

Defining Violence:
A Plausibility Probe Using Agent-Based Modeling

Ian S. Lustick
University of Pennsylvania

Prepared for LiCEP, Princeton University, May 12-14, 2006. This is an updated and corrected version of a paper prepared for Presentation for “Mapping the Complexity of Civil Wars,” Swiss Federal Institute of Technology Center for Comparative and International Studies Conference in Zurich, September 15-17, 2005. Many thanks to colleagues at the Solomon Asch Center for Study of Ethnopolitical Conflict at the University of Pennsylvania for their critical comments and suggestions, especially to Britt Cartrite, Alex DeRenzy Channer, Roy J. Eidelson, Dan Miodownik, Kaija Schilde, and Stefan Stohler. Thanks also to Dan Miodownik for technical assistance. Templates used in experiments reported here are available for replication purposes. Please contact the author at ilustick@sas.upenn.edu. PS-I executable files and operation manuals are freely available. Visit <http://www.psych.upenn.edu/sacsec/abir/>.

The Unmet Challenge of Defining Violence

Science requires observation and comparison. Implicit in the notion of comparison is an expectation of the conceptual stability of boundaries surrounding types of observable events. Determining when one sees an orange and when one sees an apple, requires clear coding rules for what counts as an orange or an apple, but the scientific value of the exercise also requires a stable notion of “fruit” so that what is learned about differences and similarities between apples and oranges can be used to generate hypotheses about other things that are to be counted as within the broader class to which apples and oranges belong. We learn more from comparing apples and oranges about plums and grapefruit, than about refrigerators or chipmunks.

Few definitional problems are more familiar than those associated with defining “terrorism.” Of course one problem (“My terrorist is your freedom fighter.”) is the politically and rhetorically fraught aspect of the word that complicates communication and triggers political struggles only thinly disguised as competitive consideration of alternative conceptual strategies. But that is only one difficulty. Another difficulty is the problem of linking any plausible definition of the term with a class of events or behaviors that is narrow enough to permit confident judgments of its absence. For example, if terrorism is defined as efforts to affect the behavior of others by scaring them, we might see the bombing of Hiroshima, the 9/11 attacks, and a parent’s threat to withhold allowance if a child’s room is not properly cleaned, as relevant observations and therefore require any theory of terrorism to account for or be tested by patterns of outcomes similar to each of these examples. On the other hand, if the definition is narrowed too much (e.g. lethal non-state actor politically motivated attacks on civilians), the theories we can build are likely to be a function of some “accidental” aspect of the particulars of the domain specified by the definition, thereby rendering the results of study incomparable with the results of other studies using equally plausible, but equally narrow definitions.

Another problem is the need to construct a definition that is strategically located with respect to theoretically developed fields of investigation. For example, this condition requires a definition so that the coding of an observation does not itself require explicit or implicit deployment of multitudes of theories (sociological, psychological, physical, biological, legal, political, etc.). Thus a definition of terrorism that would require us to establish whether an act was or was not deviant, was or was not intended, was or was not legal, was or was not “political,” etc. would tend to conflate the definitional problem with the substantive challenges confronting the disciplines required for operationalizing that definition (How is “deviant” to be defined, or “legal,” or “political?”).

Given the difficulties associated with defining terrorism, it is not surprising that in the 1980s and 1990s the term had much more currency outside of academia, than within. To be sure, this pattern has begun to change after the events of 9/11 multiplied the funding available for “terrorism” studies by a gargantuan factor. It is still true that social

scientists, psychologists, sociologists, anthropologists, political scientists, have traditionally exhibited a much greater willingness to study “violence,” than terrorism. It is rarely admitted, however, that the focus on violence not avoided the problems associated with defining terrorism. In fact students of “violence” and even “political violence” have had no more success offering a consistent or even explicit, denotative definition of violence, than others have had trying to define terrorism. True, we have an enormous amount of data about killings, reports of killings, numbers of armed insurrections or armed clashes of different sizes, etc. In political science at least, books, conferences, journal articles, and dissertations abound on the subjects such as violence and politics, the effects of political violence, the conditions that produce violence, and techniques to reduce violence. For the most part however, the question of what is actually meant by “violence” has been avoided. Thus we may learn from this literature about the circumstances that draw particular groups into the killing of others, or about the proportion of killing that is motivated within a civil war by personal as opposed to political motives, or about the amount of damage in lives or property that might be used as a proxy for the amount of violence present, or about the strategies of elites that are most likely to reduce participation in deadly riots. But we do not thereby learn about the difference that “violence,” per se, makes in a situation because the concept of “violence,” independent of the damage it causes or the particular way its effects are registered (deaths, per thousand in a population, for example), is almost never specified.

This paper identifies the absence of a coherent and analytically strategic definition of violence as an important problem for comparative political analysis. A candidate for such a definition is then advanced and operationalized in a virtual, agent-based modeling environment. This implementation, within a highly stylized context of an authority structure exposed to violent strikes, provides opportunities to evaluate the soundness of the definition by seeing whether outcomes observed virtually correspond in sensible and systematic ways with standard intuitions.

A survey of dictionary definitions of violence illustrates the problem. The single most common element in these definitions is the idea of an outsized or powerful and damaging effect. Some definitions include intention to commit the damaging effect, some do not. Some stress the requirement of physicality, some do not. Some omit the damaging effect as an element of the definition, and stress the forcefulness or vehemence of the act. Some include the suddenness or unexpectedness and/or unnaturalness of a "violent" act compared to an expected or natural process. Some exclude non-human causes of violence, others include storms, for example, as violent. Some mention illegality as a criterion, though most do not.¹ In short, dictionaries cannot be relied on to provide a "common law" definition of violence that can be expected to be stable in the minds of researchers or readers without explicit stipulation of the definition being used.

Charles Tilly has tried hard to define the term, but has more or less admitted failure. Tilly cites Hannah Arendt's interest in violence as springing from her revulsion at the 1968 student riots in France. Arendt's treatment is at pains to distinguish "violence" from "power," by which she seems to want to distinguish between illegitimate physical coercion and legitimate governmental authority; but it is difficult to distinguish her

definition of violence (never explicitly presented) from her theory of it.² Tilly quotes the French commission that investigated those disturbances, reporting that the commissioners searched for a definition of violence, but "finally threw up their hands" agreeing only that "violence consisted of force used wrongly."³

The absence of and need for a precise definition of "violence" was acknowledged by all participants in a panel at a recent American Political Science Association meeting featuring leading researchers working in the area of political violence. In answer to the author's question about the formal definition of violence, as opposed to the damage caused by violence, that participants in the panel were using in their work and commentary, each panelist, along with the chair and the discussant, responded that no specific definition had been formulated. The response that secured the support of most participants, offered by Jeremy Weinstein of Stanford, was that such a definition was not to be expected as yet since the study of political violence was "still in its infancy."⁴

To be sure, as the example of Tilly suggests, sociologists have exhibited somewhat more concern with the definition issues involved with the study of violence than have political scientists. A particularly pointed and thorough review of attempts that have been made to define violence explicitly was published in 2002 by the sociologist Mary R. Jackman, who noted the failure of social scientists to offer a clear definition of violence. Jackman's treatment emphasizes the crippling analytic difficulties she sees as identified with available, explicit, definitions of violence before offering a 'generic' substitute definition."

Much as some scholars have bemoaned the lack of cohesion in research on violence...most scholars have proceeded without hesitation as though the conceptual tangle had been cleared. Researchers commonly refer to a phenomenon called violence that implies a clearly understood, generic class of behaviors, and yet no such concept exists.⁵

Jackman's critique of existing definitions stresses what she regards as a distorting displacement of attention to one or more particular aspects of certain kinds of acts commonly viewed as violent without identifying what, generically, is to be considered "violence." These incidental aspects include malicious motivation, corporal (as opposed to psychological) injury, deviance, physical behaviors, coercive rather than self-inflicted or invited acts, legality, perpetration by individuals rather than organizations or communities. Her own solution to the problem is drastic, and to my mind, unacceptable. She offers a "comprehensive" definition of violence with which she means to avoid excluding any behaviors or events which might be considered violence. "Violence" as Jackman defines it, refers to "actions that inflict, threaten, or cause injury. Actions may be corporal, written, or verbal. Injuries may be corporal, psychological, material, or social."⁶

With this definition Jackman does avoid the problem of other definitions that exclude events analysts, or some analysts, might want to include. However, this is achieved at the cost of a conception of violence that is essentially identical to the concept

of “damage.” Two problems arise. One is the problem of coding “injury” or “damage.” Surgery involves injury before it helps to heal. Thus violence could be a matter of time frame, or, of course, perspective. Is the destruction of a mountain by a mining company injury to the beauty of the local community’s environment or assistance to the local economy or the national aspiration to achieve energy independence? Another more fundamental problem is that with a definition this broad Jackman can offer no theoretical guidance to researchers. A concept of “damage” or of actions that produce, or could produce, damage or injury subsumes so many phenomena and so many circumstances that research on “violence” would end up as virtually coextensive with research on society and politics. The definition is simply too capacious too to avoid the crippling problems associated, for example, with efforts in political science to treat “power” as a causal variable.

Jackman’s survey of research on violence, along with extensive work in political science that similarly fails to define that key term, testifies to the odd fact that despite the in principle necessity of clear definitions, social scientists using traditional methods (discursive, statistical, or formal closed-form approaches) can often pursue research programs ahead that entirely ignore basic definitional issues. We see that in the immense amount of attention given, as noted earlier, to modeling the causes or effects of violence, without ever modeling violence itself.

For agent-based modelers the challenge of defining violence is no different, it is just less avoidable. In work done with PS-I on political violence and its effects on political stability and foreign policy outcomes in the Middle East, the inability to define violence itself required that virtual worlds created for this project not incorporate that variable, but rather feature the conditions believed to give rise to “violence” as the independent variable and patterns of outcomes believed likely to be immediately produced by violence as the dependent variable.⁷ Indeed it was in part frustration with the absence of violence itself from the MEPOLITY model, and the inability to incorporate it without a clear, but analytically limited denotative concept of violence, that has produced my interest in achieving a clear, and therefore operationalizable, definition of violence.

In the balance of this paper I will provide one definition of violence whose clarity and generalizability may justify its use. It is an abstract definition and does not correspond with many intuitive/familiar uses of the term. However, given the sloppy and inconsistent ways in which the term is used, it is automatically impossible for a consistent and explicit definition to match standard impressions and usage. After presenting the definition, I will briefly present and discuss experimental results using an agent-based modeling operationalization of it.

A Definition of Violence

What is it about a particular behavior, event, or situation, independent of the damage it causes (damage that could be caused non-violently, for example) or its legality or morality that can be coded as “violent?” Are all insults violent? Are lawsuits resulting

in the destruction of whole villages violent? Are unintended casualties the result of violence? Is the prolonged application of steady, low level pain violent? Can threats, per se, be violent? Must violence be physical? Are protestors massing themselves against the movement of traffic violent? Does human action have to be involved? Was the tsunami, which struck so “violently” and killed so many, “violent?” It did have political effects, for example on the insurrection in Aceh, the Indonesian government’s policies there, and on international attitudes toward that region. Will we recode it as violent if we discover, in twenty years time, that it was actually triggered by a Chinese device that was purposely used to trigger the earthquake/tsunami?

The definition of violence, and of violent events, whose usefulness I am evaluating draws inspiration from Arendt's treatment by focusing on the element of a sudden, intrinsically unpredictable, and therefore drastic increase in the potential negativity of the stakes in an encounter. For although, as I have noted, Arendt does not precisely define violence, she does indicate a key characteristic of it. "Violence," she writes, "harbors within itself an...element of arbitrariness...this intrusion of the utterly unexpected..."⁸ Indeed the closest she comes to offering a definition of violence is that "the moment we approach the realm of violence" we encounter an "all-pervading unpredictability."⁹

Physicality, for example the throwing of a punch, because it is intrinsically difficult to be sure exactly what the effect of the blow will be, is likely to be closely associated with violence.¹⁰ Will the blow be annoying, a bit painful, very painful, incapacitating, or fatal? As the fist is traveling toward me, I just don't know, and neither really does the person throwing the punch. Insults can do damage, perhaps even violently by my definition, but since the effects are processed through social, linguistic, and emotional sieves, and can be reprocessed and re-interpreted or reframed, we can expect, or at least hypothesize, that violence will more often be associated with physical actions and threats of physical actions, than with insults drained of perceptions of the imminence of physical attack (for example, threats delivered across a telephone line or the internet, or shouted across a wide river). Accordingly I define violence, in the social world, as follows.

A situation or event is violent to the extent that a sudden and drastic increase occurs in the scale of negative values at stake. The more drastic and rapid the increase, the more negative, and the more people who experience this increase, the more violent is the situation or event.

According to this definition, large-scale destruction or damage is not necessarily to be regarded as, or to indicate, violence.¹¹ If the damage or destruction were inflicted very slowly and predictably it would certainly not be coded as violence by this definition. Thus, to take a particularly unsettling example, while the initiation of torture might well be coded as violent by this definition, torture entailing the slow prolonged, inexorable infliction of discomfort or annoyance would not be. The definition discourages use of intentionality, legality, or modality (psychological, physical, financial, etc.) of interactions, not because they are deemed unimportant, of course, but so that they can be

preserved, outside of the definition, as empirical questions. For example, to what extent is violence possible without physical destruction? What is the balance in the extent of violence in different kinds of situations that is or is not attributable to intention? How do the political effects of violence differ if it is legal as opposed to illegal?

Of course there are a variety of fascinating probes of this definition that may be considered—probes normally used to evaluate definitions in terms of their consistency, intuitive appeal, and theoretical usefulness. For example, if someone does not know a piano is about to land on his head, is the event violent when it does occur? (Yes, because an outside observer can see the sudden appearance of death in connection with a casual stroll.) Is the violence in the act or in the head of the observer? (Violence as a social fact is in the heads of observers. The extent to which this requires the actuality of particular kinds of actions or events in the world is an empirical question.) Is there violence in the confrontation between two gangs, before any action or explicit threat has occurred? (There is a difference between violence and violence potential. Unexploded dynamite has violence potential. The sudden appearance of dynamite in the hand of an interlocutor, just as the sudden confrontation of two silent gangs, might themselves be coded as violence, though the level of violence might well decrease rapidly unless actions or threats were made.)

Most questions about defining violence arise because intuitive desires to use violence as an extensive category for events and situations we somehow feel are “violent” are difficult to understand as consistent with an intensive definition of the term that remains stable across conversations or domains of inquiry. Formal modelers, however, need an intensive definition of violence, and so I strive for a formulation that, while not entirely devoid of intuitive appeal and not entirely different from what it is standardly, if incoherently, considered to be, is clear enough to be implemented unambiguously in a virtual world and is designed strategically enough to avoid solving by definition important empirical or theoretical questions.

Operationalizing This Definition of Violence in PS-I

PS-I is a sophisticated platform for producing agent-based modeling templates. It is designed with point and click and Boolean operator interfaces that do not require knowledge of Java or other programming languages and which reduces requirements for quasi-code syntax to a minimum, even in the design of relatively complex models. Key elements within PS-I include agents within a cellular grid “activated” on one of a repertoire of “states” (standardly referred to as “identities”) available to that agent. Both the activation of a particular state by a particular agent and the complexion of the repertoire of individual agents change according to updating rules that allow agents to monitor their surroundings with stipulated levels of sensitivity. In addition to monitoring proximate and remote neighborhoods, depending on agent characteristics, agents can also monitor knowledge available to all agents. These valence signals are referred to as “biases.” Biases are exogenously generated stipulations of marginal reductions or increases in the attractiveness of particular states. These stipulations can change over

time with a volatility, predictability, and within ranges of variation stipulatable by the user.

The definition of violence presented above was operationalized in PS-I as follows. A distinct “violent” agent class was created. An instruction from a script does three things simultaneously:

- At a stipulated time it selects some set of positions in the array, (randomly or specifically selected within a particular region if desired).
- It forces agents located in these positions to activate, for just two time steps,¹² on a tag (“identity” 20 in our model), to which has been assigned a permanent and very low negative bias of -10.
- It implements a routine that exposes any agent monitoring the locations affected to a one-time burst of "influence" 30 times the normal level of that location.

The violent effect is achieved by the sudden multiplication of an extremely high influence level with an extremely negative bias. From the point of view of each agent that has been seeking to adapt to marginal changes in neighborhood and environmental “bias” conditions, this is a shock—a sudden and essentially arbitrary eruption of extreme negativity into encounters between agents.

The immediate reaction of agents experiencing the violence in their neighborhood is to activate on the tag of "20," despite its extreme negative valence. Such an agent could remain for more than one or two time steps on this “identity,” but with alternatives available and with the bias assigned to that identity so low, it is very likely instead to resume or seek to resume its adaptive course. In any case the world, both internally and externally, will have been changed in possibly irrelevant or possibly crucial ways by the “shock” of the intrusion of violence and by the interruption in the pattern of activation by agents inhabiting directly affected cells and by neighbors. Local equilibria that might have been established might not be re-attained.

The “empirical” questions in this virtual world about the aftermath of “violence” correspond to the questions of crucial social scientific importance in the study of political violence. For example, under what conditions and with what likelihood will violence have different implications for perpetrators, regimes, mobilized, and/or unmobilized groups?

Plausibility Probe

To examine whether implementing violence in this way would have effects in patterns consistent with reasonable stories analysts might tell about the effect of violence in a salient kind of political situation, a template called “Violdef” was created in a square topology with 64-cell sides. A highly stylized “authority structure” was created as a symmetrical web of influential agents activated on the same “regime” identity. This web

of influential agents is located in the center of a larger landscape. Aside from the activation of regime influentials on the regime identity (“0” or red), and the exclusion of identity “20” from the repertoires of all agents, all twenty normal identities (0-19), including the regime identity, are distributed randomly to all agents every time the model is initialized.

Figure 1 shows Violdef at time 0. Cells marked by icons are influentials, radiating in a regular pattern originating with a “Great Leader” of influence 4 (circle), then “Lords” with influence 3, then two rings of Lackeys with influence 2.




Figure 1. Typical initialization of Violdef

A zoomed-in version of the screenshot in Figure 1 is displayed as Figure 2, focused on Auth_shape [a 33X33 bloc of 1089 cells surrounding and including the regime authority structure]. The different icons marking different ranks of influentials are clearly discernible: Great Leader, circle; Lord, propeller; Lackey, small central square.




Figure 2. Zoomed-in view of Auth Shape, $t=0$

In addition to their own activated state, all agents monitor their Moore neighborhoods (the eight adjacent agents, including those touching at the corners) and update simultaneously on even time steps. Influential agents have six identities including the activated identity. Other (basic) agents have six or seven identities in their repertoires. All agents operate according to the same updating rules, described in Table 1.

Updating Trigger	Meaning	Identity Weight Margin Required to Effect Change ¹³
Rotation	An identity in the agent's repertoire is activated and the activated identity is "rotated" back into the non-visible repertoire	2
Substitution	An identity not in the agent's repertoire is brought into that repertoire and a non-activated identity is discarded from the repertoire	5
Substitution and Activation	An identity not in the agent's repertoire is brought	7

	<p>into that repertoire, a non-activated identity is discarded from the repertoire, and the newly incorporated identity is activated.</p>	
--	---	--

Table 1: Updating Rules

Batches of 100 runs were generated to time 308 under twelve different conditions.¹⁴ In each of the 100 runs a distinctive but randomly produced stream of changing identity bias assignments perturbed the trajectory. These distinctive streams of perturbations were held constant across conditions so that, for example, the stream of changing bias assignments affecting run #12 in was identical across conditions.¹⁵ The conditions themselves were produced by crossing three variables:

- Violence:**
1. no violence;
 2. one violent strike on the “southwestern” corner of the regime’s authority structure entailing a two time step punctuation at t=50 entailing the transformation of 50 per cent of basic agents in the targeted area into “violent” agents;
 3. one strike on a region of identical size in the center of the regime’s authority structure entailing a two time step punctuation at t=50 entailing the transformation of 20 per cent of agents in the targeted area into “violent” agents.

**Stability¹⁶
of Political**

- Environment:**
4. stable conditions (low bias volatility “250”);
 5. volatile conditions (increased bias volatility “500”).

Challenger

- Identity:**
6. no Challenger present;
 7. Challenger present (all agents within “Auth-Shape” not having the regime identity in their repertoire at t=0 re endowed with the Challenger identity [identity 16, yellow]).

Figures 3 and 4 display the same version of Violdef used to produce the first two figures. In Figures 3 and 4 are highlighted all agents with identity 16 (the Challenger identity) in their repertoires but not identity 0 (Regime identity). Note that in Figure 3 this is a random distribution while in Figure 4, representing the “Challenger Present” condition, the overwhelming majority of agents have the Challenger identity in their repertoires (while only a normal random sample are activated on that identity). Figure 5 shows the areas subject to violence in the Southwest and Center violence conditions, respectively.

Figure 6 shows a sequence of snapshots illustrating the sudden appearance of violence (the cell positions turned to blue and marked with a small black slash in the upper left hand corner— $t=49$, $t=51$); the immediate effect of this burst of powerful negativity, agents exposed to the violence are affected (locations turning blue, but without the small black slash) $t=52$); and the aftermath as agents seek to resume their normal activities in the absence of violence— $t=53$, $t=54$, $t=55$.




Figure 3. No Challenger present condition, $t=0$. Agents highlighted have Challenger identity in repertoire but not the regime identity.




Figure 4. Challenger present condition, $t=0$. Agents highlighted have Challenger identity in repertoire but not the regime identity.




Figure 5: Regions Exposed to Violence in Southwest and Center of the Authority Structure




Figure 6: Auth_Shape in Typical Run Violence Occurs in the Southwest (Volatility 500, Challenger Present), Time Steps 49, 51, Center Violent Strike. Light-blue cells with black triangle in upper left corner are locations directly struck by violence.)

A key mechanism involved in the dynamics of these landscapes is produced by the particular initializations associated with the rapid shuffle of bias assignments in the first eight time steps of any of the 100 runs in each treatment condition (historical legacies), and by similarly distinctive but slower paced subsequent streams of perturbations. Among these 100 distinctive streams of perturbations, some may be unfavorable for the regime but may be distinctly favorable for some other identity, producing regional cascades of activation on that identity. If these cascades occur in the vicinity of influential agents those agents may be captured by the cascading identity. My expectation was that violent strikes would disrupt the regime's ability to enforce activation on red inside Auth_Shape and decrease the integrity of regime by opening opportunities for identities other than the regime identity to capture influential agents, thereby reducing the number of influentials activated on the regime identity. I expected that under volatile conditions the regime would be more successful on each measure than under stable conditions, since rival identities would not benefit from the steadying presence of their own web of influentials. I expected that the presence of the Challenger identity would multiply the effects of violence, but was not sure what the interaction would be between volatility and the presence of the Challenger. If my expectations were met, and if a plausible story could be told about the interaction between volatility and the presence of a Challenger, the usefulness and appropriateness of my definition of violence and its operationalization for studying political violence would be corroborated.

Experimental Results and Discussion

The Excel file that accompanies this document as Appendix_Violdef contains charts reporting data from experiments conducted under conditions explained above. Neither the charts in the Appendix nor this discussion treat all data collected in different treatment conditions and as part of various sensitivity studies. In any event, my intent is not to use these experiments to pose decisive questions about the effects of political

violence. The purpose here is narrow: to interpret patterns in the data as offering stories about the relationship between violence and political outcomes to assess the performance of the definition of violence as I have operationalized it. All claims about differences between means and distributions of outcomes across different treatment and conditions have been substantiated by T-tests and one-way ANOVA tests at a .001 level of significance.¹⁷

As expected, under normal conditions, violence decreased the degree to which the authority structure was able to establish and consolidate a pattern of public affiliation with the regime. Figure 7 shows that in the absence of a Challenger identity (a latently distributed predominance of sympathy for one particular non-regime political affiliation), the average prevalence score for the regime identity within Auth_Shape was highest in the non-violence condition under both stable and volatile conditions. Under stable conditions (volatility 250) violence occurring in the periphery, that is in the southwest, had a smaller impact on average regime identity prevalence than when the center of the authority structure was the locus of the violence. However, under volatile conditions (volatility 500), the regime identity fared better when violence occurred in the center than in the periphery. This is consistent with my hypothesis that where the governing structure of the regime was stronger, in the center, its advantage over unempowered identities under conditions of turbulence could show itself. Data in Figure 8 are drawn from the same experiment run with the presence of a Challenger identity as explained above.¹⁸ Although the values for average Regime identity prevalence are lower across the board in this condition, we again see that the Regime identity does significantly better under conditions of no violence than when violence struck in either the southwest or center. We also see, again, that under relatively stable conditions, the regime identity does better when faced with violence in the periphery rather than in the center where its stabilizing/repressive capacities are concentrated.

From the data in Figure 9, "Effect of Volatility, Violence, and Challenger Presence on Regime Integrity," we see that under conditions of volatility and stability violence and the strategic direction of it (i.e. center as opposed to southwest) each decrease the number of influentials activated on 0 at t=308. Since the violent strike in the center targets a region more heavily populated by more influential agents (the Great Leader, Lords, Lackeys) than the strike in the southwest, the regime suffers more damage, on average, from center than from southwest violence. This effect is relatively small when conditions are volatile, but substantial when stable conditions are more likely to provide consistently favorable conditions for the Challenger identity or consistently unfavorable conditions for the Regime identity.¹⁹ This pattern is sustained whether a Challenger identity is present or not.

Because of the nature of agent-based modeling experiments it is relatively easy to probe more deeply into the data by studying the shapes of distributions of batches of outcomes rather than simply comparing averages. Consider that in the no violence condition the two variables determining outcomes are the particulars of the random initialization of repertoire complexion and activation distribution and the randomly generated sequence of bias assignments. Across the three experimental conditions in any

particular experiment these factors are held constant, permitting us to consider the data in any particular set of three runs as pairs or even triplets. The procedure followed to compare heuristically the shapes of outcome distributions is to first order by increasing magnitude the data points measured at $t=308$ in the no violence condition. The corresponding values for one or both of the violent conditions are then plotted alongside the run numbers as determined by the rising curve of outcomes in the no violence condition. In this way it can be seen whether the difference in average outcome is due to especially large discrepancies in some portion of the rising curve of outcomes in the no violence condition.

For the sake of simplicity and clarity the data displayed in Figures 10 and 11 match the no violence under volatile conditions with center violence under volatile conditions, but with the Challenger identity not present (Figure 10) and present (Figure 11). We see that that violence in the center produces frequent reductions of regime identity prevalence whether or not a Challenger is present. But we see that on the extreme right side of these displays, reporting runs in which the general conditions (streams of bias assignment perturbations) were favorable to the regime identity, neither violence nor the presence of a Challenger had much of an effect. Thus we see the values in this section of the displays overlap with one another. On the left side of these displays, where general conditions for the regime identity were relatively poor, we see that its prevalence was significantly reduced in the no violence condition by the presence of a Challenger. In this section of the displays we also see significant variability in the effects of violence. In the Challenger present condition more runs were registered in which the regime identity is eliminated altogether but there were also more runs in which the regime identity actually benefited from the violence. We also see an interaction between violence and the presence of a Challenger such that when general conditions for the regime identity were moderate (between runs ranked here between 50 and 80) fewer runs showed significant reductions when the Challenger was absent than when it was present. Overall, increasing favorability of exogenous conditions decreases the effects of violence as well as the effects of the presence of a Challenger. Occasionally, however, depending on specific configurations of historically produced starting points (at $t=8$) or of combinations and sequences of exogenous perturbations (bias assignments) violence has a negative effect on regime identity prevalence, even under quite positive conditions (runs ranked between 75 and 85).

Very similar patterns appear in the data displayed in Figures 12 and 13 regarding the frequency, robustness, and distribution of effects of violence and the presence of a Challenger on regime integrity, i.e. on the number of influentials, activated on the regime identity at $t=0$, that remain activated on that identity at $t=308$. Again we observe that, under volatile conditions, violence has a more potent effect when general conditions are unfavorable for the regime; that violence can occasionally benefit the regime; that variability is greater when the Challenger identity is present; and that under extremely positive conditions for the regime, neither the presence of a Challenger identity nor violence has a significant effect.

By focusing only on the Challenger present condition, we can investigate the influence of violence and volatility of conditions on the success of the Challenger identity. The data in Figures 14 and 15 show that when conditions are volatile there is a noticeable bifurcation in the prevalence scores registered by the Challenger identity, more than doubling in value between the 71st and 77th most successful runs in the no violence condition. A similar, though somewhat less marked pattern, is present when conditions are stable. When conditions are not favorable the Challenger identity lies dormant. When they turn favorable we see that violence may well have little effect, but when it does have an effect it is usually to help the Challenger identity, whether the violence occurs in the center of the regime's authority structure or in the periphery. The Challenger identity is most successful when it can exploit violence in the center, when general conditions are stable, and when particular conditions are such as to strongly favor the Challenger identity. The data regarding Challenger identity capture of influentials formerly activated on the regime identity, displayed in Figures 16 and 17, show similar overall patterns. The clusters of spikes in the center violence condition between approximately the 78th and 90th most successful runs in the no violence condition suggest that a crucial element in these successful exploitations of violence by the Challenger identity is the capture of crucial, centrally important elements of the authority structure that are then harnessed for the expansion and consolidation of the Challenger identity position (see note 14).

Conclusion

My expectation, as stated earlier, was that violent strikes would disrupt the regime's ability to enforce activation on red inside Auth_Shape and decrease the integrity of regime by reducing the number of influentials activated on the regime identity. This hypothesis was in general borne out by the data presented in Figures 7, 8, and 9. However, we do observe an interaction effect in Figure 9 such that when violence is located in the periphery its impact is, on average, to enhance the Regime identity's control of the authority structure rather than degrade it. One possible explanation is for this effect is that with the center of the authority structure safely within the hands of the Regime identity, violence in the periphery interferes more with the ability of alternative identities to capture and hold low echelon influentials than it does on the ability of the Regime identity, drawing on its intact center network of influentials, to hold or recapture those southwestern Lackeys lost due to the immediate effects of the violence.

The data in Figures 7, 8, and 9 also support the expectation that under volatile conditions the regime would be more successful on each measure (prevalence and regime integrity) than under stable conditions, since rival identities would not benefit from the steadying presence of their own web of influentials. I expected that the presence of the Challenger identity would multiply the effects of violence, but was not sure what the interaction would be between volatility and the presence of the Challenger. Data in Figures 10 and 11 suggest that violence in the center interacts with poor general conditions (with volatility held constant) to create circumstances likely to decrease Regime identity success whether a Challenger is present or not. The same interaction effect is observed in data reported in Figures 12 and 13 regarding regime integrity. In

line with expectations, the presence of a Challenger identity (Figures 11 and 13) multiplied the power of this effect.

Finally, as noted prior to the presentation of experimental results, I began this investigation uncertain with regard to how volatility would interact with the presence of a Challenger. We have seen that in the presence of a Challenger volatile conditions tend to bifurcate the distribution of outcomes, leading to proportionally more extreme values of Challenger success and failure. Indeed the results in this regard are complex, appearing somewhat sensitive to the location of the violence. These interactions require further study. Nonetheless, in addition to the bifurcation in the data referred to above, prevalence data displayed in Figure 18 indicates that while generally volatile conditions assist the Regime Identity when the Challenger is present, generally stable conditions assist the Challenger identity.

Overall, these findings suggest that the definition and operationalization of violence presented here do not produce drastically unfamiliar or bizarre patterns of outcomes but do suggest stories about how violence and its strategic location, the presence or absence of latent challengers, and the volatility or stability of general conditions may interact to produce particular expectations about the political effects of violence. This is a fundamentally unsurprising result. When conducting substantive research this usually registers as uninteresting. However at this point in the research program, when definitions and operationalizations are being developed and tested, the absence of surprise may be reasonably interpreted as corroborative of the potential value of the conceptual apparatus, the analytic utility of the definition of violence here advanced, and the specific technique for its operationalization I have employed.

On a conceptual level we may ask why this way of thinking about violence, as a sudden and drastic increase in the scale of negative values at stake in an encounter, should work to capture what we want to study as political violence. One line of argument might be that the essence of what distinguishes unpleasant or unfortunate encounters that are not violent from those that are is the sense, on the part of the observer (whether victim or not) that enormous uncertainty about what might about to be lost is compressed into a small space of time. When the fist is heading to the face, neither the puncher nor the about-to-be-punched can know with any certainty just what scale to use to evaluate the stakes of this encounter. This unsettling shock will pass quickly in time. The question is the effect of this kind of shock, or of violence, after it ceases. How do traces of the existence of severe and compressed uncertainty and threat in a particular space or population impact political outcomes? This is the central question a theory of political violence must answer. The operationalization explicated and deployed in this paper is designed to help lay a firm foundation for such a theory.

¹ Dictionaries consulted include Webster's Third International Dictionary; Oxford's English Dictionary (second and revised editions); Oxford American Dictionary of Current English; The American Heritage Dictionary of the English Language; Merriam-Webster's Online Dictionary; Heinle's Newbury House Dictionary of American English; Encarta World English Dictionary, North American Edition; Cambridge Dictionary of American English; and several online dictionaries such as Merriam-Webster's Online Dictionary; Wikipedia; wordnet.princeton.edu/perl/webwn; cmpda.ca/jsp/v-information.jsp; www.austin.cc.tx.us/audit/Glossary/.

² Hannah Arendt, On Violence (New York: Harcourt, Brace & World) 1968.

³ Charles Tilly, "Memorandum: Large-Scale Violence as Contentious Politics," Workshop on Contentious Politics, March 21, 2000. I am grateful to Christian Davenport for providing me with a copy of this memorandum.

⁴ American Political Science Association Annual Meeting, Chicago, Panel in the Comparative Politics Section, "Civil Wars and Violence." Panel participants included Macartan Humphreys, Robert Bates, Ethan Buena de Mesquita, Jeremy Weinstein, and Eric Dickson. Similar answers to this author's same query were received at interdisciplinary conferences on political violence and civil wars held at the Santa Fe Institute, January 16-18, 2003; and at a workshop on "Civil War and Peace-making in Colombia" at the University of Chicago, November 13, 2004.

⁵ Mary R. Jackman, "Violence in Social Life," Annual Review of Sociology, Vol. 28 (2002) p. 388.

⁶ Ibid., p. 20.

⁷ Ian S. Lustick, Simulating the Effects of Israeli-Palestinian Violence, Fundamentalist Mobilization, and Regional Disruption on Regime Stability and USA-Friendly Outcomes in Middle East Polity, January 2003, <http://discuss.santafe.edu/files/politicalviolence/lusticksantafe.pdf>

⁸ Arendt, On Violence, op. cit., p. 4.

⁹ Ibid., p. 5.

¹⁰ I would prefer NOT to include physicality in the definition of violence in order to preserve opportunities to theorize the political implications of specifically *physical* violence and the conditions under which physical coercion, per se, might not be experienced as violence or have the same effects.

¹¹ The definition of violence operationalized here assumes that negativity associated with different kinds of events or perceived potential events can be scaled in some way analogous to the Richter scale for comparing seismic events. For a fascinating exercise in the development of a scale of "disasters" using real data to translate distributed "stress" into a logarithmic scale placing parking tickets and World War II on the same continuum, see Harold D. Foster, "Assessing Disaster Magnitude: A Social Science Approach," The Professional Geographer, Vol. XXVIII, no. 3 (August 1976) pp. 241-247.

¹² Each agent class in this simulation updates on even time steps only, so a two time step punctuation is equivalent to a single update cycle.

¹³ Updating routines require each agent to count the number of agents in its Moore neighborhood activated on different identities. Each identity receives one identity weight "point" for each agent, including self, activated on it, multiplied by the respective influence level of each agent. To the sum of these products is added the exogenously and randomly changing "bias," which varies between -3 and +3 for all identities except for identity 20 (which as explained is permanently assigned a bias of -10). The resulting sums for each identity are compared by each agent to determine which update, if any, to implement.

¹⁴ Additional variability is introduced into the futures of *Violdef* by inserting before the standard 300 time-step "run" a short, eight time-step, interval during which the bias values associated with each identity, i.e. the external signals related to their relative attractiveness, are shuffled. This maintains the distinctiveness of each of the 100 streams of perturbations used to produce experimental runs in each treatment condition while removing an artificially homogeneous beginning point at which each identity is treated as having a "0" bias. In this way the existence of distinct "histories" preceding the unfolding future is simulated.

¹⁵ For technical reasons this is specifically true across conditions using either bias volatility 250 (stable) or bias volatility 500 (volatile) but not across conditions using different bias volatility settings.

¹⁶ Stability vs. volatility of general conditions is operationalized by changing the probability that at any even time step (when updating occurs in this model), any one particular identity will be eligible for a random change in its bias assignment. When eligible, a fresh bias assignment is made via a random draw from available values. In this series of experiments the bias range was set at $-3,+3$, so available values were $-3, -2,-1, 0, +1, \text{ and } +2,+3$. Volatility settings are expressed as fractions of 10,000. Thus a "stable" setting was implemented as 250, meaning that each identity at every update had a .025 probability of being eligible for a fresh bias assignment, while a "volatile" setting, implemented at 500, means that each identity at every update had a .050 probability of being eligible for a fresh bias assignment. Note that with seven available values in the pool, each fresh assignment of a bias entailed a probability of approximately .14 that there would actually be no change in the identity's assigned bias.

¹⁷ I wish to thank Rumi Morishima for the extraordinary assistance she provided in the process of data analysis.

¹⁸ It is important to remember that this manipulation involved no additional *activation* of identity 16, only a substantial increase in the availability of the identity within the repertoires (or "subscriptions") of the Auth-Shape agents not allied with the regime by virtue of the presence in their repertoires of identity 0.

¹⁹ Indeed the rate at which high echelon influentials (Great Leader, Lord) are captured by the Challenger identity is, in both volatile and stable conditions, more than twice as high when violence is located in the center than when it is located in the periphery.