

Modeling Civil Violence: An Agent-Based Computational Approach

Joshua M. Epstein, John D. Steinbruner, Miles T. Parker*

Center on Social and Economic Dynamics
Working Paper No. 20
January 2001

Abstract

This working paper presents an agent-based computational model of civil violence. We present two variants of the Civil Violence Model. In the first, a central authority seeks to suppress decentralized rebellion. In the second, a central authority seeks to suppress communal violence between two warring ethnic groups.

* The authors thank Ross Hammond for research assistance. For comments, they thank Robert Axtell, William Dickens, Steven Durlauf, Carol Graham, Peyton Young, and members of the MacArthur Network on Inequality and Social Interaction. For financial support, the authors thank the National Science Foundation and the John D. and Catherine T. MacArthur Foundation.

OVERVIEW

This Working Paper presents an agent-based computational model of civil violence. The model is very much a work in progress, and this preliminary report is intended to stimulate discussion, garner feedback, and foster refinement. For an introduction to the agent-based modeling technique, see Epstein and Axtell (1996). We present two variants of the Civil Violence Model. In the first, a central authority seeks to suppress decentralized rebellion. Where we use the term “revolution,” we do so advisedly, recognizing that no political or social order is represented in the model. Perforce, neither is the overthrow of an existing order, the latter being widely seen as definitive of revolutions properly speaking. *The dynamics of decentralized upheaval, rather than its political substance*, is the focus here. In the second model, a central authority seeks to suppress communal violence between two warring ethnic groups. And, as in Model I, we are interested in generating certain characteristic phenomena and core dynamics; we do not purport to reconstruct any particular case in detail, though—as discussed in Appendix B—that is an obvious long-term objective.

CIVIL VIOLENCE MODEL I:

GENERALIZED REBELLION AGAINST CENTRAL AUTHORITY

This model involves two categories of actors. "Agents" are members of the general population, and may be actively rebellious or not. "Cops" are the forces of the central authority, who seek out and arrest actively rebellious agents. Let us describe the agents first. As in all agent-based models, they are heterogeneous in a number of respects. The attributes and behavioral rules of the agents are as follows:

The Agent Specification

First, in any model of rebellion, there must be some representation of political grievance. Our treatment of grievance will be extremely simple, and will involve only two highly idealized components which, for lack of better terminology, we will call Hardship and Legitimacy. Their definitions are as follows:

H: The agent's perceived hardship. In the current model, this is exogenous. We assume it to be heterogeneous across agents. Lacking further data, each individual's value is simply drawn from $U(0,1)$, the uniform distribution on the interval $(0,1)$. Of course, perceived hardship alone does not a revolution make. As noted in the Russian revolutionary journal, *Narodnaya Volya*, "No village ever revolted merely because it was hungry." (Quoted in Kuran, 1989 and deNardo, 1985). Another crucial factor is:

L: The perceived legitimacy of the regime, or central authority. In the current model, this is exogenous, is equal across agents, and in the runs discussed below, will be varied over its arbitrarily defined range of 0 to 1.

The level of *grievance* any agent feels toward the regime is assumed to be based on these variables. Of the many functional relationships one might posit, we will assume:

$$\mathbf{G} = \mathbf{H}(1-\mathbf{L}).$$

Grievance is the product of perceived hardship (H) and perceived "illegitimacy," if you will (1-L). The intuition behind this functional form is simple. If legitimacy is high, then hardship does not induce political grievance. For example, the British government enjoyed unchallenged legitimacy ($L=1$) during World War II. Hence, the extreme hardship produced by the blitz of London did not produce grievance toward the government. By the same token, if people are suffering (high **H**), then the

revelation of government corruption (low **L**) may be expected to produce increased levels of grievance.

Of course, the decision to rebel depends on more than one's grievance. For example, some agents are simply more inclined to take risks than others. Accordingly, we define

R: The agent's level of risk aversion. Heterogeneous across agents, this (like **H**) is assumed to be uniformly distributed across agents. Each individual's level is drawn from $U(0,1)$ and is fixed for the agent's lifetime.

All but the literally risk neutral will estimate the likelihood of arrest before actively joining a rebellion. This estimate is assumed to increase with the ratio of cops to already rebellious--so-called "active"--agents within the prospective rebel's vision. To model this, we define:

v: The agent's vision. This is the number of lattice positions (north, south, east, and west of the agent's current position) that the agent is able to inspect. It is exogenous and equal across agents. As in most agent-based models, vision is limited; information is *local*. Letting $(C/A)\mathbf{v}$ denote the cop-to-active ratio within vision **v**, we assume the agent's estimated arrest probability **P** to be given by

$$\mathbf{P} = 1 - \exp[-k(C/A)\mathbf{v}].$$

The constant k is set to ensure a plausible estimate (of $P=.9$) when $C=1$ and $A=1$. Notice that A is always at least 1, since the agent always counts himself as active when computing P . He is asking, "How likely am I to be arrested if I go active?" This functional form makes P a legitimate probability density. Again, the intuition is very simple. Imagine being a deeply aggrieved agent considering throwing a rock through a bank window. If there are 10 cops at the bank window, you are much more likely to be arrested if you are the first to throw a rock ($C/A=10$) than if you show up

when there are already 29 rock-throwing agents ($C/A=1/3$). For a fixed level of cops, the agent's estimated arrest probability falls the more actives there are. This simple idea will play an important role in the analysis.

Clearly, in considering whether or not to rebel, a risk neutral agent won't care what the estimated arrest probability is, while a risk averse agent will. It will therefore prove useful to define

$N = \mathbf{RP}$: The agent's net risk--the product of his risk aversion and estimated arrest probability. (This can be considered the special, $\alpha=0$, case of $N = RPJ^\alpha$ where J is the jail term, discussed below). These ingredients in hand, the agent's behavioral rule is summarized in the Table below.

Table. Agent State Transition

<u>State</u>	<u>(G - N)</u>	<u>State Transition</u>
Q	$>T$	$Q \Rightarrow A$
Q	$\leq T$	$Q \Rightarrow Q$
A	$>T$	$A \Rightarrow A$
A	$\leq T$	$A \Rightarrow Q$

If, for an agent in state Q, the difference G-N exceeds some non-negative threshold T--which could be zero--then that quiescent agent goes active. Otherwise, he stays quiescent. If, for an agent in state A, the difference exceeds T, then that active agent stays active. Otherwise, he goes quiescent. In summary, the agent's *simple local rule* is:

Agent Rule A: If $G-N > T$ be Active; Otherwise, be Quiet.

This completes the agent specification.

Bounded Rationality

It is natural to interpret this rule as stipulating that the agent take whichever binary action (Active or Quiescent) maximizes expected utility

where, in the spirit of Kuran (1989), $G-N$ is the expected utility of publicly expressing one's private grievance, and T is the expected utility of not expressing it (i.e., of preference falsification, in Kuran's terminology). Typically, we set T at some small positive value. Notice, however, that if it takes negative values, like $-G$ (i.e., the frustration level associated with preference falsification equals the grievance level itself), agents may find it rational to rebel knowing that they will suffer negative utility. It's simply worse to "sit and take it anymore." Agents weigh expected costs and benefits, but they are not hyper-rational. One might say (*pace* Olsen, 1965) that individual rationality is "local" also, in the sense that the agent's expected utility calculation *excludes* any estimate of how his isolated act of rebellion may effect the social order. Notice, very importantly, that deterrence is local in this model and depends on the local (individually visible)--not the global--ratio of cops to actives, which is highly dynamic in this spatial model with movement.

The Cop Specification

The cops are much simpler than prospective rebels. Their attributes are as follows:

\mathbf{v}^* : The cop vision. This is the number of lattice positions (north, south, east, and west of the cop's current position) that the cop is able to inspect. It is exogenous and equal across cops. The cops' \mathbf{v}^* need not equal the agent's \mathbf{v} , but will typically be small relative to the lattice size: cop vision is local also. The cops, like the other agents, have one simple rule of behavior:

Cop Rule C: Inspect all sites within \mathbf{v}^* and arrest a random Active Agent.

Cops never defect to the revolution in this model.

Movement and Jail Terms

Though the range of motion will vary depending on the numerical values selected for \mathbf{v} and \mathbf{v}^* , the syntax of the movement rule is the same for both agents and cops:

Movement Rule M: Move to a random site within your vision.

Regarding jail terms for arrested actives, these are exogenous and set by the user. Specifically, the user selects a value for the maximum jail term, **J_max**. Then, any arrested active is assigned a jail term drawn randomly from $U(0, \mathbf{J_max})$. **J_max** will affect the dynamics in important ways by removing actives from circulation for various durations. However, for the present version of the model with alpha implicitly set to zero (see discussion above), there is no *deterrent* effect of increasing the jail term. Setting alpha to a positive value would produce a deterrent effect. In addition to having no deterrent effect, we also assume agents leave jail exactly as aggrieved as when they entered.

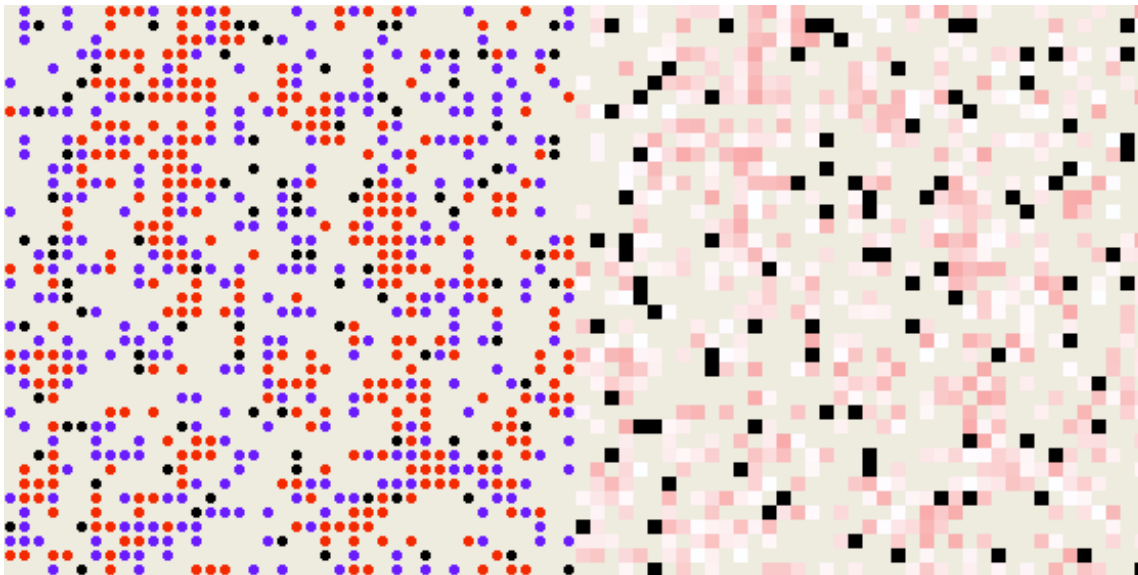
Measurement

It is important to state forthrightly that we make no pretense to *measuring* model variables such as perceived hardship or legitimacy. The immediate question is whether this highly idealized model is *sufficient to generate* recognizable macroscopic revolutionary dynamics of fundamental interest. If not, then issues of measurement are moot. So, the first issue is whether the model produces interesting output. In addition to data generated by the model, run-time visualization of output is very useful. Our graphical strategy is as follows.

Graphics

Events transpire on a lattice. Agents and cops move around this space and interact. We are interested in the dynamics of grievance and--quite separately--in the dynamics of revolutionary action. The point of separating these private and public spheres is to permit illustration of a core point in all research on this topic: public order may prevail despite tremendous private opposition to--feelings of grievance toward--a regime. Given this important distinction between private grievance and public action, we will show two screens (see Figure 1).

Figure 1. Action and Grievance Screens



On the right screen, agents are colored by their *private* level of grievance. The darker the red, the higher the level of grievance. On the left screen, agents are colored by their *public* action: Blue if quiescent; Red if active. Cops are colored black on both screens. Simply to reduce visual clutter, all agents and cops are represented as circles on the left screen and squares on the right. Unoccupied sites are gray colored on both screens.


Runs

To begin each run of the model, the user sets L , J , V , V^* , and the initial cop and agent densities. To ensure replicability of our results, input assumptions for all runs are provided in Appendix A. Agents are assigned random values for H and R , and cops and (initially) quiescent agents are situated in random positions on the lattice. The model then simply spins forward under the Rule set: $\{A, C, M\}$. An agent or cop is selected at random (asynchronous activation) and, under Rule M , moves to a random site within his vision, where he acts in accord with Rule C (if a cop) or A (if an agent). The model simply iterates this procedure until the user quits or some stipulated state is attained. What can one generate in this extremely simple model? We present 5 Runs.

Run 1. Individual Deceptive Behavior

Despite their manifest simplicity, the agents do something we had not anticipated. This behavior is displayed below in two snapshots from Animation 1.

Animation 1. Deceptive Individuals



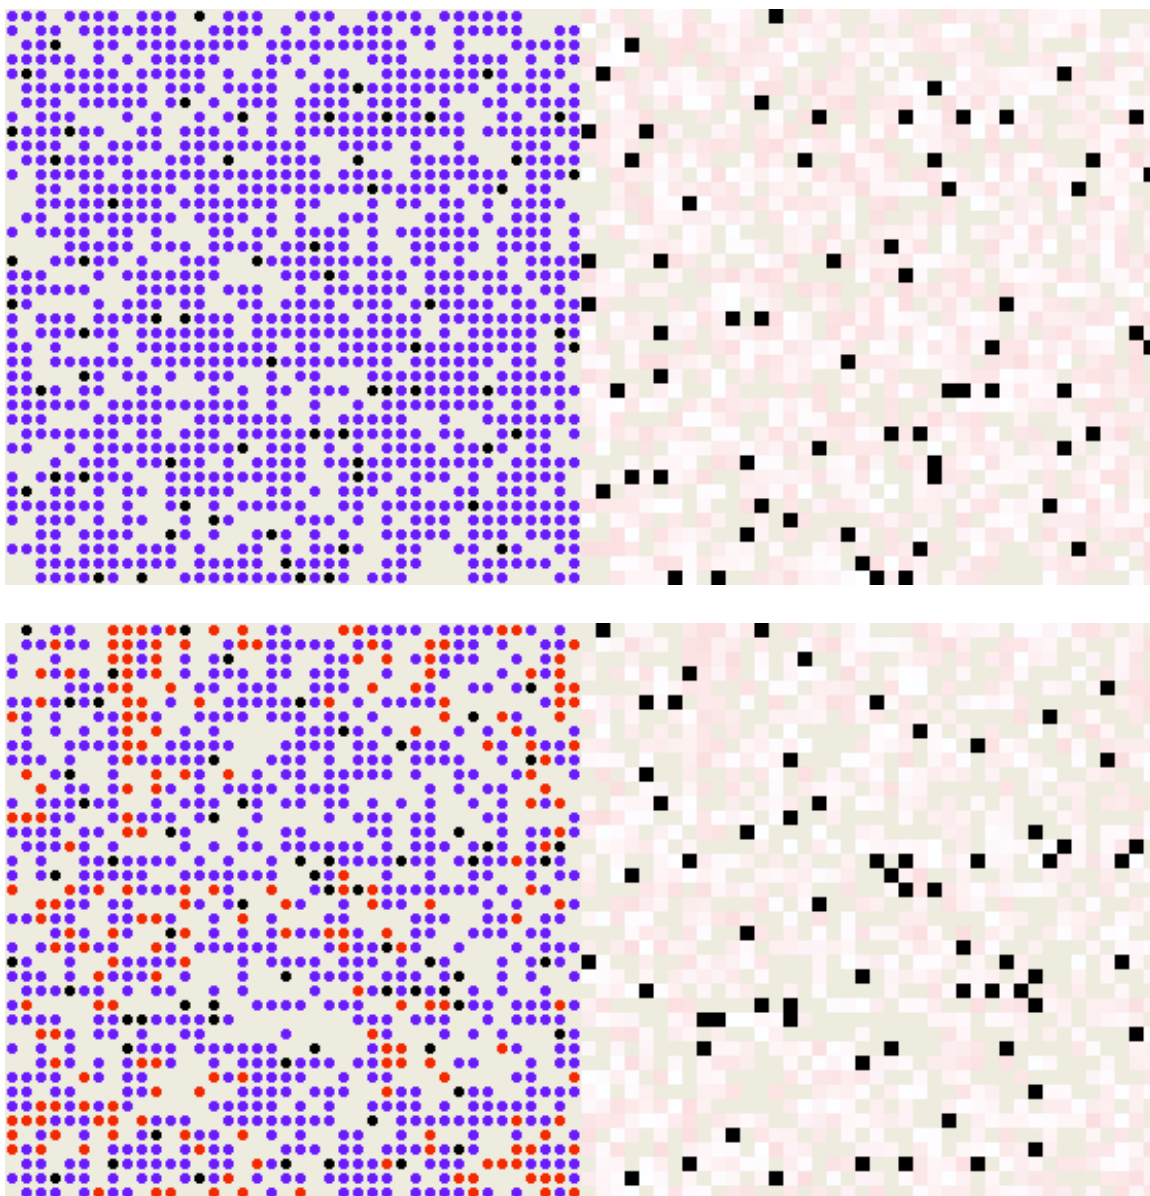
Focus then on the circled agent in the left frame. Though not shown, this agent's grievance level is high. However, there are two cops within his vision and he is Blue (Quiescent). But, he turns Red (Actively rebellious) the moment these cops move away, as shown in the right frame. In short, the

agents exhibit unexpected deceptive behavior: privately aggrieved agents turn Blue (as if they were non-rebellious) when cops are near, but then turn Red (actively rebellious) when cops move away. They are reminiscent of Mao's directive that revolutionaries should "swim like fish in the sea," making themselves indistinguishable from the surrounding population. *Ex post facto*, the behavior is easily understood: the cop's departure reduces the C/A ratio within the agent's vision, reducing his estimated arrest probability, and with it his net risk, N , all of which pushes $G-N$ over the agent's activation threshold, and he turns Red. But it was not anticipated. Moreover, it would probably not have been detected without a spatial visualization; individual deception would not be evident in a time series of total rebels, for example. With both cops *and* agents in random motion, more interesting dynamics emerge.

Run 2. Free Assembly Catalyzes Rebellious Outbursts

With both agents and cops in random motion, it may happen that high concentrations of actives arise endogenously in zones of low cop density. This can depress local C/A ratios to such low levels that even the mildly aggrieved find it rational to join. This catalytic mechanism is illustrated in Animation 2.

Animation 2. Local Outbursts



Random spatial correlations of activists catalyze local outbursts. This is why freedom of assembly is the first casualty of repressive regimes. Relatedly, it is also the rationale for curfews. The mechanism is that local activist concentrations reduce local C/A (cop-to-active) ratios, reducing (via the equation for P above) the risk of joining the rebellion. To be the *first* rioter, one must be either very angry or very risk neutral, or both. But to be the 4000th--if the mob is already big, relative to the cops--the level of grievance and risk taking required to join the riot is far lower. This is how, as Mao Tse Tung liked to say, "A single spark can cause a prairie fire." (Quoted in Kuran, 1989). Coincidentally, the Bolshevik newspaper founded by Lenin was called *Iskra*--the spark! The Russian revolution itself provides a beautiful example of the chance spatial correlation of aggrieved agents. As Kuran (1989; 63) recounts,

On February 23, the day before the uprising, many residents of Petrograd were standing in food queues, because of rumors that food was in short supply. 20,000 workers were in the streets after being locked out of a large industrial complex. Hundreds of off-duty soldiers were outdoors looking for distraction. And, as the day went on, multitudes of women workers left their factories early to march in celebration of Women's Day. The combined crowd quickly turned into a self-reinforcing mob. It managed to topple the Romanov dynasty within four days."

A random coalescence of aggrieved agents depresses the local C/A ratio, quickly emboldening all present to openly express their discontent.


A time series of total rebels is also revealing. It displays one of the hallmarks (Young 1998) of complex systems: punctuated equilibrium. Long periods of relative stability are punctuated by outbursts of rebellious activity.

Figure 2. Punctuated Equilibrium




And indeed, many major revolutions (e.g., East German) are episodic in fact.

The same qualitative pattern of behavior—punctuated equilibrium—persists indefinitely, as shown in Figure 3, which plots the data over some 20 thousand iterations of the model.

Figure 3. Punctuated Equilibrium Persists**Waiting Time Distribution**

Is there any underlying regularity to these complex dynamics? For many complex systems, it turns out to be of considerable interest to study the distribution of waiting times between outbursts above some threshold. In this analysis, we set the threshold at 50 actives. An outburst begins when the number of actives exceeds 50 and ends when it falls below 50. We are interested in the time between the end of one outburst and the start of the next. Sometimes, we have to wait a long time (e.g., 100 periods) until the next outburst. Sometimes, the next outburst is nearly immediate (e.g., a gap of only 2 periods). The frequency distribution of these inter-outburst waiting times, for 100,000 iterations of the model, is shown in Figure 4.

Figure 4. Waiting Time Distribution




In the complexity literature, one often encounters the notion of an “emergent phenomenon.” I have argued elsewhere (Epstein, 1999) that substantial confusion surrounds this term. However, if one defines emergent phenomena simply as “macroscopic regularities arising from the purely local interaction of the agents,” (Epstein and Axtell, 1996), then this waiting time distribution surely qualifies. It was entirely unexpected, and would have been quite hard to predict from the underlying rules of agent behavior. For instance, Figure 4 suggests a Weibull or Lognormal distribution. While rigorous identification is a suitable topic for future research, these data are clearly not uniformly

distributed. But all distributions used in defining the agent population—the distribution of Hardship and Risk Aversion—are uniform. In a uniform waiting time distribution, one is just as likely to wait 100 cycles as 50; that is not the case in our model, at least for these parameters. The mean of these data—the average duration between outbursts—is 60.¹ Clearly, most of the probability density is concentrated around this value: this means that we are much more likely to get successive outbursts within 60 cycles than within 100.

If one is willing to truncate this distribution—throwing out the most high frequency events (waiting times less than 30 cycles), the remaining distribution is well fit by the negative exponential. In particular, the logged truncated data are shown on Figure 5.

Figure 5. Logged Data (Truncated)




¹ The standard deviation is 55.

Ordinary Least Squares regression yields an R-squared of 0.98 with slope -0.07 and intercept 3.5. The negative exponential distribution is ubiquitous in the analysis of failure rates—the rates at which electrical and mechanical systems break down. It would be interesting if “social breakdowns” followed a similar distribution. Another obvious issue is the sensitivity of this distribution to variations in key parameters. For example, how would an increase in the jail term deform the distribution? One might conjecture that, by removing rebellious agents from circulation for a longer period, increasing the jail term would “flatten” the distribution and raise its mean. All of these issues could be fruitfully explored in the future. For the moment, the core point is that a powerful statistical regularity underlies the model’s punctuated equilibrium dynamics.

Outburst Size Distribution

A second natural statistical topic is the size distribution of rebellious outbursts. To study this, we use the same parameterization as above and adopt the same threshold: 50 actives. But there are numerous ways to measure outburst size for statistical purposes. For example, imagine a flare-up lasting five days with the following number of actives per day: 60, 100, 120, 95, and 80. This outburst has a peak active level of 120. It has average (daily) activation of 100, and total activation (the sum) of 500. The size distribution of outbursts using the total activation measure is shown in Figure 6.

Figure 6. Total Activation Distribution




The mean and standard deviation are, respectively, 708 and 230. The distributions using the peak and average data are qualitatively similar. As in the case of the waiting time distribution, one could conduct further analysis to identify the best fit to Figure 6. But the point to emphasize here is not which distribution is best, but that some macroscopic regularity emerges. A major strength of agent models is that they generate a wealth of data amenable to statistical treatment.

A Ripeness Index

Turning to another topic, we often speak of a society as being "ripe for revolution." In using this terminology, we have in mind a high level of tension or private frustration. Does the present model allow us to quantify

this in an illuminating way? As a first cut, we noted earlier that society can be bright red on the right screen (indicating a high level of grievance) while being entirely Blue on the left (indicating that no one is expressing, or "venting," their grievance). So, if this combination of high average grievance \bar{G} on the right and high frequency of Blues \bar{B} on the left were the best indicator of high tension, a reasonable "ripeness" index would be simply their product: $\bar{G} \bar{B}$. This, however, ignores the crucial question, *why* are agents Blue? If they are inactive simply because they are risk averse and have no inclination to go active, then they are not truly frustrated in the inactive Blue state. So, for fixed \bar{G} and \bar{B} a good tension index should increase as average risk aversion falls (more agents want to act out, but are nonetheless staying blue). Hence, a better simple measure is: $\bar{G} \bar{B} / \bar{R}$, where \bar{R} is average risk aversion. In Figure 7, we plot this against a curve of actives designed to exhibit high volatility. It is clear that a buildup of tension precedes each outburst.

Figure 7. Tension and Actives



In fact, on close inspection, our tension curve dips just at the outset of major explosions. The true warning indicator is this pronounced dip. Notice, moreover, that—like any good indicator—it is usefully insensitive. It does

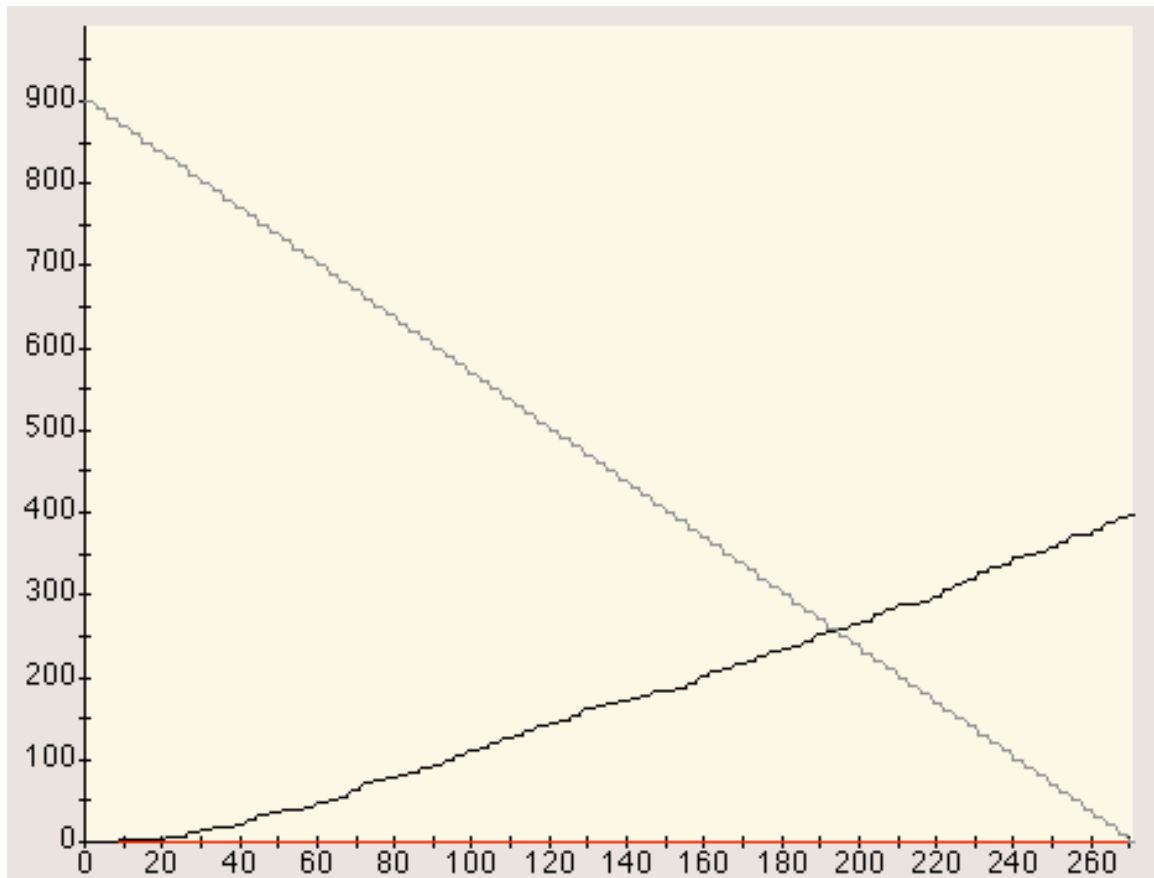
not give false alarms (i.e., there are no dips) for the minor eruptions at periods 20 or 45 for instance.

We turn now to a comparison of two runs involving reductions in legitimacy. Both runs begin with legitimacy at a high level. In the first, we execute a large absolute reduction (from $L=.9$ to $L=.2$) in legitimacy, but in *small* increments (of a percent per cycle). In the second, we reduce it far less in absolute terms (from $L=.9$ to $L=.7$), but do so one jump. Which produces the more volatile social dynamics, and why?

Runs 3 and 4. Salami Tactics of Corruption.


Figure 8 shows the results when we reduce legitimacy in small increments. It displays three curves. The downward sloping upper curve plots the steady incremental decline in legitimacy over time. (To make these graphs clear, we actually plot $1000L$). The horizontal red curve just above the time axis shows the number of actives in each time period. Even though legitimacy declines to zero, there is no red spike, no explosion, since—as discussed earlier—each new active is being picked off in isolation, before he can catalyze a wider rebellion. And this is why the middle curve—representing the total jailed population—rises smoothly over time.

Figure 8. Large Legitimacy Reduction In Small Increments




The same variables are plotted in Figure 9. However, the scenario is different. We hold legitimacy at its initially high level (of 0.90) for 77 periods. Then, in one jump, we reduce it to 0.70, where it stays. The upper legitimacy curve is a step function. Even though the absolute legitimacy reduction (of 0.30) is far smaller than before, there is an explosion of actives, shown by the red spike. And, in turn, there is a sharp rise in the jailed population, whose absolute size exceeds that of the previous run.

Figure 9. Small Legitimacy Reduction In One Jump $t=77$



The spatial situations in the two runs are compared in Figure 10. Each snapshot is for $t=77$, the time of the one-shot legitimacy reduction above.

Figure 10. Large Incremental (Upper) and Small One-Shot (Lower) Legitimacy Reductions compared at t=77.



Now, why the difference? In the incremental legitimacy reduction scenario, the potentially catalytic agents at *the tail of the grievance distribution* are being picked off in isolation, before they can stimulate a local contagion. The sparks, as it were, are doused before the fire can take off. In the

second—one-shot reduction--case, even though the absolute legitimacy decline is far smaller, multiple highly aggrieved agents go active at once. And by the same mechanism as discussed earlier, this depresses local C/A ratios enough that *less aggrieved agents jump in*. Hence, the rebellious episode is greater, even though the absolute legitimacy reduction is smaller. It is the rate of change--the derivative--of legitimacy that emerges as salient.

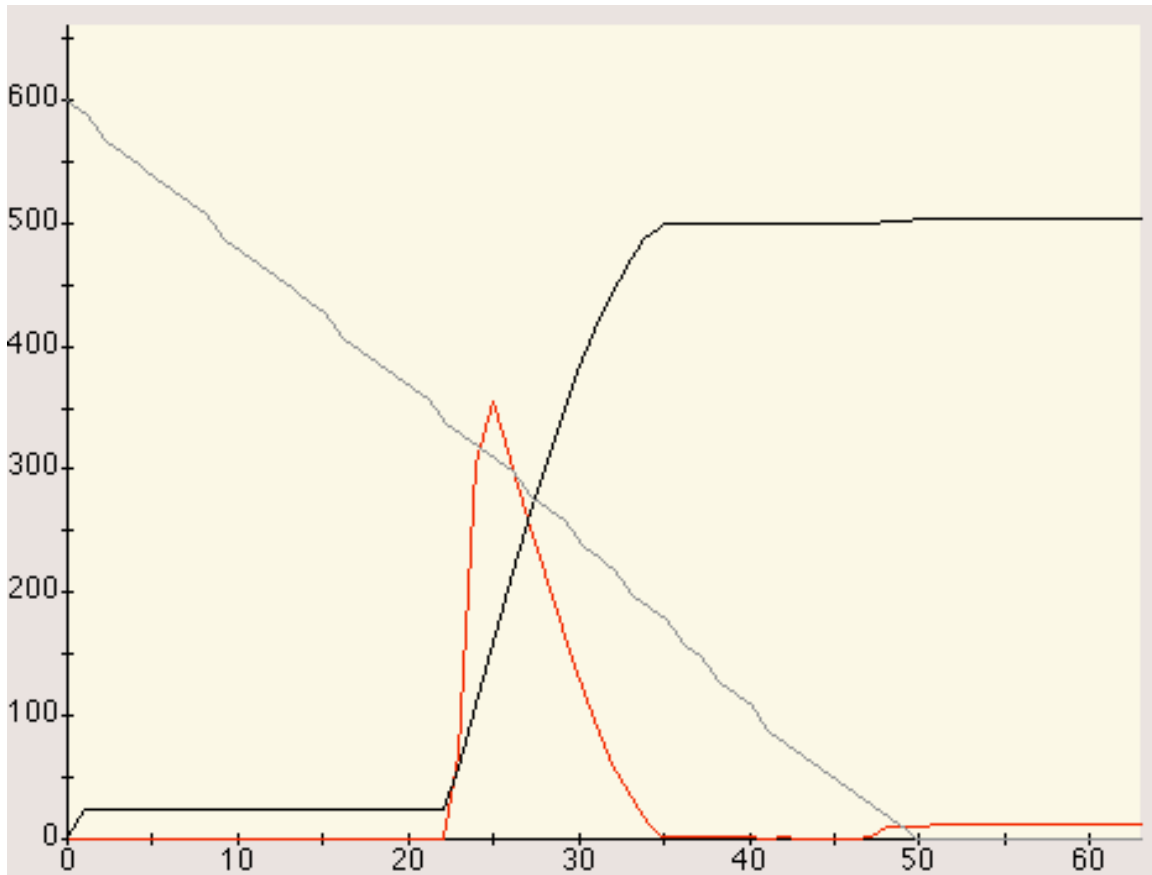
This result would appear to have important implications for the tactics of revolutionary leadership. Rather than chip away at the regime's legitimacy over a long period with daily exposes of petty corruption, it is far more effective to be silent for long periods, and accumulate one punchy expose. Indeed, the single punch need not be as "weighty," if you will, as the "sum" of the daily particulars. (The one-shot legitimacy reduction need not be as great as the sum of all the incremental deltas.) Perhaps this is why Mao would regularly seclude himself in the mountains in preparation for a dramatic reappearance, and why the return of exiled revolutionary leaders--like Lenin and Khomeini--are attended with such trepidation by authorities. Perhaps this is also why dramatic "triggering events" (e.g., assassinations) loom so large in the literature on this topic; often, they are instances in which the legitimacy of the regime *suddenly* takes a dive. While often sufficient, sharp legitimacy reductions are not the only inflammatory mechanism.

Run 5. Cop Reductions

Indeed, "It is not always when things are going from bad to worse that revolutions break out," wrote de Tocqueville. "On the contrary, it oftener happens that when a people that has put up with an oppressive rule over a long period without protest suddenly finds the government relaxing its pressure, it takes up arms against it." According to Kuran (1989), in the cases of the French, Russian, and Iranian revolutions, "substantial numbers of people were privately opposed to the regime. At the same time, the regime appeared strong, which ensured that public opposition was, in fact, unalarming. What, then, happened to break the appearance of the invincibility of the regime and to start a revolutionary bandwagon rolling? In

the cases of France and Iran, the answer seems to lie, in large measure, in a *lessening of government oppression.*" Indeed, De Tocqueville wrote that "liberalization is the most difficult of political arts." Here we interpret liberalization as cop reductions. Beginning at a high level, we walk the level of cops down. Figure 11 shows the typical result.

Figure 11. Cop Reductions



Unlike Run 3 (salami tactics), there comes a point at which a marginal reduction in central authority *does* "tip" society into rebellion. The dynamics under reductions in repressive potential (cops) are fundamentally different from the dynamics under legitimacy reduction in this model--and perhaps in societies. Since both type of reduction are emboldening to revolutionaries, one might have imagined that reductions in the regime's repressive power--the cop level--would produce dynamics equivalent to those under legitimacy reduction. As we see, however, the dynamics are fundamentally different.

Stylized Facts Generated in Model I

While Model I is embryonic and preliminary, it does produce noteworthy phenomena with some qualitative fidelity and therefore seems, at the very least, promising. We saw, first, the unexpected emergence of individually deceptive behavior, in which privately aggrieved agents hide their feelings when cops are near, but engage in openly rebellious activity when the cops move away. In general, the model naturally represents political "tipping points"--revolutionary situations in which the screen is Blue on the left (all agents quiescent) but Red on the right. Surface stability prevails despite deep and widespread hostility to the regime. When pushed beyond these tipping points, the model produces endogenous outbursts of violence and punctuated equilibria characteristic of many complex systems. For some parameter settings, the size distribution of rebellious outbursts and the distribution of waiting times between outbursts exhibit remarkable regularities. The model explains standard repressive tactics like restrictions on freedom of assembly and the imposition of curfews. Such policies function to prevent the random spatial clustering of highly aggrieved risk-takers, whose activation reduces the local Cop-to-Active ratio, permitting other *less aggrieved and more timid* agents to join in. This same catalytic dynamic underlies the intriguing "Salami Tactic" result: Legitimacy can fall much farther incrementally than it can in one jump, without stimulating large-scale rebellion. The reason is that, in the former (Salami) case, the *tails* of

the radical distribution--the sparks--are being picked off before they can catalyze joining by the less aggrieved. The model bears out DeTocqueville's famous adage that "liberalization is the most difficult of political arts," showing that (quite unlike legitimacy reductions) incremental reductions in repressive potential (cops) *can* produce large-scale tipping events. Finally, it may offer a useful laboratory to explore peacekeeping operations.

It should be added that the individual-level specification is quite minimal, imposing bounded demands on the agent's (and cops) information and computing capacity, while still insisting that the agent crudely weigh expected benefits against expected costs in deciding how to act. Agents are boundedly rational and locally interacting; yet interesting macroscopic phenomena emerge.

As noted at the outset, the model seems most promising for cases of *decentralized* upheaval. While one could argue that certain *effects of* revolutionary leadership--reductions in perceived legitimacy through rousing speeches or writings that expose regime corruption--are captured, explicit leadership as such is really not modeled. That could be a weakness in cases--for example, the communist Chinese revolution--where central leadership was important, though some would argue that, even there, the main issue was not the individual leader, but society's "ripeness for revolution." As Engel's wrote, "In default of Napoleon, another would have been found." (Cited in Kuran, 1989). Our tension index might be a crude measure of this "ripeness."

Let us turn now to situations of interethnic violence, as in Kosovo or Rwanda. We will elaborate the model to represent conflict between two ethnic groups.

CIVIL VIOLENCE MODEL II: INTER-GROUP VIOLENCE


Although distinct cultural groups have been generated in agent-based computational models (Epstein and Axtell, 1996; Axelrod, 1997), here, we will posit two ethnic groups: Blue and Green. Agents are as in Model I and turn Red when active. But now, "going active" means killing an agent of the other

ethnic group. The killing is not confined to agents of the out-group known to have killed. It is indiscriminant. In this variant of the model, "legitimacy" is interpreted to mean each group's assessment of the other's right to exist, and for the moment, L is exogenous and the same for each group. (In Appendix B, we indicate how the model can be extended to endogenize legitimacy, allowing it to vary over time, and to take different values for each group.) For Model II, we also introduce some simple population dynamics. Specifically, agents clone offspring onto unoccupied neighboring sites with probability p each period. Offspring inherit the parent's ethnic identity and grievance. Because there is birth, there must be death to prevent saturation. Accordingly, agents are assigned a random death age from $U(0, \text{max_age})$. Here, $\text{max_age}=100$. Cops are as before, and arrest--evenhandedly--Red agents within their vision. There is no in-group policing in this version of the model, though Fearon and Laitin (1996) argue convincingly that this may be important in many cases.

Run 6. Peaceful Coexistence

For the first run of Model II, we set Legitimacy to a high number, just to check whether peaceful coexistence prevails with no cops. Figure 12 depicts a typical situation.

Figure 12. Peaceful Coexistence








On the left, we clearly see spatial heterogeneity and peaceful mixing of groups with no red agents. On the right, we see only the palest of pink shades, indicating low levels of grievance. Harmonious diversity prevails. However, with no cops to regulate the competition, if L falls, even to 0.8, the picture darkens substantially. Indeed, we grow ethnic cleansing.

Run 7. Ethnic Cleansing

The sequence of four panels clearly shows local episodes of ethnic cleansing, leading ultimately to the annihilation of one group by the other: genocide.

Animation 7. Local Ethnic Cleansing To Genocide






Over a large number of runs ($n=30$), genocide is always observed. The victor is random. The phenomenon is strongly reminiscent of "competitive exclusion" in population biology (see May, 1976). When two closely related species compete in a confined space for the same resource, one will gain an edge and wipe the other out. If, however, the inter-species competition is regulated by a predator that feeds evenhandedly on the competitors, then both can survive. Peacekeepers are analogous to such predators. We introduce them in Run 8.



Run 8. Safe Havens

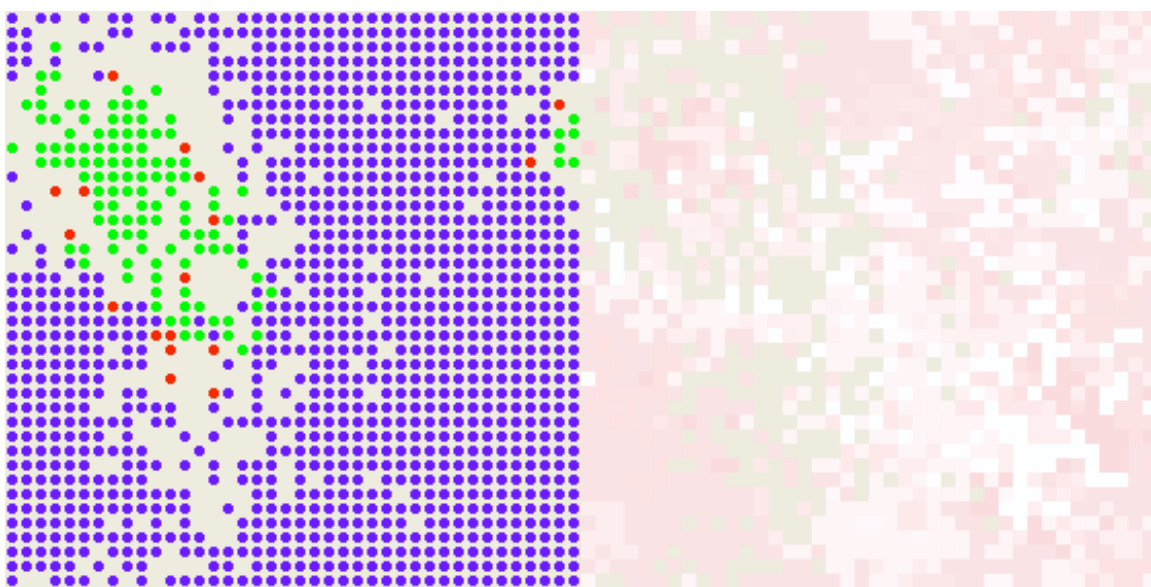

We begin this run exactly as in the previous genocide case. But, at $t=50$, we deploy a force of peacekeepers. They go to random unoccupied sites on the lattice. And this typically produces safe havens. A representative result is shown in Figure 13.


Figure 13. Safe Havens Emerge Under Peacekeeping



Rather than begin with no cops initially, as in the previous run, we now offer a case in which the initial cop density is high. Once a stable pattern has emerged, we withdraw the cops.

Animation 9. Initial Peacekeepers Withdrawn





We begin in Panel 1 with a high cop density of 0.04. With heavy authority from the start, we see the emergence of a stable but nasty regime. The presence of cops prevents either side from wiping the other out, but their co-existence is not peaceful. On the right screen, it is clear that ethnic hostility is widespread at all times. Then, beginning in Panel 2, all cops are withdrawn--the peacekeepers suddenly go home. *Ceteris paribus*, there is reversion to competitive exclusion and, predictably, by Panel 5, one side has wiped the other out.

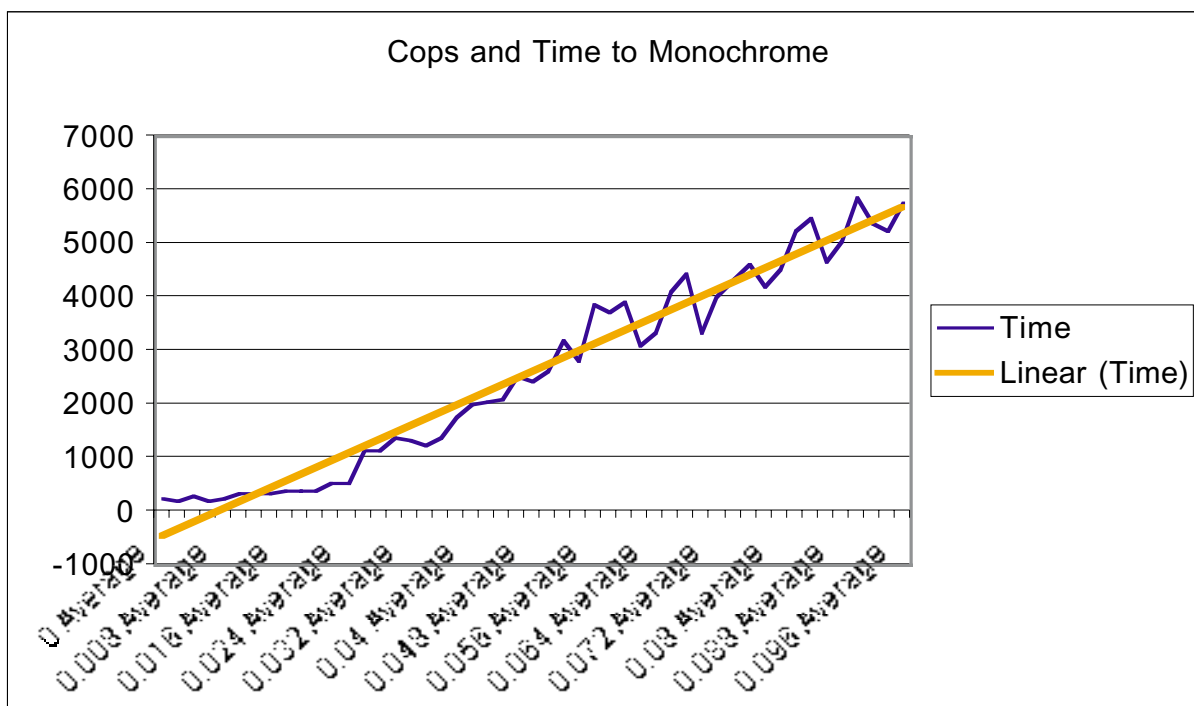
Clearly, peacekeeping forces can avert genocide. But what is the overall relationship between the size of a peacekeeping presence and the incidence of genocide? As an initial exploration of this complex matter, we conducted a sensitivity analysis on cop density.

Cop Density and Extinction Times

For purposes of this analysis, all cops are in place at time zero. But there will be random run-to-run variations in their initial positions and, of course, in their subsequent movements. These initial cop densities are

systematically varied from 0.008 up to 0.01, in increments of 0.002. For each such value, the model was run 50 times until the monochrome--genocide--state was reached. (If it was not reached after 15,000 cycles, we cut off the computation). So, at each density, there is a sample distribution of waiting times over the 50 runs. First, we plot the means of these distributions at each initial density. In Figure 14 we show that curve, and also the best linear fit to the same data.


Figure 14. Waiting Time Mean and Initial Cop Density



Both curves rise to the right overall--the bigger the initial force of peacekeepers, the more time one buys. Notice, however, that the actual mean curve (the blue curve) is *not* monotone! It rises and falls. Within certain ranges, the mean waiting time may fall--genocide may actually be earlier--


with increasing force size! This is counterintuitive and reflects the extreme path dependency and high variance of the process. Indeed, as Figure 15 shows, the variance in waiting time itself increases with density, and is very great above a density of around 0.05.

Figure 15. Waiting Time Variance and Initial Cop Density




Reading across, it is clear that one can have convergence to genocide in 200 cycles at virtually every level of initial cop density! And reading vertically, at a high cop density of 0.08, one can have high effectiveness (delays of over 15000 cycles) or extremely low effectiveness (convergence in tens of cycles). The devil would appear to be in the details in peacekeeping operations. Above a certain point, higher force levels *per se*, do not confer much higher confidence of the outcome in these spatially distributed path dependent dynamics. In figure 16, the standard deviation in waiting times is plotted against force density.

Figure 16. Waiting Time Standard Deviation and Initial Cop Density



The result is a logistic-like S-shaped curve, suggesting diminishing marginal effectiveness above the "knee" of this curve. Since the mean μ is essentially linear and the standard deviation σ is essentially logistic, their ratio σ/μ —the coefficient of variation—is as shown in figure 17

Figure 17. Coefficient of Variation



This entire analysis proceeds from the assumption that all cops are in place at time zero. The same analysis should be conducted for different arrival schedules. At the moment, our claim is simply that the agent-based methodology *permits one to study* the effects of early and late interventions of different scales, which is obviously crucial in deciding how to size, design, and operate peacekeeping forces.

Summary Of Model II Results.

With high legitimacy (mutual perception by each ethnic group of the other's right to exist), peaceful coexistence between ethnic groups is observed; no cops are needed. However, if the cop density is held at zero, and legitimacy is reduced (to 0.8), we see local episodes of ethnic cleansing, leading to surrounded enclaves of victims, leading ultimately to the annihilation of one group by the other. With early intervention on a

sufficient scale, this process can be stopped. Safe havens emerge. With high cop density from the outset, the same level of legitimacy (0.8) produces a stable society plagued by endemic ethnic violence. If cops are suddenly removed, there is reversion to competitive exclusion and genocide. The statistical relationship between initial cop densities and the waiting time to genocide was studied. That analysis showed that, above a certain level, there were diminishing marginal returns to the cop density, per se. Throwing forces at the problem, as it were, is likely to prove inefficient, if it is effective at all. And outright ineffectiveness (quick convergence to genocide) at extremely high force levels, it was shown, is not at all unlikely, due to the path dependent and highly variable dynamics of interethnic civil violence.

Conclusion

Agent-based methods offer a novel and, we believe, promising approach to understanding the complex dynamics of rebellion and interethnic civil violence, and, in turn, to fashioning more effective and efficient policies to anticipate and deal with them.

APPENDIX A. Input Assumptions for Runs*

VARIABLE NAME	MODEL ONE				MODEL TWO		
	Run 1	Run 2	Runs 3&4	Run 5	Run 6	Run 7	Run 8
Cop Vision	1.7	7	7	7	1.7	1.7	1.7
Agent Vision	1.7	7	7	7	1.7	1.7	1.7
Legitimacy	0.89	0.82	0.9	0.8	0.9	0.8	0.8
Max. Jail Term	15	30	infinite	infinite	15	15	15
Movement	none	random site in vision	random site in vision	random site in vision	random site in vision	random site in vision	random site in vision
Initial Cop Density	0.04	0.04	0.074	0.074	0	0	0.04

ALL MODELS

Lattice Dimensions	40x40
Lattice Topology	TORUS
Cloning Probability	0.05
Arrest Probability Constant, k	2.3
Max. Age	200
Agent "active" threshold (for G-N)	0.1
Initial Population Density	0.7
Agent Updating	Asynchronous
Agent Activation	Once per period, random order

* Departures noted in text

APPENDIX B: IMMEDIATE EXTENSIONS AND LONGER-TERM OBJECTIVES.

Four extensions of the present model are of immediate interest:

[1] To explore the effect of *different jail terms* on the dynamics of Model I, define:

$$J = (1 + \frac{J_{\max}}{Life}),$$

where J_{\max} and $Life$ are the maximum jail sentence and agent lifetime, respectively. Then, we generalize the equation for the agent's net risk to read:

$$N = RPJ^\alpha, \text{ where } \alpha \in [0,1].$$

The case explored in the text is $\alpha = 0$. Clearly, as we increase α , we increase the deterrent effect of jail, an effect worth exploring. The next three extensions concern the ethnic violence Model II.

[2] Endogenous Legitimacy

Define the Green-induced Blue per period attrition rate as

$$\beta_i = \frac{BluesKilledByGreensInPeriod_i}{BluesAliveAtStartOfPeriod_i}.$$

Then define Blue's (homogeneous) measure of Green's legitimacy (right to exist) as:

$$L_B(t) = L_0(1 - \sum_{i=1}^m \beta_i^i),$$

where i ranges from one period *ago* to m periods ago, and L_0 is the initial value. Thus, m is the Blue culture's shared "memory" for Green

transgressions. So, if L_0 were 1.0 and, over the last three periods, Green had inflicted on Blue the attrition rates of 0.0 (one period ago), 0.0 (two periods ago), and 0.1 (three periods ago), we would have:

$$L_B(3) = 1 - (0.0^1 + 0.0^2 + 0.1^3),$$

whereas if the attrition rate of 0.1 had been inflicted in the most recent period, Blue would compute

$$L_B(3) = 1 - (0.1^1 + 0.0^2 + 0.0^3),$$

a lower legitimacy, leading to higher Blue grievance, etc. Obviously, the Green calculation of $L_G(t)$ would be exactly analogous. Three points deserve emphasis. The agent can inherit the “cultural memory” of transgressions inflicted before his birth (m may exceed individual agent lifetimes). There is discounting, with more recent transgressions weighing more. With time, old transgressions can fade away (i.e., if $m=10$ all transgressions committed more than ten periods ago are forgotten).

Now, this extension will tend to produce spirals of violence once a few risk neutral aggrieved agents begin executing members of the other group. Each group’s evaluation of the other’s legitimacy will be revised downward, leading to escalating violence unless there are breaks on the process. Of course, peacekeepers are one such regulatory mechanism. But, a richer model would allow for the endogenous regulation of inter-ethnic violence, despite high tension, without external forces. A modification that would operate to mitigate the escalatory dynamic just noted is as follows.

[3] Agents estimate the probability of other-group retaliation.

Recall that agents compute the probability of arrest via the formula:

$$P = 1 - \exp[-k(C/A)_v],$$

where C/A is the cop-to-active ratio within the agent's vision v . For the inter-ethnic variant of the model, let A_B and A_G denote active Blues and Greens. Then, there would be a deterrent effect if Blues were to compute

$$P = 1 - \exp\left[-k\left(\frac{C + A_G}{A_B}\right)_v\right]$$

Then, even in the absence of cops ($C=0$), an explosion of Green actives would serve to deter Blue activation, and vice versa, damping the escalatory pressures generated under endogenous legitimacy above. How the two effects --escalatory and deterrent--would interact is potentially quite rich.

[4] Relative Deprivation.

A final and easily implemented extension would crudely capture *relative deprivation* (Gurr, 1972). One would assume different Hardship distributions for different ethnic groups. For example, one could draw the Hardship of Blue agents from $U(0, .4)$, and the Hardship of Green agents from $U(.5, .9)$., in which case the poorest Green would still be better off than the richest Blue.

Computational Reconstruction

The principal longer-term objective would be the computational reconstruction of the true sequence of events in, for example, Rwanda. This effort would be on the model of the Artificial Anasazi Project, which—aided by an extremely detailed environmental and demographic data base—did reconstruct important elements of a true spatio-temporal history in an agent-based computational model. For a report, see Dean, et al (1999).

REFERENCES

- Arendt, H. 1963/1965. *On revolution*. New York: Penguin Books.
- Bates, Robert. 1983. "Modernization, Ethnic Competition, and the Rationality of Politics in Contemporary Africa." In *State Versus Ethnic Claims: African Policy Dilemmas*, ed. Donald Rothchild and Victor A. Olorunsola. Boulder, CO: Westview.
- Calvert, Randall. 1995. "Rational Actors, Equilibrium, and Social Institutions." In *Explaining Social Institutions*, ed. Jack Knight and Itai Sened. Ann Arbor: University of Michigan Press.
- Collier, Paul and Anke Hoeffler. 1999. "Justice-Seeking and Loot-Seeking in Civil War," World Bank DECRG Working Paper (April).
- Collier, Paul, Anke Hoeffler, and Mans Soderbom. 1999. "On the Duration of Civil War," World Bank DECRG Working Paper (February).
- DeNardo, J. 1985. *Power in numbers: The Political Strategy of Protest and Rebellion*. Princeton, NJ: Princeton University Press.
- Doyle, Michael W. and Nicholas Sambanis. 2000. "International Peacebuilding: A Theoretical and Quantitative Analysis." Draft (March 7, 2000).
- Ellison, Glenn. 1994. "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching." *Review of Economic Studies* 61 (July): 567-88.
- Fearon, James D. and David D. Laitin. 1996. "Explaining Interethnic Cooperation," in *American Political Science Review* 90:4, December 1996.
- Graham, Carol and Stefano Pettinato. 1999. "Assessing Hardship and Happiness: Trends in Mobility and Expectations in the New Market Economies." Center on Social and Economic Dynamics Working Paper No. 7, The Brookings Institution.
- Gurr, T.R. 1970. *Why Men Rebel*. Princeton, NJ: Princeton University Press.
- Horowitz, Donald. 1985. *Ethnic Groups in Conflict*. Berkeley: University of California Press.
- Kandori, Michihiro. 1992. "Social Norms and Community Enforcement." *Review of Economic Studies* 59 (January):63-80.
- Kuran, Timur. 1989. "Sparks and prairie fires. A theory of unanticipated political revolution" in *Public Choice* 61 : 41-74.

Mao Tse-Tung 1930/1972. "A single spark can start a prairie fire.' In *Selected Military Writings of Mao Tse-Tung*, 65-76. Peking: Foreign Languages Press.

Olson, M. 1965/1971. *The logic of collective action: Public goods and the theory of groups*, Cambridge, MA: Harvard University Press.

Posen, Barry. 1993. "The Security Dilemma and Ethnic Conflict." *Survival* 35 (Spring):27-47.

Rogers, E.M. 1983. *Diffusion of Innovations*, 3rd ed. New York: Free Press.

Skocpol, T. 1979. *States and Social Revolutions: A comparative analysis of France, Russia, and China*. Cambridge: Cambridge University Press.

Snyder, D. and Tilly, C. 1972. "Hardship and collective violence in France." *American Sociological Review* 37: 520-532.

Steinbruner, John D. 2000. Principles of Global Security. Washington, DC: Brookings.

Tocqueville, A. de 1856/1955. *The Old Regime and the French Revolution*. S. Gilbert, transl. New York: Doubleday.

Tullock, G. 1974. *The Social Dilemma: The Economics of War and Revolution*. Blacksburg, VA: University Publications.

Zagorin, P. 1973. "Theories of revolution in contemporary historiography." *Political Science Quarterly* 88: 23-52.

Zonis, M. 1983. "Iran: A theory of revolution from accounts of the revolution." *World Politics* 35 (July): 586-606.