Riot Games: A Theory of Mass Political Violence

Charles Cameron and Sunita Parikh

7th Annual Conference
October 20-21, 2000

W. ALLEN
WALLIS
Institute of
POLITICAL
ECONOMY

UNIVERSITY OF
ROCHESTER

Riot Games: A Theory of Riots and Mass Political Violence

By

Sunita Parikh
Department of Political Science
Washington University

Charles Cameron
Department of Political Science
Columbia University

Prepared for the Wallace Institute Conference on Political Economy, University of Rochester, October 19-20 2000

Version 1.2: Preliminary Conference Version

We thank Peter Bearman, Randy Calvert, Rui DeFiguierdo, Jim Fearon, and participants in seminars at Columbia University and Washington University (especially Norman Schofield) for helpful comments on an early version of the paper.

I. Introduction

Riots, pogroms, and mass political violence have plagued political life, ancient and modern. At the turn of the twenty-first century, racial, ethnic, and religious violence are pervasive in heterogeneous societies around the world. Few social phenomena are as disturbing, intractable and – perhaps paradoxically – so well-studied yet still so perplexing.

In this paper, we use recent advances in game theory to study the initiation of riots and mass collective violence. Our approach unifies and to some extent reconciles two distinct approaches to the study of conflict in heterogeneous societies. On the one hand are explanations for violence that see outbreaks of violence as manifestations of long-standing cultural antagonisms, the "age-old ethnic hatred" beloved by the popular press (Connor 1994, Smith). The difficulty of explaining variations in conflict with a constant like unchanging hostility, however, has encouraged the development of a quite different type of explanation. This approach attributes conflict to the machinations of elites with material economic and political goals, who manipulate the masses by appealing to culturally salient symbols and threats (Brass).[1] The "hatred" versus "manipulation" approaches have tended to talk past one another, with each stressing a preferred causal mechanism as the "best" explanation of conflict. Neither investigates how affective

[1] In addition, more recent analyses portray cultural identity as socially and historically constructed and hence somewhat flexible, rather than completely stable and enduring. In this view, the salience of different identities varies across contexts and, since identities can overlap with each other, may be the object of personal choice (Anderson 1984, Laitin 1998).

attachments and strategically motivated manipulation interact to create outbreaks of violence. Yet it is precisely this interaction that may be most characteristic of large scale violence in racially, ethnically, and religiously diverse societies. In this paper, we pursue this insight. We first show formally and then through empirical examples exactly how affective attachments within a large group and strategic manipulation by a smaller one can interact to create riots and mass collective violence.

Our models have a close relationship with others that exist in the huge literature on the theory of collective action. Our treatment of "expressive rioters" (those with affective attachments) is closely related to the critical mass or threshold models initiated by Schelling 1978 and explored further by Granovetter 1978.[2] These ideas continue to be widely employed in the social sciences (see e.g., Laitin 1998 and Kuran 1991.) However, our analysis differs from the Schelling-style models in two respects. First, the microfoundations of those models are not elaborated in detail and remain rather murky. We supply explicit microfoundations for the Schelling-style models, modeling collective action by expressive rioters as a large Poisson game with population uncertainty (Myerson 1997a, 1997b, 1998). This is an attractive approach because the exact number of rioters is surely a matter of uncertainty for the potential rioters, and is an important part of the strategic situation they face. The Poisson game approach has the additional advantage of allowing a comparatively simple formal analysis of the situation, mirroring

---

[2] The critical mass models are quite different from models of information cascades [citation]. In the latter, individual action reflects private information about the state of the world. Late moving actors may infer the private information held by early moving actors by observing their actions, up to the point at which an "information cascade" occurs, after which no further learning is possible.

the conceptual clarity of the early threshold models. Second, the analysis of Schelling-style models often focuses on the dynamics of achieving critical mass, typically assuming an arbitrary tannonement process. Instead, we focus on the existence of multiple, qualitatively different equilibria. In our view, the coordination problem created by multiple equilibria is central to the phenomenon of rioting.

Our treatment of "riot provocateurs" is similar to Palfrey and Rosenthal's model of collective provision of a discrete public good (1984). In that model, each individual may participate by making a fixed contribution to a discrete public good. If a sufficient number contribute, the good is provided, but otherwise it is not. In other words, it is also a threshold model, but aimed at the supply of a public good. Our analysis is distinguished from Palfrey and Rosenthal's in that they assume the number of players is common knowledge. Again, we model the interactions of the riot provocateurs as a Poisson game with population uncertainty.[3] This allows a simpler treatment of the problem and leads to an arguably more plausible characterization of how the riot provocateurs overcome the free rider problem that dogs their efforts.

The distinctive contribution of this paper, however, is to link the two games – the expressive rioting game and the provocation game – into a unified whole, with an explanation for mass political violence lying in the interaction of the two games. In particular, we show how a successful solution to the provocateurs' collective action problem can reduce the available equilibria in the expressive rioting game to a single

---

[3] The provocation model is distinct from Lohmann's ingenious models of mass signaling, in which collective action aggregates private information dispersed in a group, and signals that information to an outside party (e.g., Lohmann 1994).

equilibrium characterized by a large-scale riot. The aim of the collective action by the provocateurs is to do exactly that. We know of no formal analysis of mass collective violence that employs this approach.

The huge empirical literature on mass political violence contains, we believe, many examples of the phenomenon we model. For example, several analysts have discussed the existence of riot provocateurs and link their activities to large scale political violence. Paul Brass, for instance, in a noted study of collective violence in India, records the existence of "fire-tenders" or "riot specialists" who operate within "institutionalized riot systems" (1997:16). He notes, "When the time is right for the fomenting of a large-scale disturbance, then students, hooligans, low-caste persons from slums and outlying areas, criminals, and *special squads of trained activists* such as the members of the Bajrang Dal will be brought in. (p. 282, emphasis added) He provides a portrait of one such riot specialist (Chapter 7). Ethnographies of riots and studies of extremist politics offer portraits of other such individuals. However, the empirical literature is largely silent on the strategic problems facing riot activists, and how this strategic calculus intersects with that of potential rioters. We supply this analysis.

The paper is organized in the following way. In Section II, we study riots as a large Poisson game. We show that three outcomes are candidates for equilibria in a Schelling-like threshold model. All or some of these equilibria may exist under different conditions, which we specify. We also show that a small riot begun by riot provocateurs can reduce the multiple equilibria to a single equilibrium in which there is a large riot. In Section III, we study the collective action problem facing riot provocateurs. The familiar free riding problem is a threat to successfully triggering a riot, but it is possible for rational provocateurs to overcome this problem. In Section IV, we use the models and qualitative data from histories and case studies to interpret episodes of large scale collective violence. Section V concludes.

II. A Model of Expressive Rioting

We begin with a simple version of the model, in which there is but a single type of player. We show how multiple equilibria may exist in this model, given different parameter values. Then we examine how a pre-existing riot can dramatically reduce the available equilibria. We conclude the section with illustrative numerical examples.

The Model

The number of players in the game is a random variable drawn from a Poisson distribution with mean $n > 0$. Given this parameter n, the probability that there are k players in the game is

$$p(k \mid n) = \frac{e^{-n} n^k}{k!}$$

where $e = 2.71828\ldots$ and $k! = 1 \times 2 \times 3 \ldots \times k$ with $0! = 1$.[4]

The action set C for each player is $C = \{0,1\}$, where "0" denotes "don't riot" and "1" denotes "riot." Let $x(c) \in X = \{\square \cup 0\}$ (the set of non-negative integers) denote the

[4]See standard references on probability, e.g., Mood, Graybill, and Bois. Throughout this paper we employ notation similar to that in Myerson 1998, to facilitate comparison with standard references on Poisson games.

number of players in the game who choose action *c*. We will be primarily concerned with $x(1)$, the number of players who riot.[5]

Let the gain from rioting *t*, a non-negative real number, be a private good that is acquired only through the act of rioting. In the following section, we will allow *t* to vary over players and be private information, but in the simple version of the model we assume *t* is the same for all players and common knowledge. The gain *t* may constitute the expected proceeds from looting. But it may also be a non-instrumental expression of one's ethnic, racial, or ideological identity.[6] We scale *t* so that it is bounded by zero and one. We denote the set of possible values of *t* as *T*.

Let $l : X \to \mathbb{R}_+$ be the potential cost of rioting to a player. The potential losses from the act of rioting include the possibility of beatings, arrest and prosecution by the authorities, as well as other physical risks inherent in the activity. We assume the per-rioter cost of rioting falls as the number of rioters increases, at least in some range of rioters. More precisely, we assume the loss function is a bounded, continuous function of $x(1)$, the number of rioters, and that the function is everywhere at least weakly decreasing and is strictly decreasing for some values of $x(1)$. Call the maximal cost of rioting, which occurs when only one person riots, the *ceiling cost* of rioting, $\bar{l} \equiv l(1)$. Call

---

[5] We will not need to consider "action profiles" (the vector $x = (x(0), x(1))$) nor the set of all possible action profiles $Z(C)$, although action profiles and $Z(C)$ play an important role in the analysis of canonical Poisson games (Myerson 1997).

[6] If so, *t* may be a function increasing in $x(1)$, but we abstract from this complication.

the minimal cost of rioting, $\underline{l} \equiv \lim_{x(1)\to\infty} l(x(1))$, the *floor cost* of rioting. We assume $l(\bullet)$ is scaled so that that $0 \le \underline{l} < \bar{l} \le 1$.

The utility for each player in the game is a function $U : T \times X \times C \to \mathbb{R}$. More specifically, we specify U as:

$$U(x(1), c; t) = c\left[t - l\left(x(1)\right)\right] \tag{1}$$

If c = 0, then $U(\bullet,\bullet,\bullet)$ = 0. If c = 1, then $U(\bullet,\bullet,\bullet) = t - l(x(1))$. Under the above assumptions, this function is bounded.

Strategies in Poisson games have a character that is distinctly different from strategies in traditional games in which the set of players is assumed to be common knowledge. As Myerson explains,

> In a traditional game…, we describe players' perceptions of each others' strategic behavior by strategy profiles that predict a distinct random strategy for each player in the game [and these strategies may differ with the identity of the player]. In our games with population uncertainty, however, players' perceptions about each others strategic behavior cannot be formulated as a strategy profile that assigns a random strategy to each specific individual in the game, because a player is not aware of the specific identity of all the other players…It cannot be commonly perceived that two different individuals of the same type would behave differently because, in our model with population uncertainty, two players of the same type have no commonly know attributes by which others can distinguish them. . . In effect, population uncertainty forces us to treat players symmetrically in our game-theoretic analysis. (1997:4-5).

Hence, in this game in which all the players are identical and move simultaneously, all players must all employ the same strategy $s(c) \in [0,1]$, in other words, the same probability of choosing action c.

Two other important properties of Poisson games require mention. First, when the players behave according to strategy $s(c)$, $s(c)$ >0, the number of players who choose each action c in C is a Poisson random variable with mean $ns(c)$. This is true due to the "decomposition property" of the Poisson distribution (Myerson 1997a:6).[7] Hence, $x(1)$ is a Poisson random variable with mean $ns(1)$. Second, from the perspective of any player in the game, the expected number of players *other than himself* is also a Poisson random variable with the same mean $ns(1)$, a feature known as "environmental equivalence."[8]

Given these properties of Poisson games, for a randomly selected player, the expected utility of choosing action c when all other players are expected to behave according to strategy $s(1)$ > 0 (the probability of rioting) is

$$\overline{U}(c \mid s(1)) = \sum_{k=0}^{\infty} p(k \mid ns(1))\, U(x(1), c)$$

$$= \begin{cases} 0 & \text{if } c = 0 \\ t - \sum_{k=0}^{\infty} p(k \mid ns(1))\, l(k+1) & \text{if } c = 1 \end{cases} \tag{2}$$

[7] Furthermore, the number of players who choose the action c is independent of the numbers of players who choose all other actions . So in this game, the number of rioters and non-rioters are independent random variables. This "independence-actions property" of Poisson games is proven in Myerson (1997).

[8] Myerson 1997 provides a proof of the "environmental equivalence" result (Section 5).

Equilibria

Because C is finite and U is bounded, it follows from Theorem 3 in Myerson 1997a that at least one equilibrium exists in this Poisson game. But we wish to characterize the equilibria and establish conditions on their existence. We proceed through an examination of incentive compatibility constraints, in the spirit of d'Aspremont and Gerard-Varet 1979.

Because population uncertainty imposes symmetry on strategies, there are only three classes of equilibria to consider: a pure strategy equilibrium in which all players riot; a pure strategy equilibrium in which no players riot; and mixed strategy equilibria in which some players riot and others do not, but all employ the same mixing probability.

Type 1 Equilibrium (All players riot: $s(1)=1$)

If $s(1)=1$ is to be played by all players in equilibrium, it must be the case that for a randomly sampled player the expected utility of $s(1)=1$ is weakly greater than that of $s(1)=0$.[9] In other words, this equilibrium requires

$$\overline{U}\left(1 \mid s(1)\right) \geq \overline{U}\left(0 \mid s(1)\right)$$

$$\Rightarrow t - \sum_{k=0}^{\infty} p\left(k \mid n\right) l\left(k+1\right) \geq 0$$

[9] We need not consider deviations to a non-degenerate mixing probability. If such a deviation were profitable it would be dominated by a deviation to $s(1)=0$. If the player were indifferent between such a deviation and maintaining $s(1)=1$ (a mixed strategy equilibrium) then she would also be indifferent between $s(1)=0$ and $s(1)=1$. So it is sufficient to check a deviation to $s(1)=0$.

$$\Rightarrow t \geq \sum_{k=0}^{\infty} p(k \mid n)\, l(k+1) \qquad (3)$$

In other words, the expected benefit of rioting t must be greater than the expected loss of rioting.

An obvious consequence of Equation (3) is that the equilibrium necessarily exists if $t \geq \bar{l}$. Equally obviously, the equilibrium cannot exist if $t < \underline{l}$. The more interesting situation involves $\bar{l} > t > \underline{l}$. In this case, the existence of the equilibrium depends on the mean size of the population of players, n, as indicated in Lemma 1.[10]

Lemma 1.

If $\bar{l} > t > \underline{l}$ then there exists n^* such that $t \geq \sum_{k=0}^{\infty} p(k \mid n)\, l(k+1)$ for all $n \geq n^*$.

Proof: See Appendix.

Roughly speaking, the lemma indicates that if the expected cost of rioting falls below t when a riot becomes sufficiently large, then the riot equilibrium can exist if the expected population of players is large enough.

Type 2 Equilibrium (No players riot: $s(1) = 0$)

This equilibrium requires that

$$\bar{U}(0 \mid s(1)) \geq \bar{U}(1 \mid s(1))$$

[10] If $\bar{l} > t = \underline{l}$, the equilibrium cannot exist, because for any n there will be some probability weight placed on values of $x(1)$ such that $l(x(1)) > t$. No matter how much remaining weight falls on values of $x(1)$ such that $t = \underline{l}$, Equation (3) cannot be satisfied.

$$\Rightarrow 0 \geq t - l(1)$$

$$\Rightarrow t \leq \bar{l} \tag{4}$$

This condition is very weak: If the benefit of rioting exceeds the ceiling cost of rioting, the equilibrium cannot exist. But otherwise it can. Note however, that when $t = \bar{l}$ a Type 2 equilibrium is not perfect: any chance of a deviation by any player makes $\overline{U}(1 | \boldsymbol{s}(1)) > \overline{U}(0 | \boldsymbol{s}(1))$. Hence, we replace (4) with

$$t < \bar{l} \tag{5}$$

Type 3 Equilibrium (Players riot probabilistically: $s(1) \in (0,1)$)

The logic of a mixed strategy requires that

$$\overline{U}(0 | \boldsymbol{s}(1)) = \overline{U}(1 | \boldsymbol{s}(1))$$

$$\Rightarrow 0 = \sum_{k=0}^{\infty} p(k | n\boldsymbol{s}(1)) \left[\left[t - l(k+1)\right]\right]$$

(using environmental equivalence)

$$\Rightarrow 0 = t - \sum_{k=0}^{\infty} p(\bullet) l(\bullet)$$

$$\Rightarrow t = \sum_{k=0}^{\infty} p(k | n\boldsymbol{s}(1)) l(k+1) \tag{6}$$

It is obvious that Equation (6) cannot be satisfied if $t > \bar{l}$ or if $t < \underline{l}$. Neither can it be satisfied if $t = \bar{l}$ or if $t = \underline{l}$.[11] The more interesting case is $\bar{l} > t > \underline{l}$, addressed in Lemma 2.

Lemma 2.

If $\bar{l} > t > \underline{l}$, there exists an *n*, *n**, such that for all *n*> *n**, there exists a $s(1) \in (0,1)$ such that $t = \sum_{k=0}^{\infty} p(k \mid n s(1)) l(k+1)$.

Proof: See Appendix.

In other words, if *n* is large enough, the mixed strategy equilibrium can exist, when $\bar{l} > t > \underline{l}$.

We gather these results together in Proposition One, which is illustrated in Figure Two for a specific example.

PROPOSITION ONE

In the simple game (where *t* is common to all players), the possible equilibria can be characterized in the following way: 1) If $t \geq \bar{l}$, then only a Type 1 (riot) equilibrium exists. 2) If $\bar{l} > t > \underline{l}$, then a) if *n* < *n**, only a Type 2 (no riot) equilibrium exists, but b) if

[11] Recall that $l(k)$ is assumed to be weakly decreasing in k, and strictly decreasing for some values of k. So there exists a k' such that $l(k') < \bar{l}$. For all *n* > 0 and $s(1) > 0$, $p(k' \mid n s(1)) > 0$. Hence $\bar{l} > \sum_{k=0}^{\infty} p(k \mid n s(1)) l(k+1)$. Similarly, for all *n* > 0 and $s(1) > 0$, $p(0 \mid n s(1)) > 0$. So there is some probability weight placed on $\bar{l}$, and consequently $\underline{l} < \sum_{k=0}^{\infty} p(k \mid n s(1)) l(k+1)$.

$n > n^*$, Types 1, 2, and 3 (mixed strategy) equilibria exist. 3) If $t < \underline{l}$, then only a Type 2 equilibrium exists. 4) If $n = n^*$, only Types 1 and 2 equilibria exist.

Discussion

Proposition One has the following interpretation. "Small" crowds are dangerous only if t is so high that each member of the crowd would riot regardless of what others do. Otherwise, there can be no riot in a small crowd. In contrast, "large" crowds are potentially dangerous, even when t is much lower (but not so low that it is not worth while to riot even if everyone else were to do so). But such a large crowd is only potentially dangerous, for its members face a difficult coordination problem. The problem is, no one will riot unless they are sure that enough others will too – rioting must be focal. Of course, particularly salient events can make the riot equilibrium focal. For example, crowds of Muslims all across India broke out in riot after the destruction of the Ayodhya mosque, a dramatic event that symbolized for many a governmental retreat from inter-ethnic neutrality. Race riots began in many cities across the U.S. in the hours after the assassination of Martin Luther King, as news of his death spread. Riots broke out in Los Angeles, following the announcement of verdicts in the racially charged, extremely controversial, and closely followed Rodney King trial. In general, however, when multiple equilibria exist, it is the non-riot equilibrium that is apt to be focal, because the absence of a riot is the baseline from which most crowds act.

The Effect of A Pre-existing Riot

In the simple game, the players all move simultaneously. Suppose, however, that a separate group of players acts earlier, staging a riot. Then the decision of players in the simple game (the "expressive players") is no longer whether to riot, but whether to join the on-going riot. In this section, we investigate the decision made under these altered circumstances.

Denote the size of the pre-existing riot as $\hat{x}(1) \in \hat{X} = X$. Because the expressive players move after observing the pre-existing riot, a strategy for them is now a function $s : \hat{X} \rightarrow \Delta(C)$, where $\Delta(\bullet)$ is the set of all probability distributions over a finite set. This definition of a strategy raises the possibility of a kind of "sunspot" equilibrium, in which the pre-existing riot does not materially affect the utility of the expressive rioters but makes one equilibrium focal from amid several candidates.[12] For example, multiple equilibria in the expressive game may exist but if the pre-existing riot reaches a sufficient size, the expressive rioters may shift their expectations from a Type 2 equilibrium to a Type 1 or 3 equilibrium, a shift in expectations that would be self-fulfilling. However, we concentrate here on the way the pre-existing riot *materially* affects the utility of the expressive players through its impact on their loss function.

Define the effective loss function facing the players of the expressive game as $l : X \times \hat{X} \rightarrow \mathbb{R}_+$, with $l\left(x(1); \hat{x}(1)\right) = l\left(x(1) + \hat{x}(1)\right)$. The loss function of the prior sections is simply this function, when $\hat{x}(1) = 0$. The effective ceiling cost of rioting facing the

[12] The terminology derives from certain models in macroeconomics. In these models, there are multiple equilibria. Rather arbitrary, non-payoff relevant events ("sunspots") determine which of the possible equilibria the economy resides in (citation).

expressive players is $\bar{\bar{l}} \equiv l\left(1+\hat{x}(1)\right)$. However, the effective floor cost of rioting remains unchanged, since $\lim_{x(1)\to\infty} l\left(x(1)+\hat{x}(1)\right) = \lim_{x(1)\to\infty} l\left(x(1)+0\right)$.

Under the earlier assumptions about the loss function (at least weakly decreasing everywhere and strictly decreasing somewhere), if $\hat{x}(1) > 0$ then the ceiling cost of rioting weakly falls, as does the cost of rioting at any level of riot (i.e., $\bar{\bar{l}} \le \bar{l}$ and $l\left(x(1)+\hat{x}(1)\right) \le l\left(x(1)+0\right)$). The expected cost of rioting for the expressive players, $\sum_{k=0}^{\infty} p\left(k \mid n\boldsymbol{s}(1)\right) l\left(k+1+\hat{x}(1)\right)$, strictly falls as $\hat{x}(1)$ increases, for all values of $\hat{x}(1)$.[13] These facts imply that a pre-existing riot can alter the available equilibria in the expressive rioting gam*e because the pre-existing riot reduces the effective cost of rioting*, as indicated in the following proposition.

PROPOSITION TWO

For all t such that $\bar{t} > t > \underline{t}$, there exists an $\hat{x}(1)$ such that $t \ge \bar{\bar{t}}$.

Proof: See Appendix.

Discussion

Proposition Two has the following interpretation. A pre-existing riot reduces the cost of rioting for the expressive players. Suppose the benefit of rioting is greater than the floor cost of rioting (otherwise, a riot is never possible). If the pre-existing riot is large enough, it can reduce the cost of rioting for the expressive rioters so much that the only equilibrium in the game is one in which all the expressive players riot. A corollary to the proposition is somewhat weaker. Suppose t is greater than the floor cost of rioting but the crowd is too small to be potentially dangerous (the only equilibrium is the no riot equilibrium). Then a pre-existing riot of sufficient size can make the crowd "dangerous," that is, Type 1 and Type 3 equilibria can become possible.[14]

An Example of the Expressive Rioting Game

To illustrate the model, we introduce an example. In the example, the loss function is:

$$l\left(x(1)\right) = \frac{1}{\boldsymbol{a} + (x(1))\boldsymbol{b}}$$

where $\boldsymbol{a} \geq 1$ and $0 < \boldsymbol{b} \leq 1$. Thus, $\underline{l} = 0$. Because $\bar{l} = \dfrac{1}{\boldsymbol{a} + \boldsymbol{b}}$, $0 < \bar{l} < 1$.

The parameters $\boldsymbol{a}$ and $\boldsymbol{b}$ have a substantive interpretation. The ceiling cost $\bar{l}$ can be seen as a measure of the police "presence" at a potential riot. In other words, the ceiling cost indicates the implicit response of the authorities to an individual attempting

---

[13] The derivative of the expected loss function with respect to $\hat{x}(1)$ is $\sum_{k=0}^{\infty} p\left(k \mid n\boldsymbol{s}(1)\right) l'\left(k+1+\hat{x}(1)\right)$.

All the $l'(\bullet) \leq 0$ and some $l'(\bullet) < 0$ so the sum must be less than zero.

[14] A formal statement of this corollary is: If $\bar{t} > \sum_{k=0}^{\infty} p\left(k \mid n)\right) l\left(k+1\right) > t > \underline{l}$ then there exist $\hat{x}(1)$ such that $\bar{t} > t \geq \sum_{k=0}^{\infty} p\left(k \mid n)\right) l\left(k+1+\hat{x}(1)\right) > \underline{l}$.

to initiate a disturbance. For small b, $\bar{l} \approx 1/a$, so that larger values of a indicate a smaller police presence. In contrast, b measures the resilience of the authorities in the face of a riot, their ability to resist being overwhelmed by a mob. Small values of b indicate that the authorities can continue to impose costs on rioters even during a large riot. Larger values indicate that rioters easily swamp the ability of the police to respond effectively. A degree of police collusion with the rioters or even simple incompetence among the authorities – both of which are often alleged to be an important part of ethnic rioting (see e.g., Brass 1997:273-276, 286-288) – can be represented by high values of a and b.

Closed form solutions to the model require calculating the expected loss function $\sum_{k=0}^{\infty} p(k \mid n)) l(k+1)$, a sum with infinitely many terms. However, the bulk of the density of a Poisson random variable lies near the variable's mean (recall that the variance of a Poisson distributed variable equals the mean). Thus, summing over the first, say, 10n terms in the expected loss function provides a very, very close approximation to its true value, even for small n. We use this approximation in calculating numerical examples. Figure One shows the expected loss function for values of n between 1 and 1000, given $a = 1.49$ and $b = .01$.

The left-hand panel in Figure Two illustrates Proposition One, again for the case in which $a = 1.49$ and $b = .01$. As shown, the Type 1 (riot) equilibrium exists over a large range of t's and n's but it is the sole equilibrium only for rather high values of t, at or above $\bar{l}$ = 2/3. The right-hand panel of Figure Two illustrates Proposition Two. It shows the effect of a pre-existing riot of 100. Given the riot, the effective ceiling cost has dropped to .40 and the effective expected loss function has shifted down. Some combinations of t and n that could not support a Type 1 equilibrium without the pre-existing riot now do so. In other words, some crowds that were too small to be "dangerous" become so in the face of an instigating riot. But even more dramatic is the expansion of the area in which *only* a Type 1 equilibrium can exist. In these cases, a relatively small group of instigators pushed a crowd into active rioting. The case is rather dramatic but illustrates how a small spark can ignite a conflagration.

III. A Model of Riot Provocateurs

The sparks that fire a large riot may simply reflect the presence of individuals with an extraordinary propensity to violence. In other words, if a crowd contains some individuals with values of t greater than the ceiling cost of rioting, these "high demanders" will riot. If the mass of these rioters is large enough, it may impel a large crowd of people with lower t's into active rioting. This is the essential insight of the Schelling-style tipping models, which we formalized in the previous section. But, the ability of a small group to instigate a large disturbance raises an intriguing possibility: can a small strategically-minded group deliberately instigate a large act of collective violence?

The role of strategic provocateurs is hardly academic. The 19^{th} and 20^{th} centuries saw the rise of groups that explicitly advocated domestic disorder and mass violence as tactics for advancing their political interests. Perhaps the most famous example is the Nazis in the 1930s, with their brutal, street fighting thugs of the SA. But Brass, in his study of collective violence in India, notes "…the deliberate inculcation among the RSS cadres [the RSS is an extremist Hindu nationalist organization] who provide the shock troops for the entire ‘family’ of its organizations . . . of a cult of violence aimed at the intimidation of Muslims, their selective killings, and the destruction of their properties during riots. . . Those well-skilled in the practices of violence, prepared to use them against Muslims, are portrayed as heroes” (p. 282). Surely groups other than those explicitly advocating violence have been drawn to such stratagems as well (references).

In this section, we consider the strategic problems facing a group of riot provocateurs who stand to benefit from domestic disorder. In the interest of generality, we do not explicitly model the domestic politics that can bring such benefits to a group. Instead, we subsume them in a parameter $\hat{t}$ that is a public good for members of the group, if they can instigate mass violence. For example, a large riot may topple the government, an end desired by the provocateurs. The benefits from a toppled government accrue to all who share this goal regardless of whether they undertook the costly action of sparking the riot. Hence, those who wish to spark the riot in order to topple the government must over-come a free-riding problem. The central issue for the model is the ability of the provocateurs to overcome free-riding and successfully trigger large scale riots and domestic disorder.

We focus on the most difficult set of circumstances for the riot provocateurs. If the provocateurs can solve their collective action problem under these circumstances, they are only more likely to do so under less adverse circumstances. To focus on the toughest set of circumstances, we assume the provocateurs receive no private benefits from participating in the riot (for them, $t = 0$). If they did – as in fact most surely do – then this benefit would blunt the incentive to free-ride and thus make effective collective action easier. We also assume that if multiple equilibria exist in the expressive rioting game, a Type 2 (no riot) equilibrium is focal.[15] So the provocateurs gain $\hat{t}$ only if multiple equilibria do not exist in the expressive game, and the equilibrium that exists is a riot equilibrium. This assumption places the heaviest possible burden on the provocateurs. If they can use a very small disturbance to manipulate focal points in the expressive game, this will make the provocateurs' job easier; but we do not invoke this mechanism in our analysis.

The Model

The model is broadly similar to that of the previous section. We assume the number of players in the provocation game is a Poisson random variable with mean m. The choice set of the provocateurs is identical to that of the expressive rioters. That is, a choice $c \in C = \{0,1\}$ where "1" denotes "riot" and "0" denotes "don't riot." As in the previous section, we denote the number of provocateurs who riot as $\hat{x}(1) \in X$, the set of non-negative integers. The loss function from rioting for the provocateurs is exactly the same as that facing the expressive rioters absent a pre-existing riot, with identical ceiling and floor costs. We assume the gain to the provocateurs from a riot by the expressive players is, in expectation, $\hat{t} \in T$, and that $\hat{t}$ is a public good for all of the provocateurs. We scale $\hat{t}$ so that it is bounded by zero and one.

---

[15] If multiple equilibria exist and a Type 1 or 3 equilibrium is focal absent any action by the provocateurs, then the provocateurs simply do not need to act. If multiple equilibria exist and a Type 1 or 3 equilibrium becomes focal only if the provocateurs achieve a riot of at least size w, then the analysis below goes through entirely.

We use Proposition Two to simplify the analysis, in the following way. If the riot provocateurs achieve a sufficiently large riot themselves, then a Type 1 (riot) equilibrium becomes the unique equilibrium in the expressive rioting game. The condition that defines this situation is

$$t \geq \bar{\bar{l}} \equiv l\left(\hat{x}(1)+1\right) \tag{7}$$

Many values of $\hat{x}(1)$ will satisfy equation (7). However, define w as the smallest $\hat{x}(1)$ that satisfies it: $w \equiv \min \hat{x}(1) \in \left\{\hat{x}(1) \mid t \geq l\left(\hat{x}(1)+1\right)\right\}$. Because the loss function is at least weakly decreasing for all size riots, a riot of size w or greater assures that equation (7) is satisfied. In other words, w is the threshold size of riot by the provocateurs in order to touch off a large riot by the expressive players, for sure, and thereby receive $\$$.

The utility for each provocateur in the game is a function $V: \hat{T} \times \hat{X} \times \hat{C} \rightarrow \mathbb{R}$.

Specifically, we assume $V(\hat{x}(1), \hat{c}; \hat{t}) = I\hat{t} - \hat{c}l\left(\hat{x}(1)\right)$,where I is an indicator function that takes the value 1 if $\hat{x}(1) \geq w$ and 0 otherwise. Thus, a provocateur receives $V\left(\bullet,\bullet,\bullet\right) = \hat{t}$ if she does not riot but the other provocateurs nonetheless reach or exceed a threshold riot. She receives $V\left(\bullet,\bullet,\bullet\right) = \hat{t} - l\left(\hat{x}(1)\right)$ if she riots and the provocateurs collectively reach or exceed a threshold riot. She receives $V\left(\bullet,\bullet,\bullet\right) = -l\left(\hat{x}(1)\right)$ if she riots but the provocateurs fail to reach a threshold riot. And, she receives $V\left(\bullet,\bullet,\bullet\right) = 0$ if she does not participate and the other provocateurs fail to reach a threshold riot.

If the provocateurs employ strategy $\hat{\mathbf{s}}(1) > 0$ (the probability of rioting in the provocation game) then the probability of reaching or exceeding w is

$\pi(w) \equiv 1-\sum_{k=0}^{w-1} p\left(k \mid m\tilde{\sigma}(1)\right)$. For a random player, the expected utility of choosing action $\hat{c}$ when all other provocateurs in the game are expected to behave according to strategy $\tilde{\sigma}(1) > 0$ is

$$\overline{V}\left(c \mid \tilde{\sigma}(1)\right) = \sum_{k=0}^{\infty} p\left(k \mid m\tilde{\sigma}(1)\right) V\left(\hat{x}(1), \hat{c}\right)$$

$$= \begin{cases} \hat{t}\pi(w) \text{ if } \hat{c} = 0 \\ \hat{t}\pi(w-1) - \sum_{k=0}^{\infty} p\left(k \mid m\tilde{\sigma}(1)\right) l(k+1) \text{ if } \hat{c} = 1 \end{cases} \tag{8}$$

Equilibria

Because $\tilde{C}$ is finite and V bounded, it again follows from Theorem 3 in Myerson 1997a that at least one equilibrium exists in this Poisson game. Again we characterize the equilibria through examination of incentive compatibility constraints. As before, the strong symmetry condition imposed by population uncertainty allows only the same three classes of equilibria.

Type 1 Equilibrium (All players riot: $\tilde{\sigma}(1) = 1$)

In a model very similar to this one but without population uncertainty, Palfrey and Rosenthal consider pure strategy equilibria in which exactly w players contribute to the public good (equivalent to rioting in the provocateur game) while the remaining players do not. These are the only pure strategy equilibria in which the threshold is reached, in the game they consider. Such equilibria are not possible in this Poisson game, due to population uncertainty. However, a very different pure strategy equilibrium can exist in the Poisson game. In this equilibrium, all the provocateurs in the game riot with certainty, and receive in expectation net positive benefits from doing so. As we discuss below, the equilibrium requires rather stringent conditions if it is to exist.

The equilibrium requires

$$\bar{V}\left(1 \mid \tilde{\sigma}(1)\right) \geq \bar{V}\left(0 \mid \tilde{\sigma}(1)\right)$$

$$\Rightarrow \hat{\tau}\pi\left(w-1\right) - \sum_{k=0}^{\infty} p\left(k \mid m\right) l\left(k+1\right) \geq \hat{\tau}\pi\left(w\right)$$

$$\Rightarrow p\left(w-1 \mid m\right) \geq \frac{\sum_{k=0}^{\infty} p\left(k \mid m\right) l\left(k+1\right)}{\hat{\tau}} \qquad (9)$$

The left hand side (lhs) of equation (9) is the probability that a random provocateur is "pivotal," that is, the probability that w will not be achieved if he does not riot but will be if he does. The right hand side (rhs) is the familiar expected loss function, divided by the value to the provocateurs of the public good they will receive if they succeed in sparking mass violence. The left hand term is concave in m. It reaches a maximum at $m = w-1$ and $w-2$ if m is an integer, or at one of those values if m is not an integer (see Mood, Graybill, and Boes Theorem 8 p. 98). The right hand term is strictly decreasing in m, as noted in the proof of Lemma 1. Hence, equation (9) can be satisfied in 3 different ways:

1) lhs = rhs at two points, $\underline{m}$ and $\overline{m}$, and the equilibrium exists for values of $m \in \left[\underline{m}, \overline{m}\right]$ (this is illustrated in Figure 3).

2) lhs = rhs at one point, a tangency point, and the equilibrium exists only at the value of m at the tangency point.

3) lhs = rhs at one point, and lhs > rhs at all greater m. The equilibrium exists for all m greater than or equal to the m at which lhs = rhs.

Cases 2 and 3 are obviously quite special. Case 2 is a knife-edge equilibrium. Case 3 requires the expected cost of rioting to decline extraordinarily precipitously above $w-1$.[16] Accordingly we focus on Case 1.

Focusing on Case 1, equation 9 indicates that a Type 1 equilibrium can exist only under stringent conditions. First, m must be neither "too large" nor "too small," relative to $w-1$. If m is too small relative to $w-1$, the equilibrium cannot exist because the cost of rioting will be too high, even if all riot. If m is too large, the equilibrium cannot exist because the probability of being pivotal will be too small – in other words, free riding kills the equilibrium. Second, equation 9 indicates that a Type 1 equilibrium in the provocation game can exist only if w is rather small. This is because the pivot probability declines quickly, even if $m = w-1$ (so the pivot probability attains the highest possible value given w). For example, $p(5|5)=.175$ and $p(10|10)=.125$ but $p(100|100)=.040$. The need to keep w manageable means that in the expressive rioting game t (the participatory value of rioting to expressive rioters) must be fairly close to the ceiling cost $\bar{t}$ to begin with. In other words, the situation must be a "tinder box." Third, if w is not large, then equation (9) implies that the expected loss from rioting must decline quite quickly. The "resilience" of the police in the face of a riot must be small. Accordingly, the provocateurs will have to seek places or times when the police response is weak.

[16] This might occur if, for example, the loss function $l(\bullet)$ were a step function, with the step at w, and the value of the lower step equal to zero.

Type 2 Equilibrium (No players riot: $\tilde{s}(1) = 0$)

This equilibrium requires that

$$\overline{V}\left(0 \mid \tilde{s}(1)\right) \geq \overline{V}\left(1 \mid \tilde{s}(1)\right)$$

$$\Rightarrow 0 \geq -l(1) \text{ if } w > 1$$

or

$$\Rightarrow \hat{t} \leq \bar{l} \text{ if } w = 1 \tag{10}$$

In other words, a no-riot equilibrium can always exist, except in the exceptional case in which $w = 1$. This suggests that the provocateurs themselves face a coordination problem, assuming type 1 or 3 equilibria also exist. However, solving this problem for a small group of strategically minded players may not be difficult.

Type 3 Equilibrium (Players riot probabilistically: $\tilde{s}(1) \in (0,1)$)

The equilibrium requires that $V\left(\tilde{s}(1) \mid s(1)\right) \geq V\left(1 \mid \tilde{s}(1)\right)$ and that $V\left(s(1) \mid \tilde{s}(1)\right) \geq V\left(0 \mid \tilde{s}(1)\right)$. Together these imply:

$$p\left(w-1 \mid m\tilde{s}(1)\right) = \frac{\sum_{k=0}^{\infty} p\left(k \mid m\tilde{s}(1)\right) l\left(k+1\right)}{\hat{t}} \tag{11}$$

which is similar to equation (9). However, equation (11) is an equality, so the mixed strategy equilibrium can exist only if the pivot probability exactly equals the expected cost of rioting. This condition is even more stringent than those needed for a Type 1 equilibrium. However, note that equation (11) implies that the mixed strategy equilibrium can exist for values of m that are too large (relative to $w-1$) for the pure strategy equilibrium to exist.

An Example

We continue with the example from the previous section, that is, where $l\left(x(1)\right)=\dfrac{1}{\boldsymbol{a}+(x(1))\boldsymbol{b}}$. To calculate w, we invert $t=\bar{\bar{l}}=\dfrac{1}{\boldsymbol{a}+\boldsymbol{b}(w+1)}$ to solve for the threshold w: $w=\dfrac{1}{t\boldsymbol{b}}-\dfrac{\boldsymbol{a}+\boldsymbol{b}}{\boldsymbol{b}}$. Thus, if $\boldsymbol{a}=1.49$ and $\boldsymbol{b}=.01$, w = 234 if t = .3; w = 100 if t = .4; and w = 50 if t = .5. Note that w is independent of the expected size of the crowd of expressive players, n. Instead, n effects the expected size of the expressive riot if it is achieved.

It is very easy to construct examples in which a Type 1 equilibrium does not exist. It is more difficult to construct cases in which one does exist. One case is shown in Figure 3. In this case, $\boldsymbol{a}=1$ but $\boldsymbol{b}=.9$. The high value of indicates that the police presence is very brittle – it is easily overwhelmed by a riot, either due to incompetence or tacit collusion in the face of any sizeable disturbance. Absent any disturbance by the provocateurs, the ceiling cost of rioting for the affective players is .53, as it is for the provocateurs. We assume the value of t for affective players is a low .085. Because the police presence is so brittle, however, a riot of w = 10 is sufficient to reduce the effective ceiling cost to this level. What range of m can support a Type 1 equilibrium among the provocateurs? The answer is shown in Figure 3. The green bars show the pivot probability for each value m between 1 and 20. The black curve is the expected loss function, divided by the value of the public good (assumed to be .9). If the value of the pivot probability lies above the expected loss function, the expected population size of provocateurs supports a Type 1 equilibrium. As indicated, the values supporting the equilibrium are 9-12, inclusive. So if, for example, $m = 10$, all the provocateurs riot with $\sigma(1) = 1$ and receive positive expected utility from doing so.

IV. Interpreting Riots and Collective Violence

We turn now to an example of mass violence in which strategic behavior by riot provocateurs taps into issues with high affective salience for specific groups to bring about riot events. The violence in this case has been attributed primarily to single factors without fully explaining the interplay of outcome-oriented and act-oriented goals.

Party Politics and Violence: The BJP and the Rath Yatra of 1990

In late 1990, a series of riots broke out across India after the abrupt and incomplete end of a religious pilgrimage, or rath yatra, undertaken by L. K. Advani, one of the top leaders of the Hindu nationalist Bharatiya Janata Party (BJP). This journey was precipitated by a series of political decisions taken by the minority government of the National Front, led by Prime Minister V. P. Singh. In July 1990, the government had suffered the defection of a major backward-caste leader. To attempt to compensate for the possible loss of low-caste electoral support, Singh had announced the implementation of a long-dormant affirmative action plan at the national level that would set aside places in government employment and higher education for members of the "Other Backward Classes," who comprised over half the national population.

This policy announcement was highly problematic for the BJP. On the one hand, their electoral support came primarily from high-caste Hindus in the upper and middle classes, and these groups were staunchly opposed to affirmative action. On the other hand, even the BJP had to attempt to appeal to the vast OBC electorate, and they could not afford to alienate these voters by denouncing the proposed legislation. Indeed, their platform, like almost every other party's platform, supported implementation of these policies. Therefore, the BJP chose an end run around the policy; they emphasized the unity of Hindu interests through religious symbols in order to counter the potential cleavages created by caste-based policies.

One of the BJP's allied social-religious organizations, the Vishwa Hindu Parishad (VHP), had made the demolition of a mosque at Ayodhya its major priority. The VHP asserted that it had evidence to show that the mosque had been built on the site of the birthplace of Rama, and that a Hindu temple that had preceded the mosque had been demolished during one of the Muslim invasions of the middle ages. During the first half of 1990, the VHP had been attempting to organize a grass-roots movement in Uttar Pradesh to build a new temple where the mosque currently stood, but it had not been very successful.

Barely a month after V. P. Singh's announcement on reservations, BJP leader L. K. Advani announced that he would undertake a rath yatra, or a chariot pilgrimage, from Somnath in Gujarat to Ayodhya, to mobilise public opinion
and "politically educate" the people about the Ram Janambhoomi issue" (*Indian Express*, September 13, 1990). The pilgrimage would wind its way through nine states and make frequent stops at pre-arranged locations.

Political observers assumed that this decision was politically motivated. As the *Hindustan Times* reported, The three-day biggest-ever conclave of the BJP leaders here [Bhopal] has made one thing very clear: the Hindu card will be used vigorously against PM V. P. Singh's caste card." ^Å The entire conclave was centered on Hinduism." The BJP leadership, which feels isolated after the bombshells thrown by Mr. V. P. Singh, is now convinced that only the Hindu card could work in the wake of mid-term elections (September 19, 1990).

Advani's procession received enormous publicity and proceeded with great fanfare. Although the crowds in rural areas were initially described as "sparse," the turnout in the cities was larger and more enthusiastic (*Times of In*dia, September 27, 1990). As Advani proceeded on his 6,000 mile journey, communal tensions were heightened along the route. At the same time, the VHP began another parallel procession in southern India and in other places through which Advani would not be passing, and communal tension was rising along that route as well (Jaffrelot 1998).

Advani's journey was accompanied by increasing publicity, crowds, and conflict. although the violence was relatively sporadic. The organizations most responsible for the grass-roots mobilization, the VHP and the more militant RSS, were concerned more with generating large crowds for the roadside than with fomenting riots. Nevertheless, by the time Advani reached Delhi, three weeks into the procession, there was considerable apprehension about what might happen as he approached Ayodhya. The police, while always on alert, were reluctant to wade into the crowds. At the same time, the governments of each state were either sympathetic to the BJP or afraid of making Advani into a sympathetic figure, and so refrained from encouraging a large police presence. At the beginning of the trip Advani was primarily drawing large crowds in BJP strongholds, but by the later stages he was attracting enormous attention everywhere. Finally, on October 23, the day before Advani was scheduled to enter UP, the central government, with the support of the allied Bihar state government, arrested Advani and ended the procession.

The BJP immediately called for a Bharat bandh, or national shutdown strike, which was only partially observed in some parts of the country but was extremely successful in the states through which the procession had passed. In these states, the strike degenerated into communal violence in which dozens were killed in less than two days. Newspaper accounts stressed, sometimes with helpful maps, that the riots were more prevalent and more deadly in those localities through which Advani had passed, and especially those in which he had halted for rallies. The strike actions were led by the RSS and the VHP, and many local newspaper reports also attributed the riots to them as well. The riots largely involved religious conflict between Hindus and Muslims, as opposed to the caste-based violence that had followed the affirmative action policy announcement of V. P. Singh. Meanwhile, in New Delhi, the BJP formally withdrew its support of the V. P. Singh government, which fell less than two weeks later.

Most accounts, whether by journalists or scholars, emphasize the political strategic nature of Advani and the BJP's decision to undertake the Rath Yatra. Nandy and his coauthors assert that "the politically alert “ saw the Yatra as the beginning of the BJP's election campaign" (1995: 40). Even scholars more directly concerned with religious identity as distinct from politics saw Advani's decision as wholly political:

Since the agitation around the reservation issue imperiled the Hindu agenda of the VHP/BJP/RSS, Lal Kishan Advani, the leader of the BJP, decided to start a ritual procession that would pass through ten states . . . Advani's posturing as Rama . . . took place in the context of this campaign" (van der Veer 1994: 5).

It is difficult to disagree with this explanation, especially when BJP leaders, "before throwing themselves fully into the Ayodhya movement, openly talked about their political purposes as 'playing the Hindu card' for electoral advantage" (Brass 1997: 269). The timing of the BJP's decision to make building the Ram temple at Ayodhya the highest priority; the decision of Advani, a Sindhi Hindu who claimed to be "spiritually a Sikh" (Nandy et al. 1995: 40), to dress in the clothes of an ascetic and decorate his Toyota van like a lotus-painted chariot; and the selection of Somnath as the starting point all suggest a strategy of crass electoral gain.

Yet, despite its obviously political instrumental benefits, and despite the explosive nature of the strategy, Advani's procession also appeared to generate sincere affective behavior. Certainly the VHP and the RSS were organizing support for Advani. But the crowd of nearly 100,00 that greeted him in Delhi was not entirely manufactured by these organizations, and the accounts of the violence that followed Advani's arrest made it clear that much of it was not planned in advance.

Our model allows us to account for both the strategic, outcome-oriented behavior undertaken by the BJP and its organization allies as well as the more spontaneous, affectively-oriented behavior exhibited by rank and file participants. The BJP had an explicit political agenda: it needed to divert attention from the National Front's affirmative action policy initiative because caste-conscious policies split its Hindu electorate. The rath yatra allowed the party to put that electorate back together under the banner of Hindu solidarity. It also provided a coordination mechanism through the pilgrimage itself; participants who came out to view the procession saw others and got a sense of how many others were motivated by the same issues. When Advani was arrested, it provided a second coordination point, one around which affective participants could riot with a reasonable expectation of how many others would be rioting with them. Unlike the VHP's earlier unsuccessful efforts in UP to organize Hindus around the mosque issue, the procession was a highly salient coordination mechanism.

The model directs us to look for specific types of factors that precipitate a riot. First, there have to be a small but critical group of riot provocateurs. For the riots that followed Advani's arrest, the RSS and VHP largely provided these participants. Second, the provocateurs must provide a focal point that causes participants with high value from participation to be mobilized. Both the procession itself and the arrest served this function: the procession because it assembled like-minded participants and provided information on the size of the potential crowd, and the arrest because it served the immediate function of providing an event which the provocateurs could use to begin the riot itself. Finally, the provocateurs must believe that they will achieve their desired outcome, in this case the destabilization of the incumbent government. The pilgrimage, and the increasing crowds that accompanied it, suggested that there was considerable support for the BJP and its Hindu nationalist strategy. When Advani was arrested, the potential for rioting was perceived to be relatively high, and the success of the strategy was borne out within the month, as the National Front government was forced to resign.

V. Conclusion

The models in this paper formalize ideas implicit in many writings about riots and mass violence. Accordingly, they provide a framework for organizing many well-known observations about riots and mass violence. More importantly, they integrate two hitherto disjoint strands in the literature on mass violence, the "hatred" approach and the "manipulation" approach. When the two are brought together in an explicit and careful fashion, the result is a set of new but largely intuitive and quite plausible conjectures about the circumstances under which manipulation can occur. We see this at the principal substantive contribution of the paper.

The Poisson game approach used here to study expressive riots and riot provocation can be applied to many other situations involving thresholds, cascades, and mass collective action. Examples include mass protest movements and social movements (MacAdam), transformation of ethnic norms and the development of reputational cascades (Kuran), and language politics and the strategic construction of identities (Laitin). Moreover, the interaction of the provocation and expressive games may point to similar dynamics elsewhere.

References

Anderson, Benedict. 1984. *Imagined Communities*. Verso.

Brass, Paul. 1997. *Theft of an Idol: Text and Context in the Representation of Collective Violence.* Princeton: Princeton University Press.

Connor, Walker. 1994. *Ethnonationalism*. Princeton: Princeton University Press.

D'Aspremont, Claude and Louis-Andre Gerard-Varet. 1979. "Incentives and Incomplete Information," *Journal of Public Economics* 11:25-45.

Granovetter, Mark. 1978. "Threshold Models of Collective Behavior," *American Journal of Sociology* 83(6):1420-1443.

Kuran, Timur. 1991. "Now Out of Never: The Element of Surprise in the East European Revolution of 1989," *World Politics* 44 (October):7-48.

Laitin, David. 1998. *Identity in Formation*. Cornell: Cornell University Press.

Lohmann, Susanne. 1994. "Information Aggregation Through Costly Political Action," *American Economic Review* 84:518-530.

Mood, Alexander M., Franklin A. Graybill, and Duane C. Boes. 1974. *Introduction to the Theory of Statistics* (third edition). New York: McGraw-Hill Publishing Company.

Myerson, Roger B. 1998. "Large Poisson Games." Center for Mathematical Studies in Economics and Management Science Discussion Paper No. 1189. Northwestern University (Final Revised Version May 1998).

-----. 1997 a. "Population Uncertainty and Poisson Games." Center for Mathematical Studies in Economics and Management Science Discussion Paper No. 1102 R. Northwestern University (Revised June 1997).

-----. 1997 b. "Extended Poisson Games and the Condorcet Jury Theorem." Center for Mathematical Studies in Economics and Management Science Discussion Paper No. 1103. Northwestern University (Revised June 1997).

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.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior.* New York: W.W. Norton & Company.

Smith, Anthony. 1981. *Ethnic Revival in the Modern World*. Cambridge: Cambridge University Press.

Appendix

Proof of Lemma 1

(Sketch). Expected loss $\sum_{k=0}^{\infty} p(k \mid n)\, l(k+1)$ is continuous in *n*, converges to $\underline{l}$ as $n \to \infty$, and is strictly decreasing in *n*. The lemma then follows. To see that

$\lim_{n\to\infty} \sum_{k=0}^{\infty} p(k \mid n)\, l(k+1) = \underline{l}$, first recall that the variance *n* = *n*. So if *n* is very large,

$\sum_{k=0}^{\infty} p(k \mid n)\, l(k+1) \approx \sum_{k=n-10\sqrt{n}}^{n+10\sqrt{n}} p(k \mid n)\, l(k+1) \approx \underline{l} \sum_{k=n-10\sqrt{n}}^{n+10\sqrt{n}} p(k \mid n) \approx \underline{l}$. To see that expected loss is strictly decreasing in *n*, first note that any $\frac{\partial p(k \mid n)}{\partial n} = p(k \mid n)\left(\frac{k-n}{n}\right)$. So at any *n*, increasing *n* puts less probability weight on all loss terms in the sum in which $k < n$ and greater weight on all terms in which $k > n$. But by assumption the latter must contain terms that are strictly smaller than some terms in the former, and none greater than any term in the former. So the value of the sum must fall as *n* increases, at any value of n. Q.E.D.

Proof of Lemma 2

Any expected loss function $\sum_{k=0}^{\infty} p(k \mid n s(1))\, l(k+1)$, with $s(1) \in (0,1)$, is exactly equivalent to a loss function $\sum_{k=0}^{\infty} p(k \mid m)\, l(k+1)$ where $m = n s(1)$. Lemma 1 establishes the existence of *n** for this loss function. Q.E.D.

Proof of Proposition Two

From the specified loss function and the definition of floor costs, it must be the case that

$\lim_{\hat{x}(1)\to\infty} l\left(x(1)+\hat{x}(1)\right)=\underline{l}$ for all $x(1)>0$. The proposition is immediate. Q.E.D.

Figure 1. An Expected Loss Function

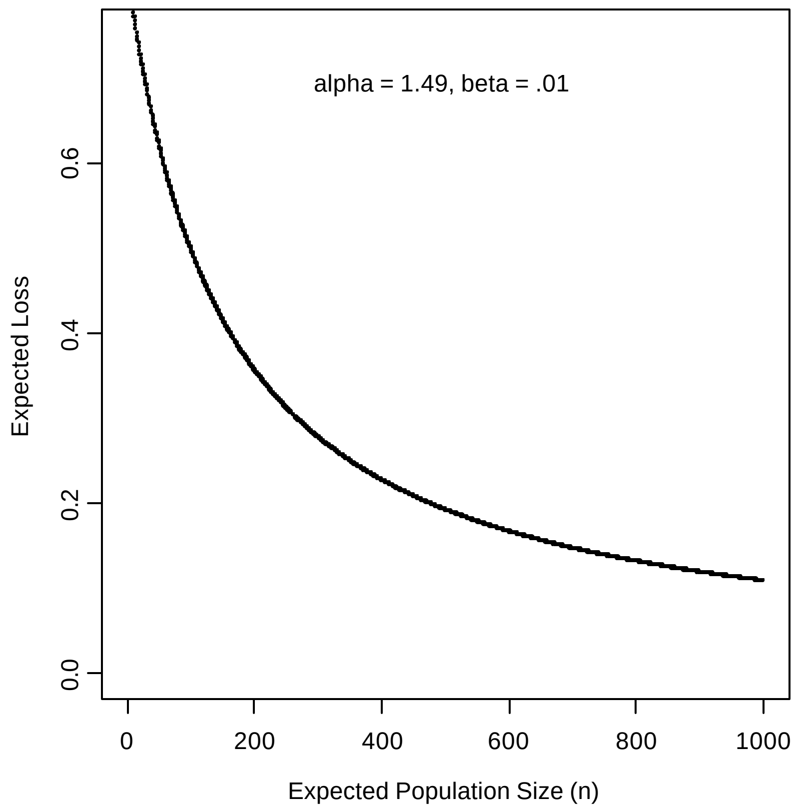

Figure 2. Equilibria in the Example

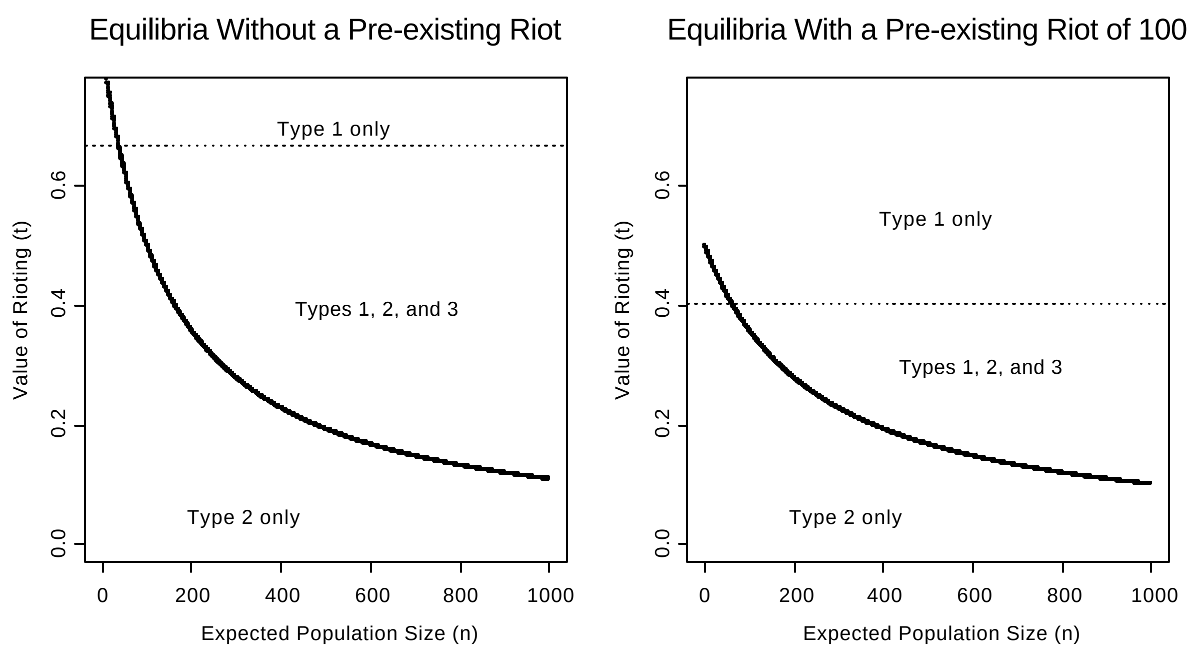

Figure 3. Loss Function and Pivot Probability in the Provocation Game