

CYCLES OF CONFLICT

**A Computational Modeling Alternative to
Collins's Theory of Conflict Escalation***

By Kent McClelland

Grinnell College

Forthcoming in *Sociological Theory*

FINAL DRAFT — December 30, 2013

Key Words: agent-based models, perceptual control theory, systems, time dynamics

***Corresponding Author:** Kent McClelland, Department of Sociology, Grinnell College, 1210
Park St, Grinnell, IA 50112

Acknowledgements: The author is grateful to Peter Hart-Brinson, Maxwell Leung, Michael Thompson, and Neil Gross, as well as to the anonymous reviewers, for their insightful comments and suggestions on earlier drafts of this manuscript. I also thank Stephanie Peterson and Dag Forssell for their assistance with the graphics.

ABSTRACT

In a new theory of conflict escalation, Randall Collins (2012) engages critical issues of violent conflict and presents a compellingly plausible theoretical description based on his extensive empirical research. He also sets a new challenge for sociology: explaining the time dynamics of social interaction. However, despite heavy reliance on the quantitative concept of positive feedback loops in his theory, Collins presents no mathematical specification of the dynamic relationships among his variables. This article seeks to fill that gap by offering a computational model that can parsimoniously account for many features of Collins's theory. My model uses perceptual control theory (PCT) to create an agent-based computational model of the time dynamics of conflict. With greater conceptual clarity and more wide-ranging generalizability, my alternative model opens the door to further advances in theory development by revealing dynamic aspects of conflict escalation not found in Collins's model.

CYCLES OF CONFLICT:
A COMPUTATIONAL MODELING ALTERNATIVE TO COLLINS'S THEORY OF
CONFLICT ESCALATION

In his 2011 Presidential Address to the American Sociological Association, Randall Collins offered a new theory of the dynamics of group conflicts (2012). Building on his own extensive investigations of episodes of interpersonal violence and warfare (2008), Collins presented conceptual models of “conflict escalation” and “de-escalation” (in his words, “C-Escalation and D-Escalation”) in violent struggles between opponents (2012:4, 11). This new theory represents a significant achievement in several respects. In a world fraught with internecine ethnic, incessant threats of terrorism, and governments deadlocked in ideological battles, the topic itself is a crucially important one for sociology.

Furthermore, his new theory is supported by substantial empirical data. In devising the theory, Collins has drawn upon his own extensive research into incidents of violence and social conflict, including investigations into the “the micro-sociology of violence,” which have involved analysis of data gathered by a variety of empirical methods, including the collection of “photos and videos, . . . participants’ detailed accounts, ethnographic observations, forensic reconstructions (*e.g.*, bullet paths and number of shots fired), data on bodily physiology, and subjective phenomenology” (2012:4), all of which he described in considerable detail in a recent book (2008; see also Collins 2009). In addition, his recent research projects have examined processes of victory and defeat in battles (Collins 2010) and displays of ritual solidarity

following the 9-11 terrorist attack (2004). In short, his theory rests on a more solid empirical foundation than much of the other theorizing in contemporary sociology.

Collins's new theory also represents a significant achievement because he has implicitly upped the stakes for sociological explanation by focusing on particular incidents of conflict escalation, not just broad social patterns, and thus has taken the time dimension of social interaction seriously. In his paper, Collins portrays the "time dynamics" of conflict as a challenging new area for sociological research (2012:13). With some exceptions,¹ the bulk of the empirical research published in the field of sociology has relied on cross-sectional analyses of statistical data, in which the time variable is not explicitly considered. Statistical studies that use longitudinal designs typically analyze cross-sectional relationships between variables at several points in time, rather than examining the ongoing flow of events. Theorists and researchers in the social micro-interactionist tradition have tended to use either qualitative observational methods, in which time is considered only implicitly, or laboratory experiments that probe relationships between variables without investigating changes over time. It has mostly been left to historically oriented sociologists to focus on issues of time, but their analytical techniques have usually involved comparison of different societies at different times, rather than a close examination of the unfolding of particular historical events (but see Sewell 2005). Thus, time, as a continuous variable, has been neglected by sociologists, and Collins, to his credit, has raised that issue.

Another way in which Collins's new theory is noteworthy is in his recourse to the language of systems theory, an approach that was regarded as a promising new development in sociological thinking sixty years ago, but has since faded from popularity among sociologists, for reasons not necessarily connected to the empirical or theoretical value of the perspective.² Collins makes use of the systems-theory idea of positive feedback loops for linking variables in

his model, depicting group solidarity and ideological polarization, for instance, as leading to an increase in conflict, which in turn increases the solidarity and polarization of the group. Thus, conflict is seen in his model as likely to result in the kind of runaway intensification that characterizes arms races and other conflict spirals, although Collins also offers a model of de-escalation with factors that lead to diminished group solidarity and, thus, the diminution of conflicts. Systems models, with their emphasis on feedback loops, are particularly useful for understanding the behavior of systems of variables in which reciprocal causation means that the changes in one variable affect another and vice versa, as when escalation by both sides in a conflict provokes answering responses from both. The conventional approach of causal modeling, still used in most statistical research (as well as most theoretical thinking) in sociology, is much less well suited for handling reciprocal causation, so Collins's resort to systems thinking makes sense for describing the incidents of conflict escalation that are his primary concern.

One advantage of framing models in terms of systems thinking is that systems theory has a well-developed mathematical basis, which allows analysts to use computational modeling and computer simulations to explore the implications of a theory. From his writings, it appears that Collins has long been intrigued by just this possibility. For instance, in discussing the positive feedback loops in his model of conflict escalation, he remarks, "Notice that all feedback loops in the model are positive. If we were to do a computer simulation, conflict would escalate to infinity" (2012:10). In an earlier book, Collins (1992) devoted a final chapter (called "Can Sociology Create an Artificial Intelligence?") to imagining how one of his earlier theoretical models would work as a computer simulation. Despite his apparent interest in this approach, however, Collins has not yet taken the step of representing any of his models in the form of

computer simulations.³ The purpose of this paper is to explore precisely that possibility by constructing a computational model for the process of conflict escalation.

In this paper, I use a variation of systems theory to construct a multi-agent computational model of dynamic social interaction that shows how the conflict-escalation processes described by Collins can be generated in computer simulations. Like his, my model relies on feedback loops, but the mathematical formulas in my model use negative feedback loops, rather than positive feedback loops, to generate the collective processes of positive feedback described in Collins's model of conflict escalation. My analysis relies on perceptual control theory (PCT), a dynamic-systems model of human behavior, which proposes that neural circuits in the brain are organized into hierarchies of negative-feedback control systems, and that individuals use these control systems to manipulate their own environments in order to control the flow of perceptual input in accordance with their internally generated preferences and expectations (Powers [1973] 2005, 2008). This psychological paradigm has provided the conceptual basis for a significant body of research in sociology and social psychology (see McClelland and Fararo 2006, Robinson 2007), and the modeling results presented here are extensions of simulations reported by McClelland (2004, 2006).

My paper has five more sections, beginning with an overview of Collins's theory of conflict escalation and de-escalation—which takes the familiar form of verbal propositions and schematic diagrams—and a critique of its conceptual adequacy for serving as the basis of computer simulations. The second section surveys literature on computational modeling, enumerating the reasons that advocates have given for using agent-based computational modeling for constructing theories of dynamic social processes. Section three describes the computational model I use in simulating conflict-escalation processes and explains the

theoretical perspective from which the model is derived. The fourth section reports the results of these simulations and discusses how my results compare to the theoretical conclusions drawn by Collins. The final section reviews convergences and divergences between the implications of the two models and asks which set of implications is better supported by empirical data. My paper concludes with an evaluation of whether the expected advantages of computational modeling have been demonstrated in this comparison.

COLLINS'S MODEL OF CONFLICT

Overview of His Model of Conflict Escalation

Positive feedback loops abound in Collins's theory of conflict escalation (2012). He illustrates his theory with seven diagrams that display hypothesized relationships between variables, all featuring one or more feedback loops; his first diagram shows only two variables, and his most complicated diagram contains twelve interrelated variables. One can get a sense his core theory by looking at Figure 1, which reproduces the third in the series of diagrams offered by Collins.

[Figure 1 about here.]

The diagram displayed in Figure 1, which Collins labels, "Escalating Conflict: Atrocities and Polarization," presents the relationships between four key variables in his model. Straight-line arrows in the diagram indicate relationships assumed to be positive, and the curved arrows show the completion of positive-feedback loops. Taking, for instance, the two central variables in his model, Collins argues, "external conflict increases group solidarity. . . . [b]ut solidarity also causes conflict" (2012:2). Similarly, he sees the occurrence of atrocities committed by one side against the other as leading to ideological polarization between the two sides, which in turn intensifies the conflict and leads to yet more atrocities (2012:3-5).

In his more complex diagrams, Collins displays this core set of variables twice, once for each side of a two-party conflict, a double-headed arrow linking the two appearances of *conflict* in the two sets of variables. The full model also contains two additional variables for each side, one labeled “mobilizing material resources,” and the other, “seeking allies, forcing out neutrals” (see Collins 2012:9, Figure 7). Causal arrows in these diagrams indicate a positive feedback loop from *mobilizing material resources* back to *conflict*, and a second positive feedback loop from *seeking allies, forcing out neutrals* back to *conflict* via *mobilizing material resources* as an intermediate variable.

Obstacles to Reformulating His Model Mathematically

Any attempt to build a simulation model directly from Collins’s model inevitably runs into the problem that his model does not easily lend itself to the precise mathematical specification necessary for constructing a computational model. While the model has considerable intuitive plausibility, Collins has relied on common-sense definitions of key variables in his model, definitions that may not always meet the test of mathematical rigor. For example, the variable *conflict*, as we see from Figure 1, is central to his model, but he offers this variable without defining it, other than by referring to the works of such venerable theorists as Simmel ([1908] 1955), and Coser (1956). While the work of these sociological pioneers offers many subtle insights about human conflict, they themselves had no ambition to render their theories mathematically—instead relying implicitly on common-sense understandings of what conflict means—and the diagrams offered by Collins do little to clarify his definition of conflict.

The variable *solidarity* presents similar obstacles to rendering Collins’s theory as a computational model. In his use of solidarity as a variable, just as with conflict, Collins has drawn on a long tradition in sociology, dating in this case back to Durkheim ([1893] 1933,

[1897] 1951), but this tradition gives scant guidance for the purpose of constructing a mathematical specification. Solidarity is evidently an attribute of a group of actors forming one side of a dispute, but is their solidarity to be defined in the instrumental sense of unity of purpose and action, or in the emotional sense of identification with the group and its shared cause? Is Collins talking about unity in goals, in actions, in emotions, or perhaps all of these? Collins cites Durkheim in arguing that solidarity “makes one willing to sacrifice oneself for the group” (2102:2), but is willingness to sacrifice solidarity or just an effect of solidarity? A micro-level schematic model included in his article, in a figure labeled, “Conflict as an Interaction Ritual,” fails to clarify the ambiguity, since it gives the causes of group solidarity as both “mutual focus of attention through common action” and “shared emotional mood” when the group is assembled face to face (2012:3, Figure 2). This ambiguity in the definition of solidarity makes the concept problematic for use in a computational model.

The other core variables in Collins’s model also lack of the kind of conceptual precision needed to support mathematical reformulation. Collins explicitly defines the variable *atrocities* as a matter of the perceptions of people on one side of the conflict interaction: “Atrocities are opponents’ actions that we perceive as especially hurtful and evil . . .” (2012:2). However, he does not offer any operational definition that would enable observers to judge whether a particular action committed by one side is likely to be defined as an atrocity by the other.

He provides no definition at all for the variable *ideological polarization*, and, although his more elaborate schematic models do not contain double-headed arrows linking the dual appearances of this variable in the core models for the two sides of the conflict, it is not entirely clear how this variable can be conceptualized as applying to one side of a dispute separately from the other. In its common-sense meaning, polarization refers to positional differences between the

two sides of a conflict, not just the position of one side or the other, although perhaps by polarization Collins instead means the extent to which each side takes an extreme position along some ideologically defined continuum.

In sum, by constructed his theory primarily in verbal terms, Collins has relied for his definitions of key variables on a sociological tradition of theorizing that has emerged from common-sense discourse about social problems, and the resulting definitional ambiguity seriously limits the possibility of constructing more tightly defined computational models based on his theory. Models in the form presented by Collins simply do not provide suitable starting points for mathematical reformulation and computer simulation. Because the traditional approach of verbal definitions and schematic diagrams lacks the conceptual precision needed for modeling dynamic processes in terms of computer simulations, we must turn instead to some other form of theoretical modeling.

An Alternative Approach: Agent-Based Computational Modeling

My alternative approach to modeling the processes of conflict escalation described in Collins's model uses agent-based computational modeling, specifically a multi-agent mathematical model in the systems-theory tradition based on perceptual control theory (PCT) (Powers [1973] 2005, 2008). My model diverges in many respects from the model offered by Collins, most notably in substituting pairs of *negative* feedback loops for the *positive* feedback loops that figure so prominently in his model. Nevertheless, I will argue that my computational model covers much the same ground substantively as his theory and that several of his key variables can be subsumed into the parameters of my model.

Despite rapid advances in computer technology in recent decades, the approach of using agent-based computational models to build theories has not yet gained much popularity among

sociologists, as can be seen from a scan of recent articles in the field's leading journals.⁴

Nevertheless, proponents of agent-based modeling argue that this approach has several advantages over more traditional styles of modeling. The benefits attributed to agent-based modeling include greater dynamism, added insight into macro-micro links, greater realism, and scientific merits in comparison to other approaches to theory construction. All of these advantages contribute to making agent-based computational models an attractive choice for modeling the processes of conflict escalation that Collins described.

Proponents of agent-based modeling argue, first, that these models are dynamic, not static, in contrast to both cross-sectional models of statistical relationships—still the norm for reporting sociological research results—and the mathematical, equation-based models favored by most economists (Gilbert 2008). Because agent-based models focus on dynamic interactions between components of systems over time, advocates describe them as “process oriented” (Miller and Page 2007:80). The models can also apply to social systems that are far from equilibrium, unlike conventional economic models of markets. This dynamism of agent-based models, as well as their process orientation and usefulness for describing situations in flux, makes them especially useful for the analysis of volatile episodes of conflict escalation like those described by Collins.⁵

A second argument put forward by proponents of agent-based models is that they provide a way to explore macro-micro links (Epstein 2006; Raub, Buskens, and van Assen 2011; Conte et al. 2012). Because these models allow multiple scales of analysis, the simulations of micro-level interactions of agents can reveal emergent social patterns at the macro level (Gilbert 2008, Miller and Page 2007). Statistical models that focus on averages among groups of individuals do not allow for comparably detailed analyses, nor do other macro-level modeling techniques, such

as economists' mathematical models of market equilibrium or dynamic simulations employing classical systems analysis (e.g., Meadows et al. 1972). Since episodes of conflict escalation emerge from interactions at the micro level, the potential of agent-based models to combine micro and macro scales of analysis provides another compelling reason to choose this method for modeling processes of social conflict.

A third advantage cited by advocates of agent-based computational models is that they can be constructed to be more realistic than other kinds of models. The agents in these models do not have to be modeled as entirely uniform in their knowledge and capabilities and thus can be given heterogeneous characteristics. Furthermore, in contrast to game-theory models and most other economic models, the agents need not embody the psychologically unrealistic assumptions of rational choice theory, but can be modeled as having bounded rationality and limited knowledge (Epstein 2006, Gilbert 2008). Of course, every approach to modeling relies on simplifications and abstractions, and degrees of realism exhibited by agent-based models will vary with the level of detail built into the modeled agents and assumptions made about the ways they interact. The agents I offer in this paper are based on perceptual control theory, a psychological perspective that is arguably more plausible than rational choice and better supported by empirical evidence (e.g., Bourbon 1990, Marken 1980, Powers 1978). They are modeled as having heterogeneous characteristics, and, although they embody an extremely simple mathematical model, they are capable of realistically generating some of the empirically observed behavior patterns that Collins describes.

Scientific rigorousness is yet another advantage of agent-based models, according to their advocates. Because the models are implemented as mathematical algorithms in computer programs, this approach forces analysts to be more precise in defining variables and describing

relationships than is true of nonmathematical approaches (Gilbert 2008). Some proponents argue that agent-based models actually produce better-quality explanations, because the social patterns being studied are generated by the model, not merely described (Epstein 2006), and the models offer viable mechanisms by which empirically observed patterns could have occurred (Hedström and Ylikoski 2010). The flexibility of agent-based models also makes virtual experiments possible, in which the analyst, by exploring the results of simulations with various combinations of parameters, investigates what-if scenarios impossible to test in the field (Marchioni and Ylikoski 2013). And the same flexibility allows investigation of the robustness of a model's results across various combinations of parameters. Moreover, one can fit models against empirical data, assessing the closeness of the fit (Epstein 2006). Finally, advocates of agent-based models say that this approach to modeling parallels recent trends in the hard sciences, as scientists are "opening up the black box" (Hedström and Ylikoski 2010:51) by using computational modeling tools to explore the inner mechanisms of the processes they study. They argue that the image of science as a quest to find broad covering laws to explain observations is largely passé among natural scientists, even while retaining its currency as an assumption implicitly underlying the statistical research techniques used by social scientists (Powers [1973] 2005:11). In sum, if the goal is to construct scientifically rigorous models, its advocates contend that agent-based modeling offers a better means than conventional approaches for reaching that goal.⁵

Not every sociologist, of course, thinks that making sociology more scientific is a worthwhile goal, but it appears that Collins himself seeks to move the discipline in that direction. Because of the scientific rigorousness of agent-based modeling, and because it offers additional advantages such as potentially greater realism in its micro-level explanations of dynamic

processes, agent-based modeling offers a scientifically attractive alternative to the conventional style of theoretical modeling employed by Collins. After presenting my computational model and examining the substantive results of my simulations, my article concludes with a discussion the comparative advantages and disadvantages of these two approaches to theoretical modeling in light of the results of this exercise in model building.

THE PERCEPTUAL CONTROL THEORY MODEL

The Basic Model

Perceptual control theory (PCT) offers a conceptual model of an intentional human actor, capable simultaneously of rational calculation and emotional response. The mathematical basis of the model comes from control-system engineering, and the model portrays the neural organization of the human brain as a complex arrangement of nested layers of negative-feedback control loops. This neuropsychological theory of behavior was developed by William T. Powers ([1973] 2005, 2008), a control-systems engineer rather than a psychologist by training, but the theory has found applications in a variety of psychological research areas, ranging from animal behavior studies (Bell and Pellis 2011; Pellis, Gray and Cade 2009; Pellis and Bell 2011) to clinical psychology (Mansell and Carey 2009; Mansell, Carey and Tai 2012). In sociology, the theory has provided the inspiration for two prominent research programs in interactional social psychology, affect control theory (Heise 2007) and identity control theory (Burke and Stets 2009), as well as for research on collective behavior (McPhail 1991).

The negative feedback loop central to the PCT model describes a behavioral process that allows humans and other animals to maintain control of important variables in their own environments by acting to reduce discrepancies between their perceptions of what is occurring around them and their own goals and expectations. The theory hypothesizes that a hierarchical

structure of control systems in the human brain allows for the control of a person's perceptions of many kinds of variables, from concrete occurrences in the physical world, to rational perceptions such as logical categories and programs of action, and ultimately to highly abstract perceptions of human values and of personal and group identity (see Powers [1973] 2005, 2008; McClelland 1994).

Conflict has a significant place in this model, and it is expected to occur whenever two control systems operating in the same environment have incompatible goals. Conflicts may emerge either between two or more control systems in the same brain, as when a person dithers between equally attractive options, or between the control systems of two or more individuals, as when people acting together don't share the same goals. This second type of conflict—social conflict between interacting individuals—is explored in my models.

Figure 2 shows a schematic diagram of the building block of the PCT model: a negative feedback loop. In the figure, the area with the gray-shaded background represents the brain and nervous system of a person, while the area outside the gray shading represents the person's environment. The segment of the loop lying within the person's body corresponds roughly to the "reflex arc" as conventionally understood, with *perceptual input* from sense organs, then information processing in the brain, in which the input is *compared* to memories that serve as *reference* values for these perceptions, and finally motor *output*, as discrepancies between the perceived input and the reference values produce *error signals* that activate the muscles involved in physical responses.

[Figure 2 about here.]

The lower half of Figure 2 depicts the segment of the negative feedback loop that passes through the organism's environment. The diagram shows a person's physical *actions* as having

feedback effects on *environmental variables* that are the sources of the person's perceptions, feedback effects that are generally overlooked in conventional psychological models. Because a person's physical actions are driven by error signals, the actions tend to compensate for *disturbances*, factors in the environment that have an impact on the variables perceived. Compensation for the effects of disturbances reduces the discrepancies between the person's perceptions and the references for those perceptions, thus keeping both the person's perceptions and the perceived environmental variables in *control* by eliminating much of the variation that would otherwise occur. Of course, a person's physical actions may also have *unintended effects* on other variables in the environment besides those that are controlled.

One can demonstrate by mathematical analysis that the variables controlled by the actions of this negative feedback loop are the input variables—perceptual signals—rather than the output variables—the organism's physical actions—which must fluctuate freely to counter the effects of disturbances (see Powers [1973] 2005). Thus, disturbances created by a given stimulus may lead to many different physical responses from the organism, depending on changes in its own internal reference conditions and in other environmental disturbances, which is why psychological experiments seeking to find lawful relationships between stimulus and response may often have inconsistent results.

A PCT Model of Conflictive Interaction

The PCT model offered as an alternative to Collins's theory of conflict (2012) comprises four simulated "agents," each modeled as negative-feedback control system, all controlling the same variable in a shared environment. Figure 3 presents a schematic diagram of the model. All four simulated agents in the model have the same computational structure, a control-system model of the simplest kind: one level of control of a one-dimensional variable. In Figure 3 each of the

gray-shaded rectangles represents the “brain” of one agent, and comparing Figure 3 to Figure 2 reveals, one can see that the agents all have the same interior components as the control-system model in Figure 2, but with the orientation of components rotated 90 degrees to the left for Agents 1 and 2 and 90 degrees to the right for Agents 3 and 4. In the simulations of conflict to be presented, Agents 1 and 2 will represent one side of the conflict and Agents 3 and 4, the other.

[Figure 3 about here.]

The interactions between the two opposing sides, and also the interactions between agents on the same side of the conflict, are modeled as mediated entirely by their joint attempts to control a single variable in their shared environment. This variable represents the “stakes” of the conflict—the focus of the fighting—and we might think of it as representing money or territory or comparative prestige. From the PCT perspective, conflict occurs whenever the two sides use different reference standards in their attempts to control an environmental variable, and victory or defeat in the contest is represented by the extent to which side or the other succeeds in bringing the variable into line with its own preferred reference conditions.

Compared to the human protagonists portrayed in Collins’s theory, these simulated agents are radically simplified. They have neither behavioral mechanisms for communicating with each other nor any ability to monitor the actions of agents on their own side or the other, except for the impact of such actions on the environmental variable on which their attention is fixed. These models do not provide the complexity needed for simulating any of the higher mental powers, such as those involved in forming a perception of personal or group identity or in perceiving the actions of the other side to be an atrocity. Clearly, with such rudimentary models one cannot expect to reproduce the subtlety or sophistication of Collins’s analysis.

Nevertheless, will be shown, this drastically simplified model of interaction is sufficient to generate patterns characteristic of the basic dynamics of conflict described by Collins. Empirical research has demonstrated that even very simple control-system models can explain patterns of behavior involving perceptions as complex as coordination of movements by groups of individuals across a field (McPhail and Wohlstein 1986) or maintenance of a sense of personal identity (Burke and Reitzes 1991). And repeated studies have shown that the moderately complex actions and perceptions involved in tracking experiments—as for example when a subject uses a mouse to pursue a target on a computer screen despite disturbances—can be predicted with remarkable precision (r^2 often in excess of .98) using simple control-system models (Bourbon 1990, Bourbon et al. 1990; Marken 1980, 1986, 1988; Powers 1978; McPhail and Schweingruber 2006). In principle, it may be possible to construct more complex PCT models, capable of simulating more of the kinds of sophisticated behaviors that Collins is describing, but even simple models can capture recognizable patterns of dynamic behavior.

Each individual agent in these simulations embodies the following mathematical model of a negative feedback control system with a total of seven variables.⁶

- p* The *perceptual input signal* for the control system. For these simulations the perceptual signal is always set equal to the value of the environmental variable, as if the agent possessed perfect perception of conditions in its environment. In other words, the input function for each control system (see Figure 2) is simply an identity function.
- r* The *reference signal* for the control system. In my simulations of conflict, the reference signals for agents on Side 1 are set to positive values (or zero), while reference signals for agents on Side 2 are set to negative values (or zero).
- o* The *output signal* for the control system, a function of the difference between *r* and *p*. In these simulations the output action for the agent is taken as exactly equal to the output

signal, as if the agent's body were perfectly efficient in translating neural impulses into muscular responses. Thus, the output function for each control system, like the input function, has been modeled as an identity function.

- v* The value of an *environmental variable*, which is affected by the combined output of the four agents, as well as by the disturbance *d* (defined below). In my simulations of conflict, I will refer to this variable as the *contested variable*. Like the input and output functions, all of the feedback functions have been modeled as identity functions in order to simplify the model, as if these agents were perfectly efficient in translating their actions into physical impact on the environmental variable.
- d* The *disturbance* acting on the environmental variable, that is, the sum of all other environmental forces affecting *v*, net of the outputs of the four control systems. For the purposes of the simulations reported here, the disturbance variable is set to a constant value of zero, although it would be more realistic to model the disturbance as a vector of changing values over time, since actual confrontations between people do not take place in environments in which everything else is static. Substantively, running these simulations with a non-zero disturbance vector would not make any appreciable difference in the conclusions to be drawn from them, but graphs of simulations with a zero disturbance (as reported here) display the patterns of interactions between agents more clearly.
- g* The *loop gain* of the control system, which is proportional to the speed at which the control system corrects its errors and also corresponds to the precision with which the control system matches its perception to its reference signal.

- s* A constant *slowing factor* (also described as a “leak”) introduced to allow the accurate representation of a continuous (analog) process in the form of a discrete (digital) simulation. In all of these simulations, $s = 0.0025$.

These mathematical simulations are iterative, and the changing values of the system of variables are recalculated at each iteration over a “run” of 100 equal-time intervals from $t = 0$ to $t = 100$. Some of the variables described above, including the reference signal, the disturbance, the loop gain, and the slowing factor, are *parameters* of the model; they are either constant values or vectors of changing values that are set arbitrarily in advance of a given run of the simulation.⁷ By adjusting these parameters from run to run, the analyst can investigate the behavior of the simulation model in terms of the dynamically changing values of the remaining variables, including the perceptual signal, the output signal, and the environmental variable, which are recalculated at each iteration. Formulas for each of the recalculated variables for agent k at time t are as follows, where o_{kt} is the output value, p_{kt} is the perceptual signal, and r_{kt} is the reference signal:

$$p_{kt} = v_{t-1}$$

$$o_{kt} = o_{kt-1} + s \{g (r_{kt} - p_{kt}) - o_{kt-1}\}$$

$$v_{kt} = o_{1t} + o_{2t} + o_{3t} + o_{4t} + d_t$$

The variables in this simulation model do not correspond exactly to any of the variables in Collins’s model of conflict, but, in any case, his variables were not defined precisely enough to support the construction of a computational model. Thus, to create simulations descriptive of the conflict processes modeled by Collins requires some creative interpretation of the parameters

of my PCT model. Here are my definitions for computational analogs for the core variables in Collin's model:

Conflict: In models of control-system interactions, conflict is not a variable but rather an interactive outcome that inevitably occurs when two or more control systems attempt to control the same environmental variable using different reference standards. In his recent article on "the micro-sociology of violence" Collins (2009) argues that violence can result when "individuals (or groups) confront one another at cross purposes" (p. 571), a description of conflict that fits well with the definition of conflict that I am using in these simulations. When such conflicts occur, PCT models show the output of the two interacting systems diverging, with one system, in effect, pulling in one direction to bring the environmental variable into line with its preferences, while the other system pulls in the other (see McClelland 2004, 2006). Hence, the variable used in my simulations is the *intensity* of the conflict, measured by the degree of divergence in system outputs.

Solidarity: This variable has often been taken to refer to feelings of unity or agreement within a group, as well as the willingness of group members to take action in support of the group. Collins, as I have noted, leaves the definition of solidarity implicit in his model, although he notes that in conflict situations solidarity "makes one willing to sacrifice oneself for the group" (2012:2). While a definition in terms of self-sacrifice is well beyond the scope of my rudimentary simulation model, a definition of solidarity in terms of agreement between agents is possible. For purposes of these simulations, the *solidarity* of the agents on one side of the conflict in these simulations will be defined as the extent of convergence of their reference signals. Furthermore, if the agents have similar reference signals, we will say that solidarity increases as the loop gain of the lower-gain agent increases to match that of the higher-gain

agent, which might be taken as an increase in the agent's willingness to act on behalf of the group.

Polarization: For purposes of these simulations, I will take polarization to refer to the extent of differences in reference values between the two sides, rather than their positions along some ideological continuum.

Atrocities: Because of the complexity of this variable, I have not tried to model it directly in these simulations. Rather, it is modeled indirectly in two ways: first, in terms of its presumed effects on the polarization variable, by increasing (in absolute value) the reference values of agents presumed to be reacting to a perceived atrocity by the other side; and second, by increasing the loop gain, and therefore the "effort" expended in pursuing the conflict, of agents presumed to be reacting to an atrocity.

SIMULATION RESULTS

Conflict

My first goal in presenting results of simulations using the PCT model is to show how this model generates the positive-feedback loops that in Collins's view are fundamental patterns of conflict. Figure 4 shows the behavior of the PCT model in a conflict situation. In this simulation, the conflict occurs between Agent 1 and Agent 3, with Agents 2 and 4 assigned a zero loop gain and thus, in effect, sitting on the sidelines. The horizontal axis indicates time, and the units of time shown here are the 100 iterations over which the computational model is recalculated. The vertical axis indicates units of the *Contested Variable*, the environmental variable that the agents are jointly controlling. Because the input, output, and feedback functions of the control systems have been modeled as identity functions, the vertical axis can also be used to graph the values of the reference signals of the agents, as well as their outputs and the joint impacts of those outputs

on the contested variable. In this simulation Agent 1 has been given a loop gain of 40, and Agent 3, a loop gain of 10. Agent 1 has been assigned a reference signal of 20 points in the positive direction, while Agent 3 has a reference of negative 20. Because Agents 1 and 3 both attempt to control the same variable, but with different reference values, the interaction produces conflict. Agents 2 and 4 are inactive, so that their output values remain at zero throughout the simulation.

[Figure 4 about here.]

The sharp divergence in Figure 4 between the outputs of the two active agents is a characteristic signature of control-system conflict. As one agent pulls in the positive direction, the other pulls in the negative direction, so that the efforts of one agent are largely counteracted by the other. Looking at the first 30 iterations of the simulation, we see that Agent 1, because it has higher loop gain, begins more quickly than Agent 3 to correct the discrepancy it perceives between the initial value of the contested variable and its preferred reference value; in effect, Agent 1 has gotten the jump on Agent 3 and has gained the upper hand in the struggle for control of the contested variable. However, as the gap between the value of the contested variable and Agent 3's preferred reference value grows larger, Agent 3 begins to pull harder in the negative direction, as seen in the slight downward concavity of the curve for Agent 3's output. At the same time, the gap narrows between the contested variable and Agent 1's reference, so that the rate of increase of Agent 1's output slows down, indicated by a downward concavity of the curve. An important point to note is that a control system's response, in terms of output, is always proportional to the size of its perceived error, so that, as the difference between its perceptual signal and its reference signal decreases, the rate of error correction decreases, as well.

At about iteration 40, interaction of the two agents reaches an equilibrium point, with the value of the contested variable ending up considerably closer to Agent 1's reference than to Agent 3's, and the contested variable stays at virtually the same value for the remainder of the simulation. Even after this equilibrium point has been reached, and Agent 1, the higher-gain agent, has seemingly gotten the better of the contest, their outputs continue to diverge, thus continuing to intensify the conflict between them, with both agents still striving to bring the contested variable more nearly into line with their own preferred positions. After the contested variable has reached its equilibrium point, each increase in the output of one agent is matched by a nearly equal increase in the output of the other, and the value of the contested variable remains almost unchanged, so that, in terms of "facts on the ground," the conflict has reached a stalemate, even though the intensity of the conflict continues to escalate. Neither party can relax, because each side still perceives a gap between its aspirations and the current situation. Moreover, the contest is still precarious at every moment, because if either side were to stop the escalation unilaterally, its position relative to its own goal would begin to erode.

In sum, a lot is happening in this simple simulation. Winning and losing occur, in addition to the stalemate. Agent 1 emerges the winner, because it does much better than Agent 3 in reaching its goal of closing the gap between the position of the environmental variable and its reference value. By the same token, Agent 3 ends up a clear loser, far less successful in reaching its goal. The stalemate that ensues favors Agent 1, because that agent has more nearly succeeded in bringing the shared environment into line with its preferences. Finally, from the point of view of an outside observer, this interaction looks exactly like a positive feedback loop, as escalation by one side provokes counter-escalation by the other and the intensity of the conflict continues to grow.

Although this positive-feedback pattern is exactly what Collins (2012) predicts with his model of conflict, my simulation demonstrates that a PCT model can generate this pattern of escalation without any mathematical specification actually involving positive feedback. The positive-feedback pattern emerges directly from the conflictive interaction between the two opponents. The other variables that Collins includes as mediating variables in his positive feedback loops, such as group solidarity, atrocities, and ideological polarization, appear not to be necessary for producing escalation of conflict within the PCT model.

Although the PCT model, without any of the intervening variables from Collins's conflict model, produces the positive-feedback pattern of conflict, these additional variables were developed from his own extensive investigations into incidents of face-to-face conflict (2008). Thus, his contention that these variables make a difference in conflict interactions has strong empirical support. I turn next, then, to exploring PCT analogs of the other variables in Collins's model.

Solidarity

The solidarity variable is also central to Collin's model, but the *solidarity* variable is perhaps the least clearly defined of the model's core variables, and this conceptual ambiguity makes finding an analog for the variable within the PCT model less than straightforward. As I noted earlier, solidarity can connote either unity of purpose or else, as Collins puts it, "willingness to sacrifice" (2012:2). In the context of the PCT model, when agents on one side of a conflict all share the same or closely similar reference signals, we can regard them as displaying unity of purpose, and I will take this as my operational definition of solidarity. For an alternative definition of solidarity, willingness to sacrifice might be indexed by the extent to which the agents involved

on one side of a conflict all have high loop gain, and therefore devote their attention and energy to the conflict, rather than pursuing other goals.

My previous simulation involved only two agents, one on each side of the conflict. To illustrate group solidarity, however, one must have a group, or at least more than one participant on each side. But given the mathematics of the PCT model, an agent modeled as a single control system can be interpreted to represent either an individual or a collective actor. In terms of the control exerted on an environmental variable, one high-gain system working alone can have exactly the same effect on a contested variable as several lower-gain systems working together. The three panels of Figure 5 illustrate this principle.

In panel 5A, we see the individual behavior of Agent 1 (defined as in the previous simulations), when it faces no opposition, because the loop gain of the opposing agent has been set to zero. Encountering the initial gap of 20 points between the condition of the contested variable and its own preference, Agent 1 moves rapidly to correct its perceptual error and bring the variable into line with its reference value. As the value of the variable nears the goal of 20 points on the scale, the speed of its error correction slows down, until the variable comes to a stable value that almost, but not quite, reaches 20 points. Specifically, by iteration 100 in this simulation, the contested variable has been brought a value of 19.51176, with each successive iteration bringing the value closer to the goal of 20 points by a tiny increment in the fifth decimal place. Because Agent 1 has been assigned the relatively high loop gain of 40, the control system can do a good job of bringing the controlled environmental variable to its reference value, but control systems can never eliminate error entirely. Nevertheless, within certain limits of stability, the higher the loop gain, the more nearly a control system can succeed in reaching its goal.

[Figure 5A about here.]

With the PCT model, when multiple control systems use different reference values in attempting to control the same variable, their joint efforts will succeed in stabilizing the variable in spite of the ensuing conflict between them. The equilibrium point we saw in Figure 4 illustrates this stability of a contested variable despite conflict. The equilibrium position emerging from this kind of collective control, however, is a compromise value based on an average of the reference values of the participating systems, weighted by their loop gains. Comparison of Figures 4 and 5A can further illustrate this point. In Figure 5A, as was just mentioned, Agent 1 brings the environmental variable into control at a value of approximately 19.51, just shy of the reference value of 20 points. In Figure 4, in which Agent 1 must contend with Agent 3 for control of the contested variable, the equilibrium point reached is 11.76. Although Agent 1 with its higher loop gain has done a better job of approximating its goal of 20 points than Agent 3, with its goal of -20 points, the value reached as a compromise falls far short of the value that Agent 1 can attain when possessing unfettered control (as in Figure 5A). Thus, the compromises emerging from conflictive interactions satisfy none of the participants.

Panel 5B illustrates the unsatisfactory nature of compromises in another way. In this simulation, all four of the agents in the model have been enlisted to work together. Each agent has been given a loop gain of 10, and Agents 1 to 4 have been assigned reference values of 30, 25, 15, and 10, respectively, a set of values that averages to 20 points. Although the agents all have positive reference values and thus are all ostensibly working together, we see in Figure 5B that Agents 1 and 2 appear to be in conflict with Agents 3 and 4, the two agents with the lowest reference values. As soon as the value of the contested variable exceeds their own reference values, those agents start pulling in the negative direction. One of the apparent paradoxes of the PCT analysis is that agents cooperating on a common task are predicted nevertheless to come

into conflict, unless all their reference values are precisely identical (see McClelland 2004, 2006). If we consider how frequently in ordinary situations individuals who attempt to work together come into conflict, and how difficult it can be to get everyone in a group of people on the same page, this finding begins to look less paradoxical.

[Figure 5B about here.]

The most important thing to notice about Figure 5B is that the curve for the value of the contested variable is exactly the same as it was in Figure 5A. The four agents with an average (loop gain-weighted) reference value of 20 have the same impact on the environmental variable as a single agent with that reference value and a combined loop gain equal to the sum of the four. Thus, conflict does not preclude stability of outcome.

Figure 5C demonstrates the effects of collective control in yet another way. As in 5B, all four agents are working together, this time with perfect cooperation. All four have been given the same reference value of 20, and once again the curve for their control of the contested variable is identical to those in panels 5A and 5B. However, in this case the outputs of all four agents coincide at level that is markedly lower than the output of the single agent in Figure 5A. This graph clearly demonstrates the benefits of cooperation in the sense that when the four agents are working together each one has to contribute only one-fourth as much output as would be necessary for a single agent working on its own to reach the same goal.

[Figure 5C about here.]

Returning to the question of how to represent solidarity within the PCT model, it seems reasonable to argue that panel 5C shows an instance of high group solidarity, since the agents' reference values for control are all identical, even though none of them has a particularly high loop gain. If all four agents had been given a higher loop gain, such as that assigned to Agent 1

in panel 5A, they would be even more effective in working together to bring the contested value up to their agreed-on reference value and then in holding it close to that value; in such a case, one might say that solidarity had increased. Thus, solidarity in the PCT model seems to be related both to agreement in reference values and to total loop gain contributed to the joint action. Looking back at panel 5B, we see that if all four agents were to come together on their most extreme reference value, 30, their solidarity, as well as the impact of their actions on the contested variable, would clearly have increased. By the same token, if Agent 1, with the highest reference value, had been assigned higher loop gain and nothing else had been changed, their performance as a group in raising the level of the contested variable would have improved, but without any obvious improvement in solidarity. And if all four agents had adopted the reference value of 10, the lowest value in their group, this unified goal would mean an increase in their solidarity, but a decrease in their collective impact on the contested variable.

My conclusion is that, in terms of the PCT model, solidarity has no clear effect on conflict intensity or on the effectiveness of a group in pursuing its goals. In any group with scattered reference points for the contested variable, solidarity improves effectiveness in conflict only when the group comes together by agreeing on a more extreme reference value than the current group average. When those with more extreme reference values are able to increase the loop gain that they contribute to the group effort, effectiveness will increase without any clear increase in solidarity. If a group achieves greater solidarity by coming together at a less extreme position, their effectiveness as a unit in conflict will diminish, although the good news is that their internal conflicts will also diminish at the same time.

This analysis suggests that Collins's model of a positive feedback loop connecting solidarity with conflict is oversimplified. Much depends on the kind of solidarity achieved.

Bringing members of a group into agreement will do nothing to escalate a conflict if they coalesce around a shared position that is less extreme than the initial average of their reference positions. It appears in this theoretical framework that the effects of solidarity depend on the level of polarization, and thus we turn next to the effects of the polarization variable.

Polarization

In exploring Collins's core variable of *ideological polarization*, I must concentrate on a simpler version of his variable, since the concept of an ideology is far too complex to be captured by my rudimentary model. However, the basic idea of polarization has a straightforward PCT analog: polarization can be indexed by the distance between the reference conditions sought by the opposing parties in a conflict. Simulations with the PCT model show that, when two control systems try to control the same contested variable, the greater the gap between their reference values, the more intense their conflict, with intensity measured by the divergence of their outputs. Figure 6 illustrates this principle by showing a two-party conflict between the same two agents as were shown in Figure 4, but this time with a smaller gap between their reference values, and thus a reduced level of polarization. In Figure 6, Agents 1 and 3 have reference values of 10 and -10, respectively, instead of 20 and -20. The vertical scale of the graph in Figure 6 remains the same as in Figure 4, and by comparing Figure 6 to Figure 4 one can easily see the reduction in the divergence of the output curves, and thus the intensity of the conflict. It is worth noting, however, that if the loop gains of the agents are high enough, even small differences in reference values between opponents can lead to large differences in their outputs (see McClelland 2004).

[Figure 6 about here.]

The conclusion to be drawn from these PCT simulations is that polarization between the parties of a conflict contributes to its intensity by increasing the rate of escalation and counter-escalation of outputs, a conclusion that agrees well with Collins's theory, when he puts the *ideological polarization* variable into one of the positive feedback loops in core of his model (see Figure 1). Nevertheless, my examination of the effects of polarization, in conjunction with my simulations of solidarity (Figures 5A, B, and C), suggests that some rearrangement of the variables in his core model might be in order. Solidarity, as we have seen, works best to increase the group's effectiveness in conflict when its members coalesce around more extreme positions. Solidarity without polarization has little effect. The model presented by Collins puts the polarization and solidarity variables into separate feedback loops (see Figure 1), but my analysis suggests that the two factors might more properly be represented as working together in the same feedback loop.

Atrocities

The final core variable in Collins's model, *atrocities*, also poses significant challenges for translation into PCT terms. In his words, atrocities are "opponents' actions that we perceive as especially hurtful and evil, a combination of physical and moral offense that we find outrageous" (2002:2). These complicated perceptions are obviously far beyond the capability of my rudimentary PCT agents to simulate. However, Collins has put the atrocities variable into the same positive feedback loop as the polarization variable. Atrocities, according to Collins, increase the polarization of members of a group, who take more extreme positions when they feel themselves to have been victimized, and the effects of such polarization on the intensity of the conflict are easy to model with PCT.

Figure 7 presents a scenario in which tit-for-tat atrocities by opposing sides in a conflict increase the polarization in reference values between the two sides and lead to more intensified conflict. For clarity in presentation, the simulation is restricted to only two actors, Agents 1 and 3, as were the simulations shown in Figures 4 and 6. The substantive conclusions would be the same, however, for simulations in which all four modeled agents were active.

[Figure 7 about here.]

In Figure 7 the simulation is based on the assumption that unspecified atrocities take place at regular intervals, every 20 iterations, first on the part of one side and then the other. While the perceptions of the incidents of atrocity are not modeled in the simulation, the hypothesized effects of the atrocities in terms of increased polarization are shown. Each instance of an atrocity is followed by an increase (in absolute value) of 10 points in the reference signal for the agent supposed to have perceived the atrocity. The first atrocity, then, takes place at iteration 20, and the reference value for Agent 3, representing the victimized side, jumps from -20 to -30 . Atrocity number two, at iteration 40, results in a jump in the reference value for Agent 1, from 20 to 30 , and so forth. After two more exchanges of atrocities, the reference values of Agents 1 and 3 end the simulation at 40 and -40 , respectively.

As expected, the pace of intensification of the conflict in Figure 7 increases after each of the designated atrocities. These increases are easiest to discern in the output of Agent 1, modeled to have a loop gain of 40, in comparison to Agent 3's loop gain of 10. The effects of the jumps in polarization can also be seen in the curve for the contested variable. From iterations 1 to 20, the efforts of Agent 1 prevail over those of Agent 3 and bring the contested variable almost to the equilibrium point that was achieved in the simulation shown in Figure 4. From iteration 21 to 40, the jump to a more polarized reference value by Agent 3 slightly increases its output (in the

negative direction), which means that the value of the contested variable begins to decline, only to start moving upward again toward a higher equilibrium point when Agent 1 jumps to its more polarized reference value at iteration 41. The same patterns repeat at iterations 61 and 81. If atrocities increase polarization, the greater the polarization, the more intense the conflict.

An alternative theoretical possibility is that atrocities could be represented in a PCT model as leading to increases not in polarization, but in loop gain. The idea here is that, having an emotional response to a perceived atrocity, members of the victimized side would redouble their efforts to win the contest, without necessarily changing their goals of what winning would mean. Thus, each atrocity by the other side would be followed by a jump in the loop gain that actors are applying to their control loops. The outcome of this alternative specification, when tested, is almost indistinguishable from the outcome shown in Figure 7. A simulation (not shown) in which increases in the “victim’s” loop gain following each hypothesized atrocity were substituted for jumps in polarization produced patterns of increasing intensity in conflict and shifts in the value of the contested variable that were only subtly different from those shown in Figure 7. Since neither the reference values nor loop gains would be directly visible to outside observers, a researcher investigating the effects of perceived atrocities in empirical instances of conflict would have difficulty distinguishing among these theoretical possibilities: polarization, or loop gain, or perhaps both.

DISCUSSION

As I said at the outset, Collins (2012) has accomplished much of significance with his theory of conflict escalation and de-escalation. His theory builds on substantial empirical support, challenges sociologists to take the dynamics of time seriously in their theoretical models, and points to a fruitful approach for dealing with time dynamics by invoking the systems-theory

concept of feedback loops. Using a variation of dynamic-systems theory, I have presented an alternative to Collins's model: an agent-based computational model capable of actually charting the time dynamics of episodes of conflict escalation, one of the ultimate goals of the modeling project undertaken by Collins (2012:13).

To construct my model of the time dynamics of conflict, I have turned to agent-based computational modeling, because of several advantages for theory building put forward by its proponents: its inherent dynamism, its simultaneous focus on micro interactions and emergent macro patterns, its greater realism at the micro level, and its enhanced degree of scientific rigorousness. However, many sociological theorists still find this more mathematical form of theoretical modeling less intuitively appealing than the conventional approach of verbal theorizing supplemented by schematic diagrams used by Collins. By comparing these two theory-building approaches in application to the same empirical question, we can evaluate whether the purported advantages of computational modeling have made it a useful approach in this case.

To compare the implications of the two models, I will look first at similarities in their theoretical implications and then consider the ways that the computational model either contradicts or goes beyond the conclusions Collins drew from his model. I will also present evidence from two studies of conflict escalation episodes that can throw light on the theoretical differences between the two models. This comparison will demonstrate that my computational model displays many of the advantages attributed to the computational approach.

The two models considered in this article generally agree in several of their theoretical implications, although the computational model's greater precision in definitions of variables and greater micro complexity make the areas of agreement less than complete. Both models imply

that conflicts tend to escalate in a pattern of positive feedback, so that increases of aggressiveness by one party to the conflict are answered by similar increases from the other side. However, the computational model shows how this interactional pattern of positive feedback results from negative-feedback processes at the micro level, rather than from a self-reinforcing feedback loop *per se*.

Both models also agree that increased polarization leads to escalation of conflicts. However, in Collins's model the positive feedback loop linking the two variables implies that the reverse is also true: conflict escalation causes an increase in polarization. The computational model, by contrast, suggests that, although an increase in polarization will produce more rapid escalation, conflicts may continue to escalate rapidly regardless of whether polarization also increases simultaneously.

In Collins's model, the *atrocities* variable intervenes in the positive feedback loop between *conflict* and *ideological polarization*, indicating that escalating conflicts make atrocities more likely, and the atrocities will then provoke increases in polarization. While the computational model does not make any explicit prediction about the link between conflict escalation and atrocities, my simulations are consistent with the implication of Collins's model that, if atrocities increase polarization, conflict will also escalate. Alternatively, the computational model suggests that atrocities may increase the rate of conflict escalation more directly by increasing the loop gain—in other words, the effort devoted to conflict—of the victimized side. Despite these differences in detail, the two models are broadly in agreement about relationships between conflict escalation, polarization, and atrocities.

One important implication of the computational model, however, departs in an unexpected direction from the model Collins presented. The computational model implies that in

the absence of limits on the combatants' output—limits to resources, energy, or attention—conflicts tend to move toward stalemates in which the reciprocal escalations and counter-escalations have little effect on the contested variable or, in other words, the stakes of the conflict. One side might get its way more nearly than the other, but as long as their reference conditions for the contested variable diverge, the two sides will be caught in an unsatisfactory compromise that fails to halt the further escalation of the conflict.

While this image of escalating conflicts creating stalemate may seem counter-intuitive at first, one can easily bring to mind empirical examples of stalemated but still-growing conflicts. From the entrenched armies of World War I, to the Cold-War nuclear buildup, to gridlock in the United States Congress, to centuries-old ethnic and religious conflicts, such stalemated conflicts abound. Sociological accounts of episodes of escalating but stalemated conflict are much less plentiful, however, perhaps because our theoretical perspectives have not prompted us to be on the lookout for them. One notable exception is Beth Roy's *Some Trouble with Cows* (1994), a richly detailed ethnographic account of the escalation and dénouement of a communal "riot" in 1954 in an obscure village in East Pakistan (now Bangladesh). Roy's research examines the time dynamics of an escalating conflict and shows escalation and stalemate in tandem.

Roy interviewed both Hindu and Muslim participants in a conflict provoked when a Muslim farmer's cow got loose and started eating the crops in the field of his Hindu neighbor. According to Roy's informants—both Muslim and Hindu—in the next three days this quarrel between two farmers escalated into a massive confrontation between Hindus and Muslims, involving many thousands of men, some from neighboring communities, armed with primitive weapons like swords and scythes. The episode ended abruptly when the police, who had been

summoned to restore order, shot into the crowd and killed at least two of the combatants (Roy 1994; see also McClelland 2006).

According to Roy, the squabble between two neighboring farmers quickly escalated into a conflict that pitted groups of men against each other. On the first night of the episode, the Muslim farmer whose cow ate the neighbor's crops was confronted by not only the farmer whose field had been violated but also by several of his Hindu relatives, who seized the offending cow (Roy 1994:48-53). Having managed to free his cow and run home, the Muslim farmer recruited several of his own kinsmen and neighbors as backup the next day, when he again tethered his cows near the Hindu neighbor's fields. A series of fights and skirmishes broke out between two families, with others from the village joining in on both sides. Each side snatched cows belonging to the other, and the Muslim farmer was hacked on the arm by a Hindu scythe (pp. 53-56).

On the second day, the conflict grew even more serious, as messengers on horseback carrying mikes and loudspeakers recruited bands of Muslim and Hindu men from neighboring villages to join the fray (Roy 1994:56-57, 67-68). By the end of that day, large and roughly equal numbers of men on both sides had converged on the village. The quarrel over cows and crops had turned into a struggle between local Hindus—who had been dominant until the 1947 partition of Pakistan from India—and Muslims, who felt empowered in the newly created Muslim state—and it was this higher-stakes conflict that drew in men from the neighboring villages (p. 65). As the men confronted each other, and aggressive actions by one side were parried by counter-thrusts from the other, a stalemate was developing.

The conflict came to a climax on the third day, with the climactic events presenting a striking picture of simultaneous escalation and stalemate. Participants later described it as a

“riot,” but, as Roy notes, by American standards it was a remarkably orderly riot (1994:72).

Thousands of armed men from each side staked out positions in a field and sat down in long lines opposite each other in an “aggressive face-off” (p. 80). One informant recalled, “It was as if there were a canal, with two parties sitting on the two sides of it. A group sat on that side of the canal, a group sat on this side of the canal, and the space between them was empty” (p. 79).

Despite rock-throwing, some “chasing and counter-chasing” (p. 81), and a few injuries, nobody was killed, and the “main form of the battle was sedentary” (p. 80), until the police arrived and dispersed the rioters by shooting indiscriminately at both sides. Neither side had gained the upper hand, but further hostilities were postponed. This vivid image of virtual immobility in the midst of intense conflict is exactly the outcome implied by the computational model for a struggle between adversaries who are reasonably well matched.

The computational model, by highlighting the way that escalating conflicts can produce static outcomes, has yielded an important theoretical implication not found in Collins’s model.⁸ The example of the communal riot shows how struggles between opponents can paradoxically stabilize the stakes of the conflict, a dynamic that clearly applies in many other empirical cases. This confirmation of a prediction from the computational model has underlined two of the expected advantages of computational modeling. First, the dynamic capabilities of the computational model can reveal what changes over time and also what *doesn’t* change, the stalemated situation on the ground. Second, the micro-level focus of the computational model has allowed the analyst to examine the interaction from the participants’ points of view, bringing the things combatants care most about, the stakes of the conflict, into focus. By contrast, conventional modeling takes the view of an outside observer looking only at opponents’

collective actions and reactions and thus missing the motivations fueling the conflict and the situational results of the struggle from the combatants' own perspective.⁹

Turning next to the implications of the two models about solidarity and conflict, we see a sharp divergence in their predictions. Collins argues that conflict and solidarity “cause each other to rise” (2012:2) and his diagrams place the two in a positive feedback loop. His definition of solidarity, however, is ambiguous: Does he define solidarity instrumentally as unity of purpose and action, or emotionally as feelings of unity with the group? His discussion of the variable suggests that he means to include both senses of the term. But what if the instrumental and emotional senses of solidarity are analytically separable and affect the dynamics of conflict differently?

My computational model defines solidarity instrumentally as unity of purpose—agreement on reference conditions for the contested variable—because the model in its rudimentary form doesn't address emotions. Given this limitation, however, my simulations demonstrate that, contrary to Collins's assertion, an increase in instrumental solidarity does not inevitably increase conflict escalation. It depends, instead, on the degree of polarization. If agents on one side of a conflict resolve a disagreement about goals (reference conditions) by uniting at a reference point less polarized than the average of their original opinions, the rate of conflict escalation will decrease, not increase. Thus, the simulations imply that a group unified around limited goals may gain in solidarity, while becoming less aggressive as a fighting force. Do solidarity and conflict always go hand in hand, as Collins argues, or do the effects of solidarity vary with the degree of polarization, as the computational model suggests?

Kathleen Blee's recent study of grassroots activist groups (2012) provides some relevant data on the relationship of solidarity to conflict. Studying dozens of newly forming activist

groups in Pennsylvania, Blee observed meetings, interviewed members, and examined documents. Her study focuses on “sequences of action and interpretation” as groups grew or, more often, withered away (p. 14). While these groups did not ordinarily engage in violent conflict, their agendas of social change inevitably led to friction with defenders of the status quo. One “case comparison” of the trajectories of two similar groups, Planet Protection Society (PPS) and Animal Liberation League (ALL) as Blee calls them, speaks directly to the relationship of solidarity to conflict (pp. 128-131).

Emotional solidarity was high in PPS, and Blee describes this group of “young, gender-diverse, and mostly university students” as “inclusive, tightly knit, and mutually supportive . . . a ‘fun’ group” (2012:128). Their meetings even included cheerleading and group hugs (p. 129). But the group settled upon a timid agenda of action, shying away from any confrontation with authorities, and they accomplished little in the way of radical change (p. 130). Thus, the PPS displayed strong solidarity in both senses, instrumental and emotional, but the group avoided any real conflict.

The members of ALL came from similar backgrounds, but their emotional solidarity was low. Blee describes the “rigid, sober, and tense emotional style” of their meetings (2012:128) and reports that ALL had a hard time attracting and retaining new members (p. 130, 72-73). After lackluster campaigns against fur clothes, meat eating, and wool production, a campaign to get foie gras off restaurant menus helped members to redefine themselves as part of a national movement for animal rights. They began to adopt more aggressive tactics (pp. 42-46), seeking to become a “big, major annoyance” in order to force restaurants to stop serving goose-liver paté (p. 47). Conflict with restaurant owners escalated, eventually reaching the point that the state legislature passed an “eco-terrorism” bill, which ALL’s members saw as an attempt to curtail

their protests (pp. 41-42). Even though ALL meetings continued to have only “minimal emotional content” according to Blee (p. 131), ALL’s members, by gaining instrumental solidarity around a polarized set of reference values, had embroiled themselves in a rapidly escalating conflict.

These contrasting examples of activist groups from Blee’s study make clear the inadequacies and oversimplifications in Collins’s account of the relationship between solidarity and conflict, which lies at the heart of his model of conflict escalation. The computational model required a more precisely nuanced definition of solidarity than the traditional concept offered by Collins, and the empirical examples of activist groups confirmed the computational model’s implication that the effects of solidarity depend on the degree of polarization in a group’s reference values for a contested variable. Clearly, Collins’s model would have better represented the effects of solidarity on conflict, if the solidarity variable—preferably split into two or more variables—were placed in the same feedback loop as the polarization variable, instead of a separate loop as Collins depicted it.

The advantages of the computational model emerge clearly from this comparison. Not only does the computational model provide a more dynamic picture of changes and constancies over time by generating the macro processes from the micro, but the computational model also provides more realistic empirical predictions, despite its still-rudimentary form. Thus, three of the four main advantages proclaimed by advocates of computational modeling have been demonstrated in this comparison. Whether this computational model of conflict escalation will also prove more scientifically illuminating than the conventional alternative awaits further research. The PCT model’s success in revealing unexpected insights about conflict escalation,

despite its rudimentary form, suggests its promise as a scientific tool, but fulfillment of that promise will require additional work, both theoretical and empirical.

The PCT model takes a modular form, which allows for more complex models that are more realistic and widely applicable. One possibility is to construct each simulated agent as a multi-level hierarchy of control systems, in line with the neural organization envisioned by perceptual control theory. Constructing agents with multi-level control capabilities would then allow researchers to simulate multidimensional conflicts. The stakes in real-world conflicts are rarely simple, as combatants struggle for the control of many contested variables at once, and with multi-level PCT models researchers could simulate these complicated struggles.

Another possibility for revising and expanding the PCT model is to add self-reorganization features. Perceptual control theory implies that failure to control, which can happen when stalemated conflicts prevent combatants from reaching their goals, is inevitably frustrating, and that such emotional reactions set in motion reorganization processes in the brain, as individuals cast about for alternative ways to get back in control. Multi-level PCT models with these reorganization features have been constructed (Powers 2008), and the application of self-reorganizing models to simulations of conflict could help to reveal how behavioral innovations are related to conflict.

Rigorous testing of more complex simulation models against real-world data must await improvements in data collection. Ethnographic accounts, like those of Roy (1994) and Blee (2012), give us the outlines of conflict escalation sequences, but lack the strict time accounting necessary for testing a computational model. Tests of dynamic models require time-specific data that can be quantified, including data on combatants' own goals for the conflict. Thus, to see

more of the scientific advantages of computational modeling, researchers will need to develop better data-collection techniques, perhaps based on innovations in digital technology.¹⁰

The relative abstraction and flexibility of the PCT model have opened the door to new theoretical developments in the study of real-world conflicts, not just the violent conflicts described by Collins, but also political struggles waged by social-movement organizations like those studied by Blee, as well as many other kinds of social conflict. In short, this agent-based computational model, which re-conceptualizes conflict as a struggle for control, goes beyond conventional theoretical approaches and offers new possibilities of scientific advancement in the sociological study of the time dynamics of conflict.

NOTES

1. Pierre Bourdieu is one notable exception to the generalization that sociologists have neglected the time dimension in their analyses. In his field theory, Bourdieu uses the analogy of a football or tennis player to describe the way that an individual interacting in a social field must adjust to moment-by-moment changes in his environment (1990:66-67). Despite his theoretical emphasis on the importance of time, Bourdieu's empirical work (*e.g.*, 1984), like most sociological analyses, does not focus on close analysis of the time dynamics of specific incidents.
2. Systems thinking emerged in the first half of the 20th century, as American engineers worked to stabilize the newly developing technology of long-distance telephone service, and investigators supported by the armed forces sought to improve systems for aiming guns in naval warfare and shooting down enemy aircraft (see Mindel 2002).
Mathematician Norbert Wiener, who was involved in war research himself, helped to lay the mathematical foundation for the understanding of negative feedback control systems

and gave the name “cybernetics” to this emerging field (Wiener 1948). Cybernetics enjoyed a burst of popularity in the middle of the 20th century, influencing such prominent sociologists as Talcott Parsons, but then faded for a variety of reasons, including mistakes made by Wiener and other advocates in presenting the approach, the emergence of a competing perspectives like artificial intelligence research based on analogies between the brain and the newly developing digital computer, and the drying up federal funding for cybernetics research in the Cold War after Russian scientists embraced Wiener’s work enthusiastically (see Conway and Siegelman 2005, Hayles 1999).

3. Fletcher et al. (2011) performed simulations with a computational model based on a theory of battle dynamics proposed by Collins (2010), but their approach was substantially different from the agent-based modeling that I use in this paper. Another systems-modeling approach to the analysis of conflict makes use of “dynamical social psychology” (*e.g.*, Liebovitch, Vallacher and Michaels 2010; Nowak *et al.* 2010; see also Vallacher and Nowak 1997).
4. Macy and Willer ([2002], 2010) also commented on the seeming reluctance of sociologists in comparison to other social scientists, such as economists and political scientists, to embrace the agent-based modeling approach.
5. With all its positive feedback loops, the schematic model of conflict escalation offered by Collins might also be regarded as having a dynamic element, but, in contrast to agent-based models, the feedback loops in Collins’s model conform to an older tradition of theory construction called “reverse causal effects” (Stinchcombe 1968, see also Turner 1993), which represents one variable as affecting another and then the second as having

reciprocal effects on the first at some unspecified subsequent time. *Time*, as such, is not a parameter of these models.

6. An *Excel* file containing the model used for the simulations in this article is available on request from the author.
7. Both the loop gains and the slowing factor for these simulations have been selected arbitrarily, but when loop gain is sufficiently high (greater than 10) control systems perform in a predictably stable manner over a wide range of loop-gain values (Powers 2008). Within this range of plausible values, running these simulations with other combinations of loop-gain and slowing-factor parameters does not change the substantive results.
8. Collins describes stalemate as one possible outcome for an escalating conflict, but his model does not directly imply this outcome. Instead, he remarks, “How long [stalemate] goes on and why has not been carefully studied” (2012:11).
9. Although Collins offers a micro-level model in conjunction with his model of conflict escalation (2012:3), neither his micro- nor his macro-level model is mathematically specified, and thus the micro-level model cannot be used to “generate” the macro model (see Epstein 2006).
10. In his article, Collins points to one creative strategy for collecting the necessary kinds of data: he describes how he traced the escalation of an on-line controversy, drawing his data from time-stamped e-mails (2012:6-8). In another creative example of finding data with strict time accounting, Heise (2006) has applied Affect Control Theory, a dynamic model related to PCT, to the analysis of data on diplomatic exchanges between Israel and its Arab neighbors.

REFERENCES

- Bell, Heather C. and Sergio M. Pellis. 2011. "A Cybernetic Perspective on Food Protection in Rats: Simple Rules Can Generate Complex and Adaptable Behaviour." *Animal Behaviour* 82:659-666.
- Blee, Kathleen M. 2012. *Democracy in the Making: How Activist Groups Form*. New York: Oxford University Press.
- Bourbon, W. Thomas, Kimberly E. Copeland, Vick R. Dyer, Wade K. Harman, and Barbara L. Mosley. 1990. "On the Accuracy and Reliability of Predictions by Control-System Theory." *Perceptual and Motor Skills* 71:1331-1338.
- Bourbon, W. Thomas. 1990. "Invitation to the Dance: Explaining the Variance When Control Systems Interact." *American Behavioral Scientist* 34(1):95-105.
- Bourdieu, Pierre. 1984. *Distinction: A Social Critique of the Judgement of Taste*. Translated by Richard Nice. Cambridge, MA: Harvard University Press.
- Bourdieu, Pierre. 1990. *The Logic of Practice*. Stanford, CA: Stanford University Press.
- Burke, Peter J. and Donald C. Reitzes. 1991. "An Identity Theory Approach to Commitment." *Social Psychology Quarterly* 54(3): 239-251.
- Burke, Peter J. and Jan E. Stets. 2009. *Identity Theory*. New York: Oxford University Press.
- Collins, Randall. 1992. *Sociological Insight: An Introduction to Non-Obvious Sociology*. Second Edition. Oxford University Press.
- Collins, Randall. 2004. "Rituals of Solidarity and Security in the Wake of Terrorist Attack." *Sociological Theory* 22:53-87.
- Collins, Randall. 2008. *Violence: A Micro-Sociological Theory*. Princeton, NJ: Princeton University Press.

- Collins, Randall. 2009. "The Micro-Sociology of Violence." *The British Journal of Sociology* 60: 566-576.
- Collins, Randall. 2010. "A Dynamic Theory of Battle Victory and Defeat." *Cliodynamics: The Journal of Theoretical and Mathematical History* 1:3–25 (<http://escholarship.org/uc/item/5mv6v0r1>).
- Collins, Randall. 2012. "C-Escalation and D-Escalation: A Theory of the Time-Dynamics of Conflict." *American Sociological Review* 77(1):1–20.
- Conte, R., N. Gilbert, G. Bonelli, C. Cioffi-Revilla, G. Deffuant, J. Kertesz, V. Loreto, S. Moat, J.-P. Nadal, A. Sanchez, A. Nowak, A. Flache, M. San Miguel, and D. Helbing. 2012. "Manifesto of Computational Social Science." *The European Physical Journal Special Topics* 214:325–346.
- Conway, Flo and Jim Siegelman. 2005. *Dark Hero of the Information Age: In Search of Norbert Wiener, the Father of Cybernetics*. New York: Basic Books.
- Coser, Lewis. 1956. *The Functions of Social Conflict*. Glencoe, IL: Free Press.
- Durkheim, Emile. [1893] 1933. *The Division of Labor in Society*. Translated by George Simpson. New York: Free Press.
- Durkheim, Emile. [1897] 1951. *Suicide: A Study in Sociology*. Translated by John H. Spaulding and George Simpson. New York: Free Press.
- Epstein, Joshua M. 2006. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton University Press.
- Fletcher, Jesse B., Jacob Apkarian, Anthony Roberts, Kirk Lawrence, Christopher Chase-Dunn, and Robert A. Hanneman. 2011. "War Games: Simulating Collins' Theory of Battle

Victory.” *Cliodynamics: The Journal of Theoretical and Mathematical History* 2: 252-275
(<http://escholarship.org/uc/item/7hk4279k>).

Gilbert, Nigel. 2008. *Agent-based Models*. Thousand Oaks, CA: Sage.

Hayles, N. Katherine. 1999. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics*. Chicago: University of Chicago Press.

Hedström, Peter and Petri Ylikoski. 2010. “Causal Mechanisms in the Social Sciences.” *Annual Review of Sociology* 36:49–67.

Heise, David R. 2006. “Sentiment Formation in Social Interaction.” Pp. 188-211 in *Purpose, Meaning, and Action: Control Systems Theories in Sociology*, edited by Kent A. McClelland and Thomas J. Fararo. New York: Palgrave Macmillan.

Heise, David R. 2007. *Expressive Order: Confirming Sentiments in Social Action*. New York: Springer.

Liebovitch, Larry S., Robin R. Vallacher and Jay Michaels. 2010. “Dynamics of Cooperation-Competition Interaction Models.” *Peace and Conflict* 16:175-188.

Macy, Michael W. and Robert Willer. [2002] 2010. “From Factors to Actors: Computational Sociology and Agent-based Modeling.” Pp. 29-53 in *Computational Social Science*, Vol. II, Nigel Gilbert, ed. Los Angeles, CA: Sage.

Mansell, Warren and Timothy A. Carey. 2009. “A Century of Psychology And Psychotherapy: Is an Understanding of ‘Control’ the Missing Link Between Theory, Research, and Practice?” *Psychology and Psychotherapy: Theory, Research and Practice* 823:337–353.

Mansell, Warren, Timothy A. Carey and Sara J. Tai. 2012. *A Transdiagnostic Approach to CBT using Method of Levels Therapy*. New York: Routledge.

- Marchioni, Caterina and Petri Ylikoski. 2013. "Generative Explanation and Individualism in Agent-Based Simulation." *Philosophy of the Social Sciences* 43(3):323-340.
- Marken, Richard S. 1980. "The Cause of Control Movements in a Tracking Task." *Perceptual and Motor Skills* 51: 755-758.
- Marken, Richard S. 1986. "Perceptual Organization of Behavior: A Hierarchical Control Model of Coordinated Action." *Journal of Experimental Psychology—Human Perception and Performance* 12: 267-276.
- Marken, Richard S. 1988. "The Nature of Behavior: Control as Fact and Theory. Behavioral Science 33, 196-206.
- McClelland, Kent A. 2006. "Understanding Collective Control Processes." Pp. 31-56 in *Purpose, Meaning, and Action: Control Systems Theories in Sociology*, edited by Kent A. McClelland and Thomas J. Fararo. New York: Palgrave Macmillan.
- McClelland, Kent. 1994. "Perceptual Control and Social Power." *Sociological Perspectives* 37:461-496.
- McClelland, Kent. 2004. "The Collective Control of Perception: Constructing Order from Conflict." *International Journal of Human-Computer Studies* 60:65-99.
- McClelland, Kent A. and Thomas J. Fararo, eds. 2006. *Purpose, Meaning, and Action: Control Systems Theories in Sociology*. New York: Palgrave Macmillan.
- McPhail, Clark and David Schweingruber. 2006. "The Closed-Loop Negative Feedback Model: Some Quantitative Evidence." Appendix 1, pp. 317-320 in *Purpose, Meaning, and Action: Control Systems Theories in Sociology*, edited by Kent A. McClelland and Thomas J. Fararo. New York: Palgrave Macmillan.

- McPhail, Clark and Ronald T. Wohlstein. 1986. "Collective Locomotion as Collective Behavior." *American Sociological Review* 51: 447-463.
- McPhail, Clark. 1991. *The Myth of the Madding Crowd*. New York: Aldine de Gruyter.
- Meadows, Donella H. et al. 1972. *The Limits to Growth: A Report for the Club of Rome's Project on the Predicament of Mankind*. New York : Universe Books.
- Miller, John H. and Scott E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, NJ: Princeton University Press.
- Mindell, David A. 2002. *Between Human and Machine: Feedback, Control, and Computing Before Cybernetics*. Baltimore, MD: Johns Hopkins University Press.
- Nowak, Andrzej, Morton Deutsch, Wieslaw Bartkowski and Sorin Solomon. 2010. "From Crude Law to Civil Relations: The Dynamics and Potential Resolution of Intractable Conflict." *Peace and Conflict* 16:189-209.
- Pellis, Sergio M. and Heather C. Bell. 2011. "Closing the Circle Between Perceptions and Behavior: A Cybernetic View of Behavior and Its Consequences For Studying Motivation and Development." *Developmental Cognitive Neuroscience* 1:404-413.
- Pellis, Sergio M., David Gray and William H. Cade. 2009. "The Judder of the Cricket: The Variance Underlying the Invariance in Behavior." *International Journal of Comparative Psychology* 22:188-205.
- Powers, William T. [1973] 2005. *Behavior: The Control of Perception. Second Edition*. New Canaan, CT: Benchmark Publications.
- Powers, William T. 1978. "Quantitative Analysis of Purposive Systems: Some Spadework at the Foundations of Scientific Psychology." *Psychological Review* 85:417-435.

- Powers, William T. 2008. *Living Control Systems III: The Fact of Control*. Bloomfield, NJ: Benchmark Publications.
- Raub, Werner, Vincent Buskens and Marcel A. L. M. van Assen. 2011. "Micro-Macro Links and Microfoundations in Sociology." *The Journal of Mathematical Sociology* 35:1-25. DOI: 10.1080/0022250X.2010.532263
- Robinson, Dawn T. 2007. "Control Theories in Sociology." *Annual Review of Sociology* 33:157–174.
- Roy, Beth. 1994. *Some Trouble with Cows: Making Sense of Social Conflict*. Berkeley: University of California Press.
- Sewell, William H. 2005. *Logics of History: Social Theory and Social Transformation*. Chicago: University of Chicago Press.
- Simmel, Georg. [1908] 1955. "Der Streit." Chapter 4 in *Soziologie*. Translated in *Conflict and the Web of Group Affiliations*, edited by R. Bendix and K. H. Wolff. Glencoe, IL: Free Press.
- Stinchcombe, Arthur L. 1968. *Constructing Social Theories*. New York: Harcourt, Brace & World.
- Turner, Jonathan H. 1993. "A 'Pure Theory' Research Program on Microdynamics." Pp. 104-135 in *Theoretical Research Programs: Studies in the Growth of Theory*. Joseph Berger and Morris Zelditch, Jr., eds. Stanford, CA: Stanford University Press.
- Vallacher, Robin R., and Andrzej Nowak. 1997. "The Emergence of Dynamical Social Psychology." *Psychological Inquiry* 8:73-99.
- Wiener, Norbert. 1948. *Cybernetics: Control and Communication in the Animal and the Machine*. New York: Wiley.

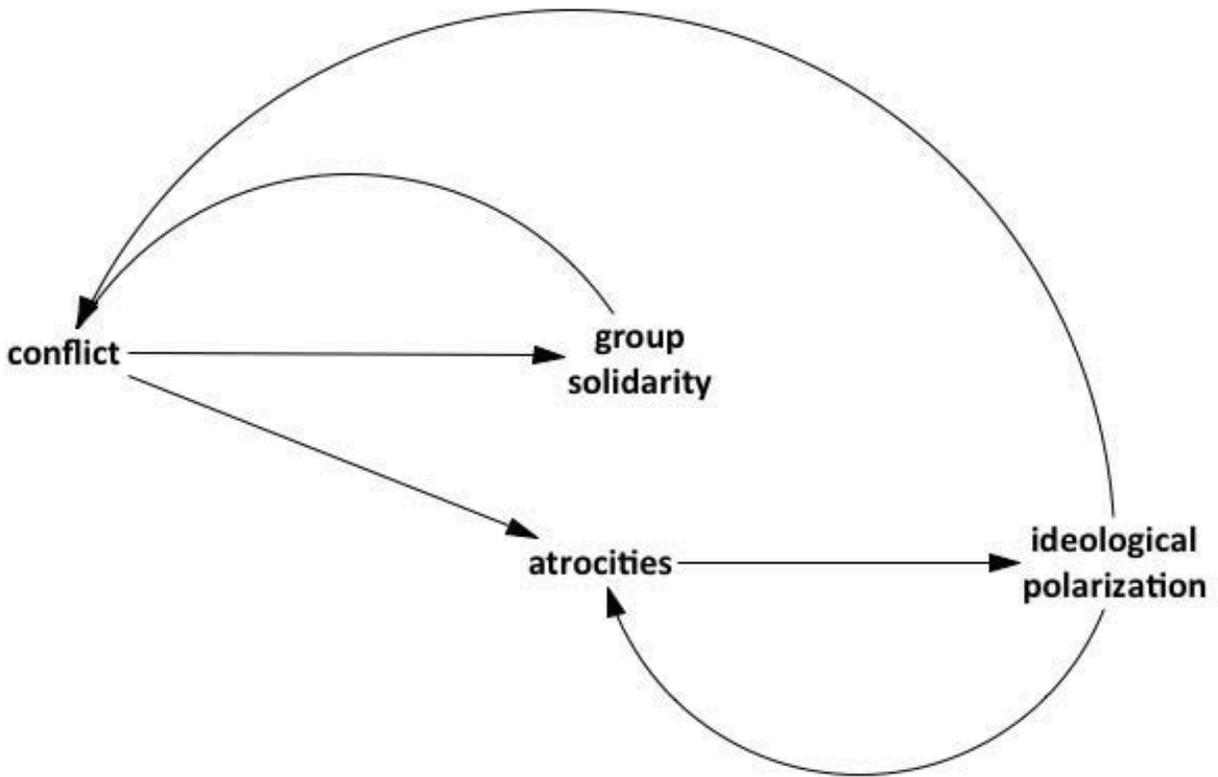


Figure 1. Positive Feedback Loops in Collins’s Theory of Conflict Escalation (Reproduction of Collins [2012] Figure 3, p. 4).

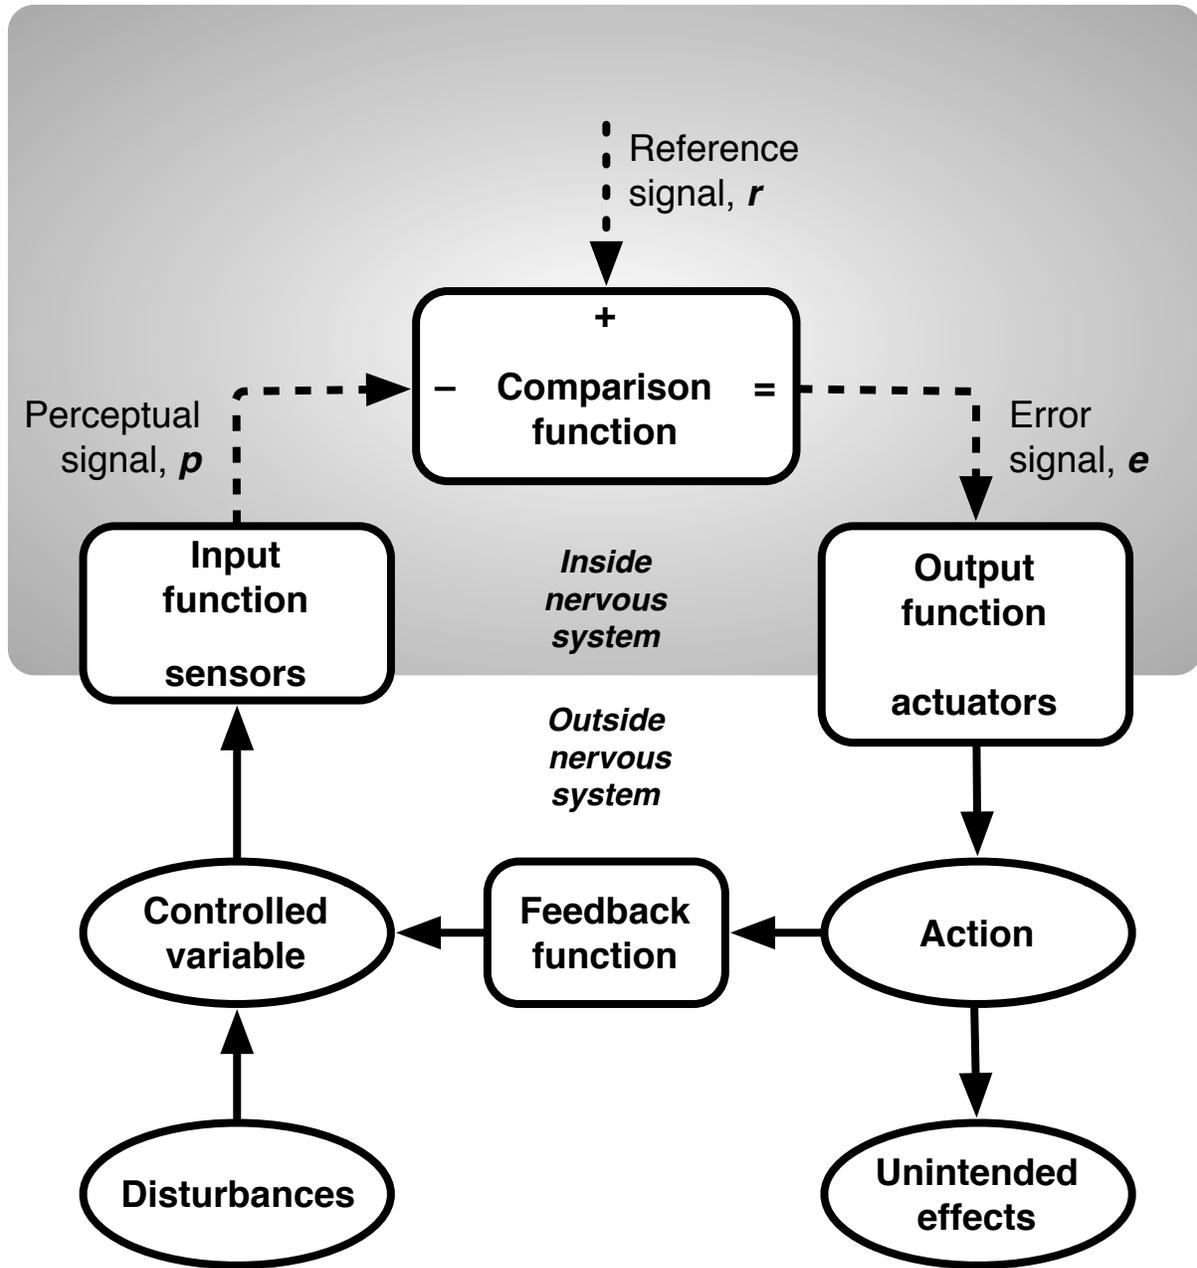


Figure 2. The Perceptual Control Theory Model of a Negative Feedback Loop

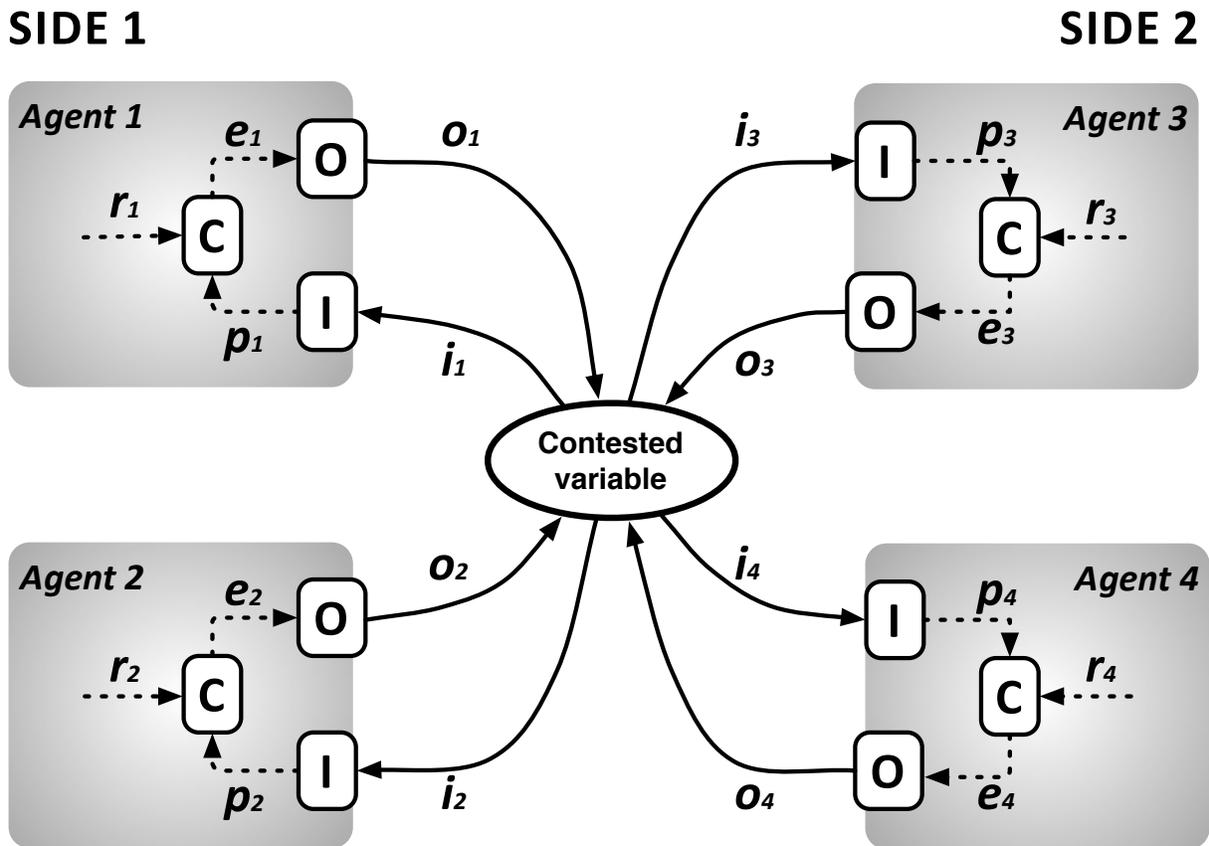


Figure 3. The Simulation Model Used in this Study

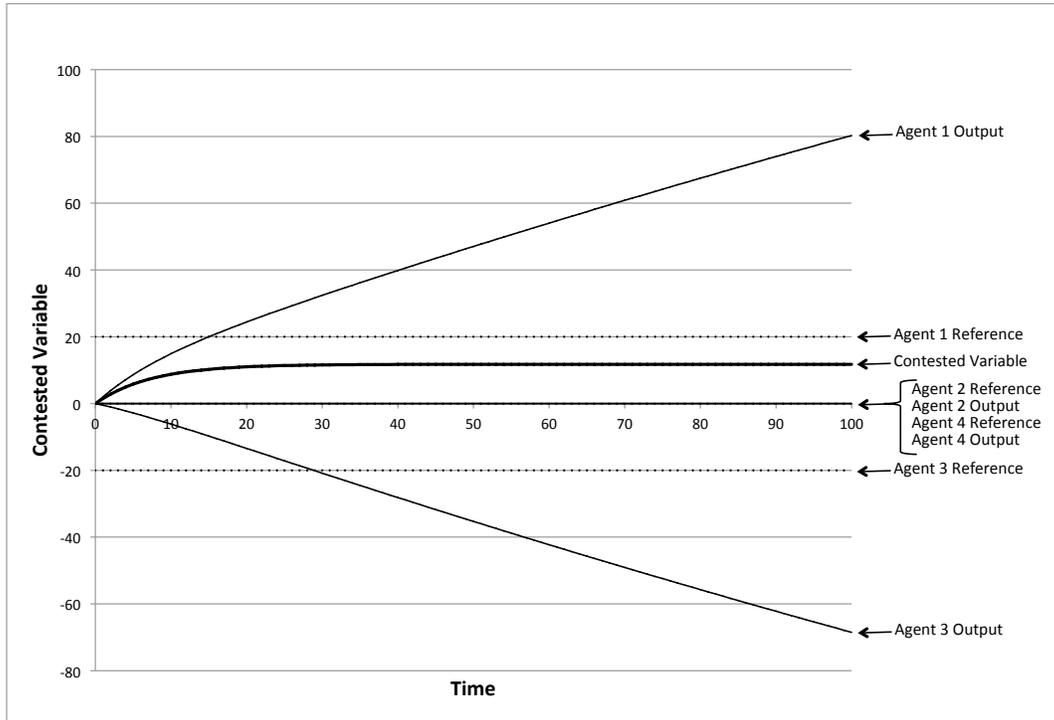


Figure 4. Simulation of Two-Party Conflict

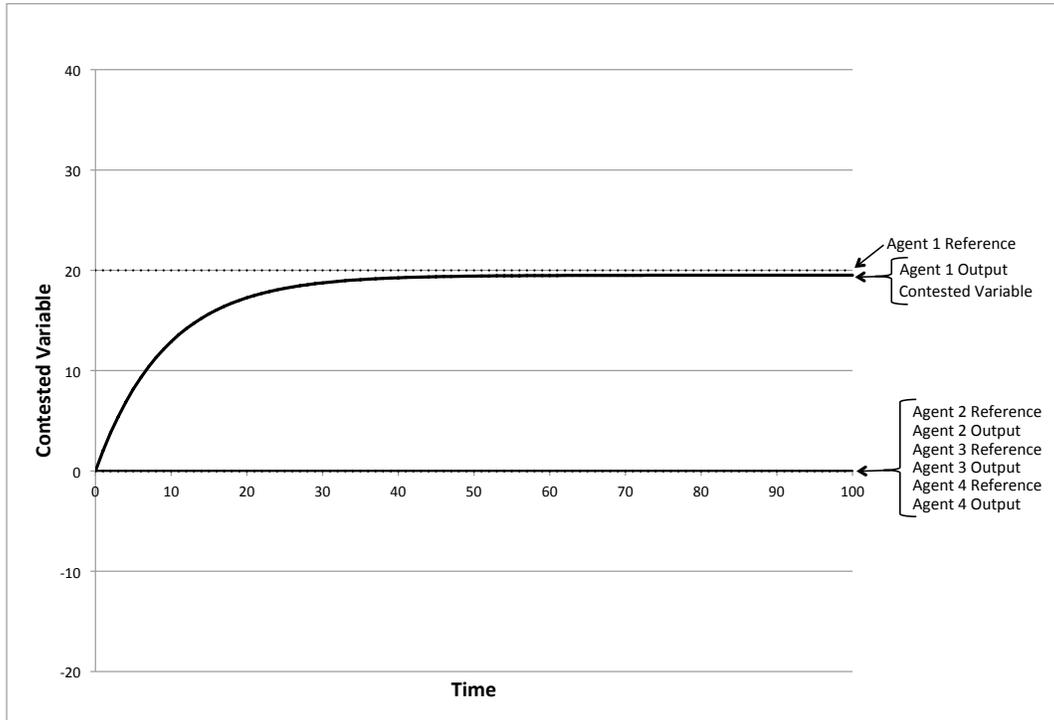


Figure 5A. Control by a Single Agent with Loop Gain = 40

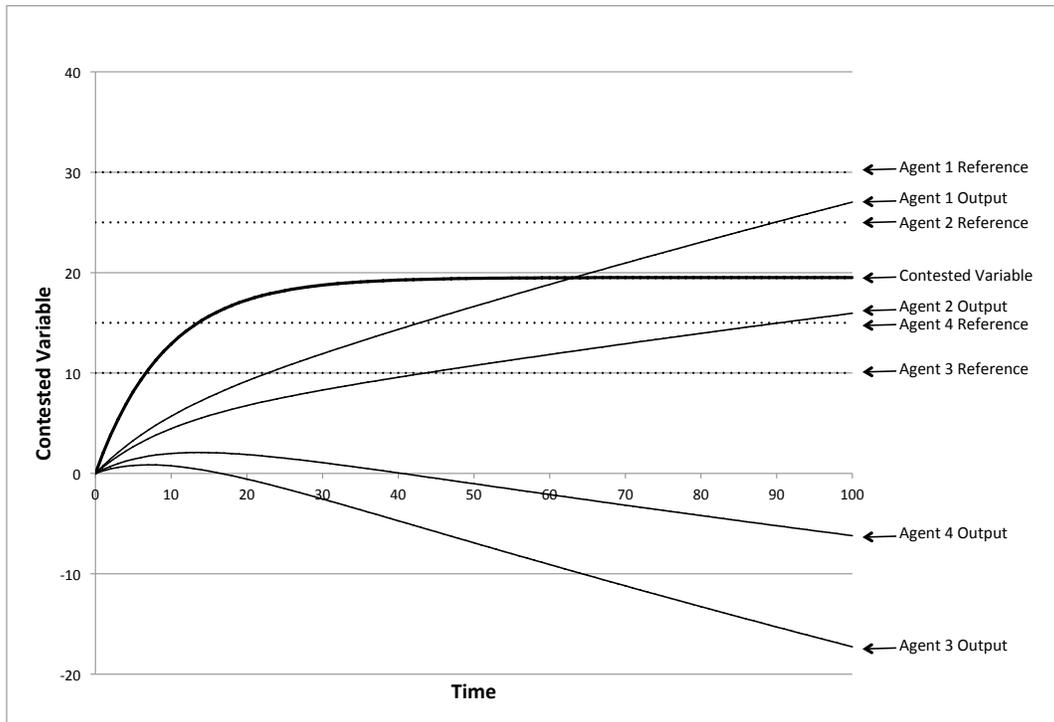


Figure 5B. Collective Control by Four Agents, Each with Loop Gain = 10

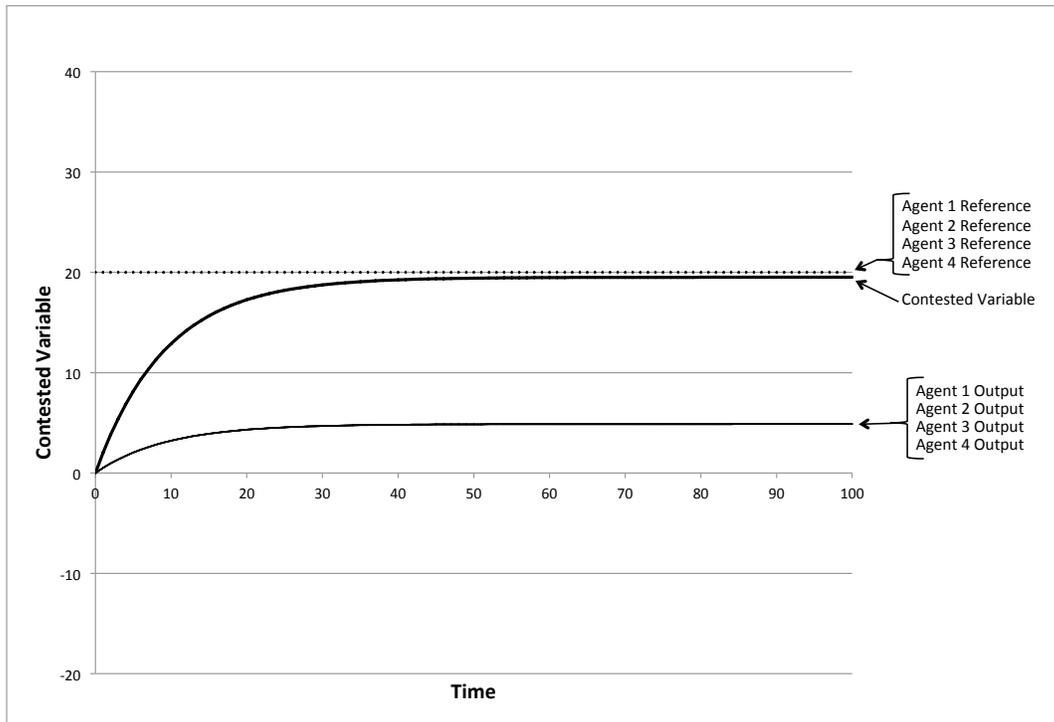


Figure 5C. Collective Control by Four Agents with Identical Reference Values and Loop Gain =10

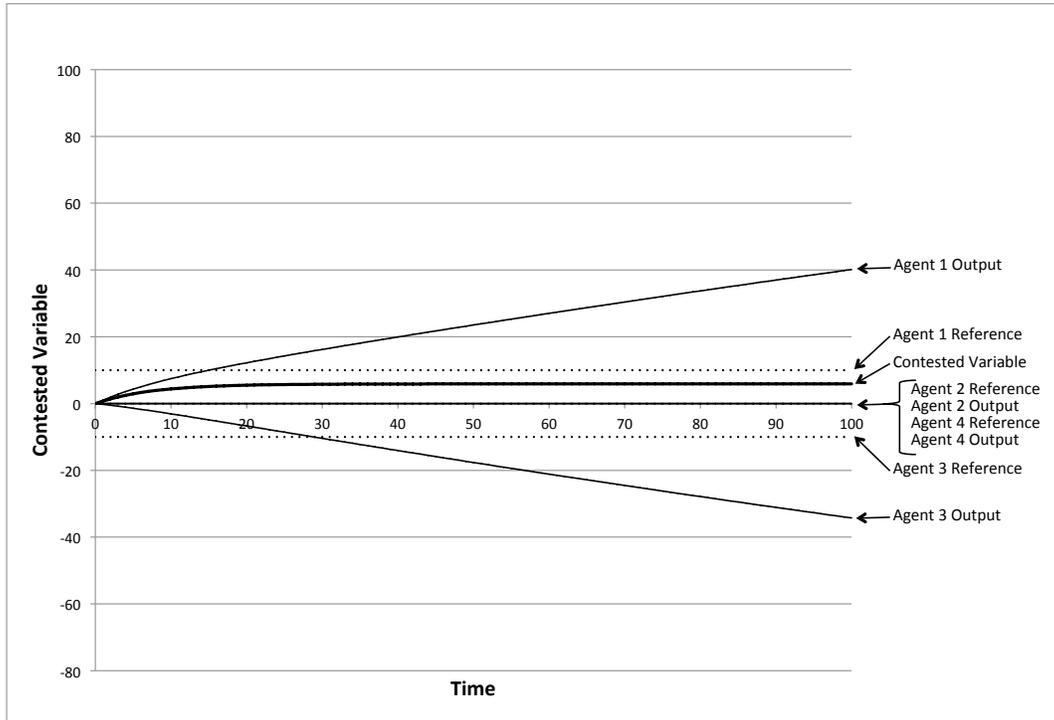


Figure 6. Simulation of Two-Party Conflict with Reduced Polarization

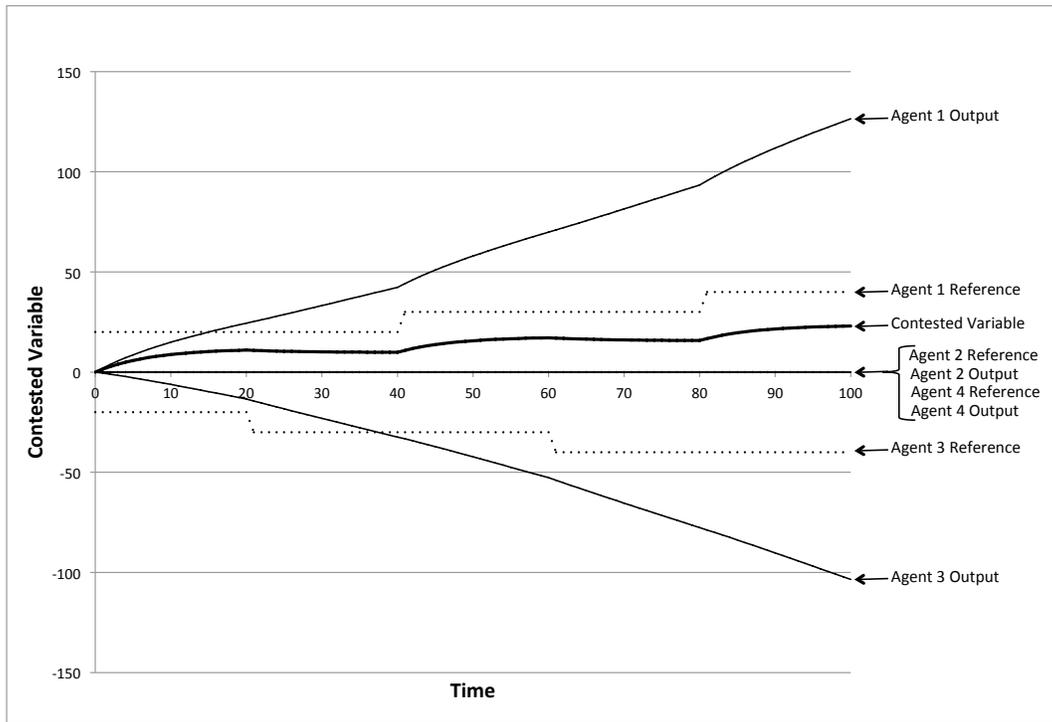


Figure 7. Simulation with Atrocities Represented as Increases in Polarization