloor = [n^*]$ 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

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

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

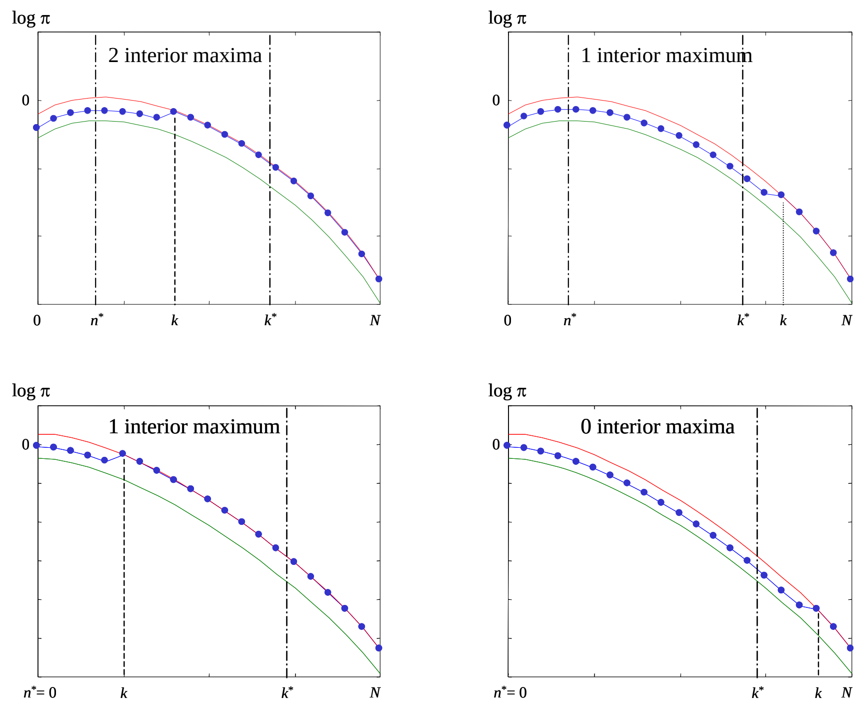

Figure 1 The limiting distribution π assumes one of four possible cases, depending on the parameters n^*, k^* and k.

reduce randomness at the individual level so that $\beta \to \infty$ (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 ($\beta \to \infty$) 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.

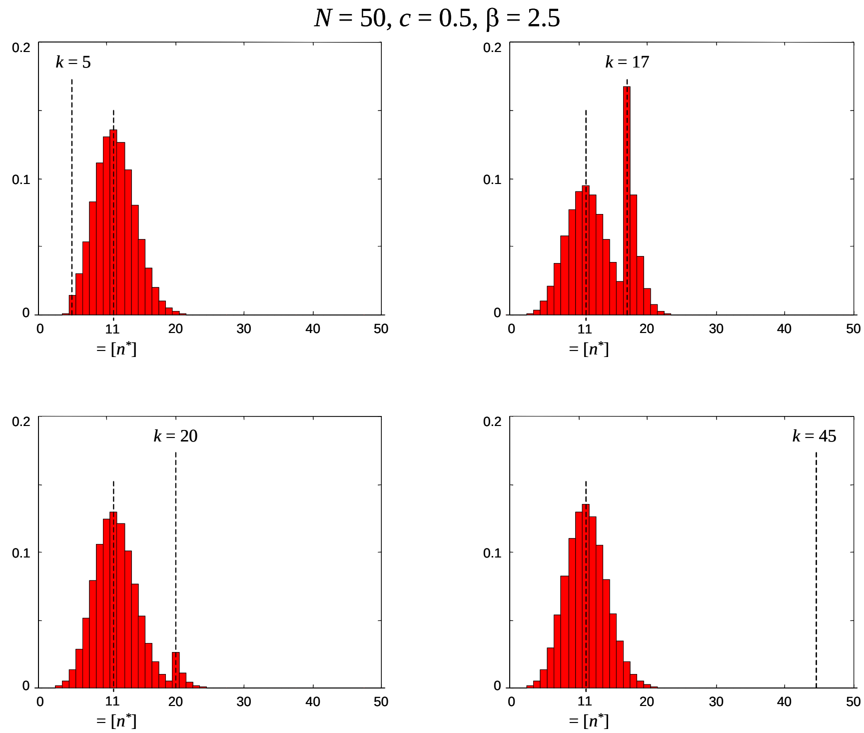

Figure 2 Four cases for the distribution π_n depending on the threshold level k. Other parameters are fixed at $N = 50$, $\beta = 2.5$, and $c = 0.5$.

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 $\beta \to 0$, the critical numbers $n^* \to (N-1)/2$ and $k^* \to (N+1)/2$. 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

This article is published in Mathematical and Computer Modelling 48 (2008) 1497-1509.

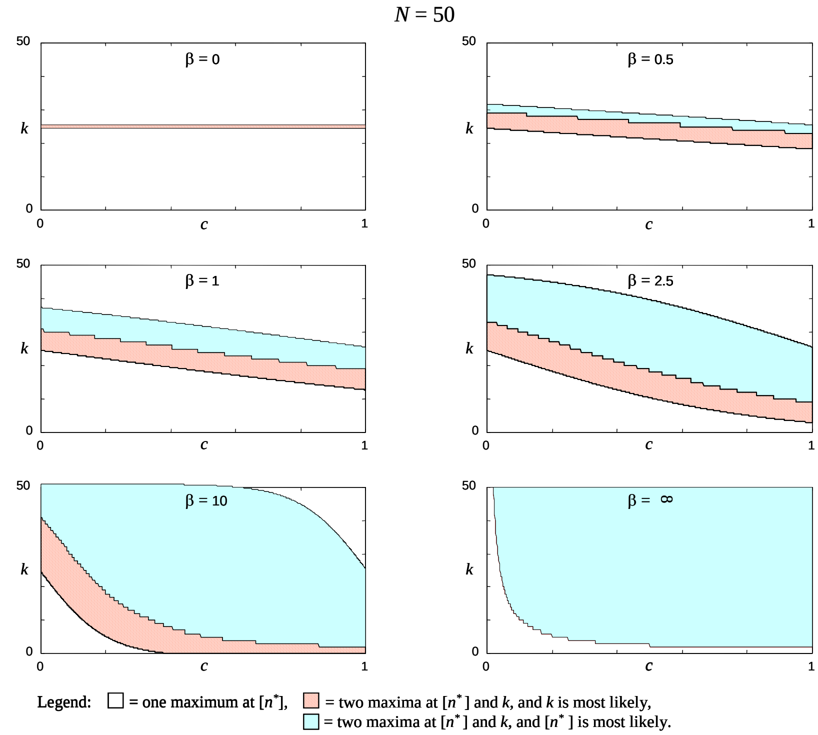

Figure 3 Strategy regions in (k, c)-space for different values of β for $N = 50$.

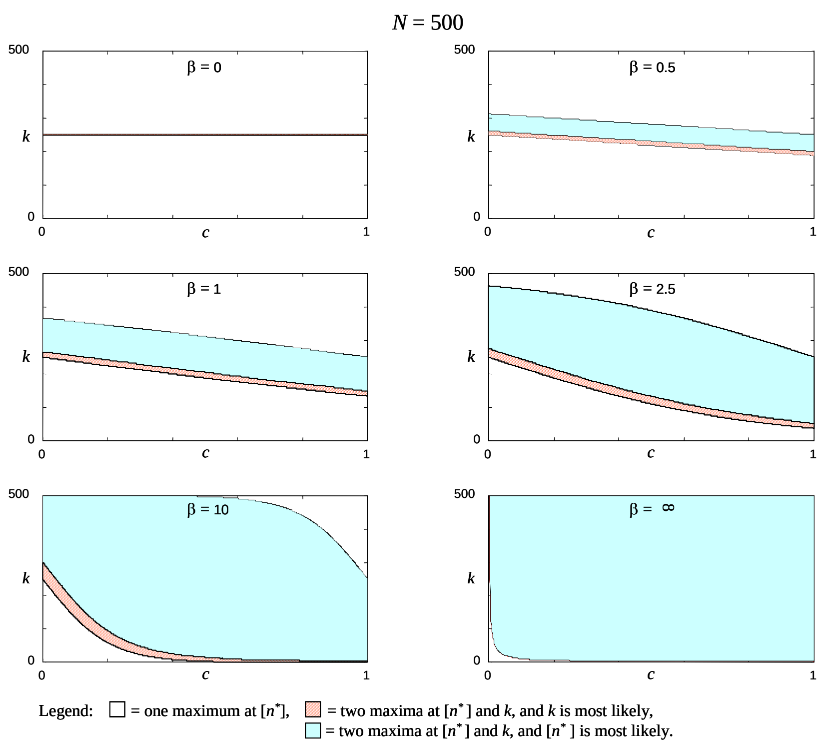

Figure 4 Strategy regions in (k, c)-space for different values of β for $N = 500$.

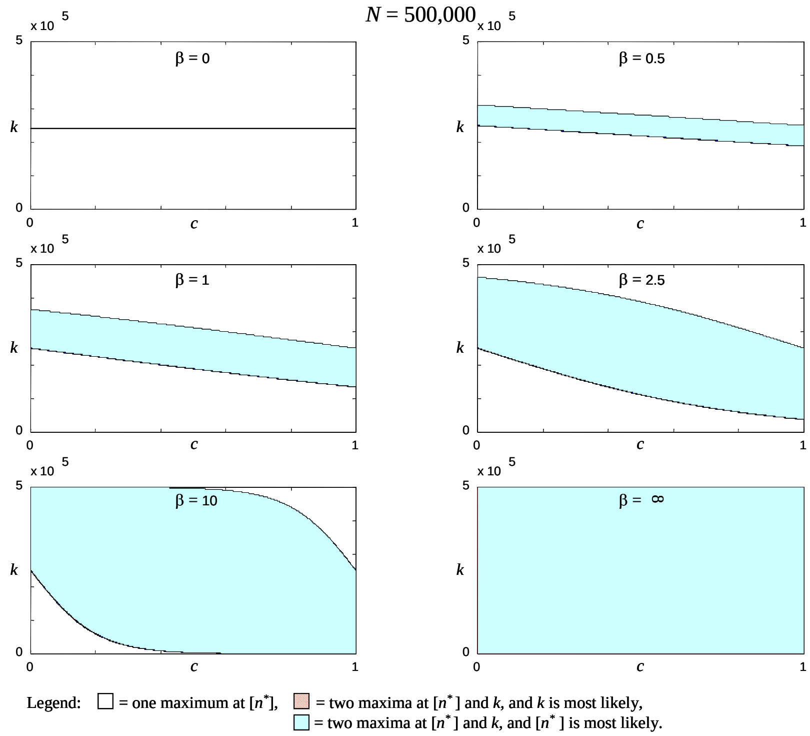

Figure 5 Strategy regions in (k, c)-space for different values of β for $N = 500,000$.

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

[11] Note 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.

[12] This 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 $\beta \to \infty$, 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 $\beta \to \infty$, then $n^* \to 0$ and $k^* \to N$, 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 $\beta \to \infty$, 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 $\beta \to \infty$, 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, form 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

This article is published in Mathematical and Computer Modelling 48 (2008) 1497-1509.

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 activits 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.