Generating Panic Within Populations[°]

Carter Butts[•]

Department of Social and Decision Sciences

and

Center for the Computational Analysis of Social and Organizational Systems

Carnegie Mellon University

Abstract

Previous research on the problem of belief panics - episodes in which numerous actors develop highly divergent beliefs for a brief period in the absence of direct evidence - has demonstrated the plausibility of belief feedback mechanisms as an explanation for panic. Building on this work, a model is here developed which allows for the emergence and dissolution of panic phenomena within structured populations of individual actors. The behavior of this model is then analyzed using a virtual experiment in order to identify the primary determinants of the rate of panic occurrence. Assumptions regarding saliency and communication are shown to be crucial aspects of the panic model as well as predictors of panic rates, along with network density and the rate at which external signals are introduced. Network clustering, while examined, is not found to be related to the panic rate.

[°] The author would like to thank Kathleen Carley and Thomas Fararo for their input and encouragement.
[•] This work is based on research supported under a National Science Foundation Graduate Fellowship, and was supported in part by the Center for the Computational Analysis of Social and Organizational Systems and the Institute for Complex Engineered Systems at Carnegie Mellon University.

1. Introduction

The notion of panic as a sociologically interesting form of collective behavior is not a new one. Indeed, since Charles Mackay's 1841 study of *Extraordinary Popular Delusions and the Madness of Crowds* (at least) there has been an awareness of the volatility of popular opinion, and of the potential for social actors to fall prey to a variety of panic phenomena. While the term "panic" (in its social sense) has been applied to behaviors including breakdown of social order in emergency situations, speculative bubbles in financial markets (Orléan, 1989), (Lux, 1995), moral panics (Goode and Ben-Yehuda, 1994), and deviance production mechanisms (Erikson, 1966), we are here interested in a particular class of social phenomena which are characterized by the rapid emergence and subsequent dissolution[1] of one or more "islands" of minority belief[2] within a "sea" of majority opinion. Sudden, localized outbreaks of concern over supposed acts of witchcraft (Erikson, 1966) or Satanism (Victor, 1993), (Hicks, 1991), would satisfy this definition of panic; so too might rumors of downsizing and acquisition in poorly performing firms, or (exaggerated) fears regarding disease outbreaks.

One example of a series of panic events which we might wish to address can be found in Jeffery Victor's study of satanic cult panics within the United States and Britain during the late 1980's and early 1990s (Victor, 1993), (Victor, 1995). During these events, men and women in a number of geographically disparate locations were confronted by terrifying stories of mutilation and murder, stories which were purported to be both factually accurate and to represent a real and immanent danger to community members (Victor, 1993). As appears to be typical of such circumstances[3], few if any of the panic victims were in a position to verify the veracity of these tales firsthand; instead, they turned to other community members for information. These alters - friends, relatives, co-workers, and the like - passed along information which they had acquired

[1] Although, as we shall see, this dissolution need not be total.
[2] Here, we shall focus on "belief" (in the informational sense) rather than "attitude," although the two are not unrelated (Anderson, 1971).
[3] And, perhaps, much of human knowledge in general.

via similar means, leading, ultimately, to a high level of redundancy which reinforced belief in the reality of the threat (Victor, 1993). Although subsequent investigation by law enforcement authorities revealed no evidence of dangerous cult activity (Victor, 1995), convictions during the panic episodes were strong enough to prompt parents to keep their children home from school and to provoke demands for police protection. Such intense, materially unsubstantiated beliefs are emblematic of a more general process of reality construction (Mehan and Wood, 1975), and hence require explicit sociological attention.

Following Victor's account, Butts (1997) presented a formal model of panic development which demonstrated a process in which signaling and inference behaviors among socially embedded actors lead to rapid polarization of belief. This model belongs to a more general class of social influence models, including those of Krackhardt (1996), Carley (1990a, 1990b), Latané (1981), Friedkin and Cook (1990), and Anderson (1959), which attempt to represent shifts in actors' beliefs and/or opinions via a structurally mediated process of diffusive change. (For foundational work in this area, see Asch (1952), Deutsch and Gerard (1955), and Sherif (1936)[4]). Similarly, it is related to work in the economic literature on herd behavior (see Shiller, 1995) which considers situationally constrained interactions between Bayesian-rational actors. While Butts' model does not assume rational behavior per se, it does represent actors' knowledge in terms of subjective probabilities (Jeffrey, 1983) and assumes that these probabilities are updated via a Bayesian inference process (as in (Banerjee, 1992), (Bikhchandani et al., 1992), (Orléan, 1989)); it may thus be thought of as a model of influence among boundedly rational actors (March and Simon, 1958), although other interpretations are possible.

This Bayesian Belief Feedback Model (BBFM), as specified in (Butts, 1997), was shown to produce a qualitatively reasonable depiction of the reinforcement process described in (Victor, 1993); this, however, is only one of the requirements for understanding belief panics. In particular, our definition specifies that panics must *emerge* from a majority configuration, and that participants' beliefs must ultimately return to some approximation of their original state. The spontaneous creation - and eventual dissolution - of panic phenomena was not within the scope of the original model, and requires a somewhat different approach. The present work attempts to remedy this by extending the original BBFM to allow for the generation of belief panics within populations of socially embedded actors.

2. The Bayesian Belief Feedback Model

Before discussing the extensions to the Bayesian Belief Feedback Model which we shall use to simulate the generation of panics, it behooves us to briefly revisit the its basic construction. As presented in Butts (1997), the BBFM is built around a core learning model. This model is, in essence, a very simple and direct application of Bayes' rule to a hypothetical situation in which an actor observes a categorical "belief" or "disbelief" signal in reference to some possible event[5]. In particular,

given:

$B \equiv$ a belief signal

$T \equiv$ a true event

$\overline{A} \equiv$ the complement (logical NOT) of event A

By Bayes' Rule:

$$[1] \quad p(T|B) = \frac{p(T)p(B|T)}{p(T)p(B|T) + p(\overline{T})p(B|\overline{T})}$$

As indicated in [1], the actor's posterior probability for event T depends on three parameters: the *á priori* probability that the signaler would send a "belief" signal were the event to be true; the *á priori* probability that the signaler would send a "belief" signal were the event to be false; and, finally, the actor's prior belief concerning the probability of T. For simplicity of discussion,

---

[4] These works verified the fact that individuals use each other as sources of information even in cases where direct and unambiguous evidence regarding the nature of reality is available; they also considered other sources of group influence over individual action (i.e., normative sanctions).

[5] In fact, it is not necessary to assume dichotomous - or even categorical - signals. We shall restrict ourselves to this simple case, however.

the first of these parameters will henceforth be referred to as the actor's "credibility rating", or C, and the second will be referred to as his or her "error rating", or E:

$$[2]\ \ C \equiv p(B|T)$$

$$[3]\ \ E \equiv p(B|\overline{T})$$

$$[4]\ \ \theta \equiv p(T)\ \ \text{(prior)}$$

By substituting these into equation [1], we arrive at

$$[5]\ \ p(T|B) = \frac{\theta C}{\theta C + (1-\theta)E} = \frac{\theta C}{\theta(C-E)+E},$$

and the alternative possibility,

$$[6]\ \ p(T|\overline{B}) = \frac{\theta(1-C)}{\theta(1-C)+(1-\theta)(1-E)} = \frac{\theta(1-C)}{\theta(E-C)+(1-E)}.$$

The posterior probabilities of not T, then, are simply

$$[7]\ \ p(\overline{T}|B) = 1 - p(T|B),$$

$$[8]\ \ p(\overline{T}|\overline{B}) = 1 - p(T|\overline{B}).$$

Intuitively, C represents the actor's belief that a signaler would know (and reveal) the event as true if it were true; E, on the other hand, represents his belief that a signaler might err in indicating the event to be true even if it were not. For purposes of the current learning model, these evidentiary weights are considered exogenous parameters; they are set, fixed, and unrelated to θ. While the nature of the inference model is always such that extreme beliefs (i.e., those near 0 or 1) are difficult to change, it is also worth noting that belief states interact nonlinearly with evidentiary weights: in particular, low priors lead to a great dependence on low E values, while high priors lead to sensitivity to low C values.

Superimposed on the actor's belief model is the signaling model, which is likewise quite simple: during an interaction, an actor will signify "belief" in T when p(T) is believed to be greater than .5, and "disbelief" otherwise[6]. This allows us to specify B as follows:

---

[6] While it may be objected that the choice of .5 as the threshold is somewhat arbitrary, it should be noted that A) p(T)=.5 represents a state in which a T is as likely to be true as not and B) the choice of threshold affects only the quantitative, not the qualitative, behaviors of the model. (In any case, since the actor's belief can lie on any point in [0,1], it is generally true that p(θ=0.5)=0.)

$$[9]\quad B(\theta)=\begin{cases} TRUE\ (B) & \theta > 0.5 \\ FALSE\ (\overline{B}) & \theta \le 0.5 \end{cases}$$

From the above, it is fairly trivial to specify the full model for a simple, two-person interaction. Combining [5] and [6] with [9] above, and making time explicit, we obtain the following:

$$[10]\quad \theta_{1,T}=\begin{cases} \dfrac{\theta_{1,T-1}C}{\theta_{1,T-1}(C-E)+E} & \theta_{2,T-1} > 0.5 \\ \dfrac{\theta_{1,T-1}(1-C)}{\theta_{1,T-1}(E-C)+(1-E)} & \theta_{2,T-1} \le 0.5 \end{cases},\quad \theta_{1,T=0}=\theta_{1,0}(prior)$$

$$[11]\quad \theta_{2,T}=\begin{cases} \dfrac{\theta_{2,T-1}C}{\theta_{2,T-1}(C-E)+E} & \theta_{1,T-1} > 0.5 \\ \dfrac{\theta_{1,T-1}(1-C)}{\theta_{2,T-1}(E-C)+(1-E)} & \theta_{1,T-1} \le 0.5 \end{cases},\quad \theta_{2,T=0}=\theta_{2,0}(prior)$$

In (Butts, 1997), it was assumed that the relation C>E was universally true among actors and that each actor had the same values for C and E. In the analysis which follows, the second condition will be relaxed; the first will be retained, however. While it is possible that, in some real-world cases, E *is* greater than C, this seems unlikely to be a common occurrence (recall that C<E implies that actors' signals carry an *inverted* meaning)[7]. The dynamics of the belief system alter radically at C=E; exploration of this region of parameter space is left as a task for future research.

The above formulation describes the special case in which two actors engage each other repeatedly and exclusively; as one might expect, this is somewhat removed from the types of real world scenarios in which we are interested. In Victor's account, participants in the panic event exchanged information with alters who shared relations of one sort or another. Typically, these relations have been conceived of as social networks (Wasserman and Faust, 1994), (Scott, 1991) in which actors are embedded (Granovetter, 1985). As conduits for information and influence (Brass, 1984), (Burt, 1992), (Granovetter, 1982), network connections enable actors to coordinate effort; presumably, this may be the case even when such coordination is neither desired, nor even fully understood by its participants. For this reason, the basic BBFM model was extended to deal with larger sets of actors, as described below.

Following convention in the field of social network analysis (Wasserman and Faust (1994), Scott (1991)), we shall here treat communication networks as sociomatrices (**M**) such that $\mathbf{M}_{i,j}=1$ iff *i* sends signals to *j*, and $\mathbf{M}_{i,j}=0$ otherwise. While these communication matrices may be more or less symmetric, they are not required to be so: this allows for the representation "broadcast" agents, who influence alters from whom they receive no signals[8]. As with (Butts, 1997), we shall here continue to take social structure as exogenous and static; while real structures are unlikely to fulfill either condition, it seems fairly reasonable to make this approximation for the short time spans (hours, days, or perhaps a few weeks) which characterize panic phenomena[9]. Future relaxation of this requirement could permit investigation of the institutional legacy of panic events (Goode and Ben-Yehuda, 1994), but will require the integration of the BBFM with a model of structural change over time (e.g., Carley (1991)).

In a simple, dyadic analysis, the low dimensionality of the BBFM permits a detailed analytic treatment (Butts, 1997). In the more general network case, by contrast, the system of interest is of dimension N (where N is number of actors); similarly, the particular coupling of these N equations is dependent on the structure the adjacency matrix **M**. Given the nature of such a large, nonlinear system, it is sensible to utilize simulation methods to understand its behavior. Here, we will follow the form of the analysis found in (Butts, 1997); using monte carlo methods, we will sample across the space of configurations to attempt to uncover general rules governing the way in which panic is produced under the BBFM.

[7] An exception might be found in the arena of political claimsmaking, in which actors may be motivated to develop beliefs which are in direct opposition to those expressed by certain alters.
[8] Broadcast actors are extremely versatile, being able to represent not only opinion leaders and media sources, but also technologies (Carley, 1995). For our present purposes, however, we shall tend to interpret such actors in terms of human beings.

3. Extending the BBFM

In order to model the emergence and dissipation of panic within a population of actors, it is necessary to extend the basic framework described above to account for two important sets of issues: novel information and saliency.

The original formulation of the BBFM was intended to demonstrate convergence in a closed system; actual social systems, however, are rarely closed with respect to novel information. As persons interact with the world at large, they necessarily make observations and draw inferences. Whether ordered or arbitrary, such new information affects both the observer (by altering his or her beliefs) and the social system as a whole. To represent this flow of novel information into the system, we assume that actors receive signals from the environment as well as from each other. These signals appear at a fixed rate S_r, and are assumed to be "belief" signals with probability S_B (otherwise, they represent "disbelief"). External signals are randomly allocated to actors in a uniform fashion; thus, each actor has an equal probability of receiving any given signal, and actors may receive more than one signal during a given time period[10]. This last feature of the signal allocation process is, as we shall see, an important one: due to the number of incoming signals resulting from incoming communications, it is easy for external information to be drowned out by social influence.

In addition to the problem of new information, it is important to address the issue of saliency in actor communication. The original formulation of the BBFM implicitly assumed that the panic's focal event was salient to all actors in all cases: that is, all actors constantly communicated (in parallel) with all other actors to whom they were connected regarding the controversial issue. While this was not unreasonable for a group in the throes of panic, our current task is to show how panic emerges in the first place. Given that the timing of social

[9] Empirical findings seem to indicate considerable stability in a variety of networks over time periods which are longer than those considered here (e.g., Morgan et al (1997), Wellman et al (1997), Suitor and Keeton (1997)).
[10] Hence, the expected number of signals per actor is S_r/N, and the probability of receiving a given number of signals is binomially distributed.

interactions is known to affect system-level outcomes in a number of cases (e.g., (Macy, 1991), (Granovetter, 1978)), this requires us to consider alternatives to the constant saliency rule.

While many possible solutions to the saliency problem might be proposed, we shall here consider four stylized models, each of which reflects a somewhat different set of theoretical assumptions regarding actor behavior. By comparing results across these sub-models, we may better understand how they affect the evolution of belief; in particular, this should provide evidence as to whether or not certain rules are required in order to produce panic.

Constant Excitation

Under the constant excitation rule, all actors will interact with all other actors to whom they are connected at each time step. These interactions are assumed to take place in parallel, and actors are limited to one such set of interactions per time step. This is the "classic" BBFM as presented in (Butts, 1997), and assumes that the focal event of the panic is salient to all actors at all times.

Signal Excitation

Unlike the constant excitation rule, signal excitation causes actors to communicate with others only when they receive new information. In particular, after receiving a new signal from the environment and updating his or her belief, an actor will immediately send a signal (using the standard rule given in [9]) to all actors to whom he or she has an outgoing tie. Because of this, a given actor may send and/or receive multiple signals (or no signals) during the course of a single time step, and these signals are resolved serially. The signal excitation rule, then, corresponds to the notion that actors regard the focal event of the panic as salient only when receiving an external signal regarding that event; otherwise, they do not trouble their neighbors with their opinions.

Belief Change Excitation

As we have seen, one way of approaching the saliency issue is by assuming that saliency is always present; another way is by assuming that the focal event is salient only when new information from the environment is received. Yet another possibility in this vein is that it is not the presence of new *information* which is important *per se*, but rather the presence of a change in orientation of belief. The belief change excitation rule implements this idea by first applying external information, and then by causing actors to communicate with alters if and only if their belief has been altered with respect to the 0.5 probability threshold (e.g., from 0.3 to 0.6) and if they have not already communicated during the current time step. These communications are both serial and recursive: that is, actors whose beliefs have not crossed the 0.5 threshold prior to the implementation of the rule may be activated by incoming signals from other actors, and will at that point communicate with their own alters. When all actors with changed beliefs have communicated, a stability condition has been achieved. At this point, the rule terminates for the current time step.

Signal Transfer

All of the above rules, following the original BBFM framework, utilize the standard signaling rule given in [9] by which actors' signals are determined by their beliefs. An alternative conception, however, is that actors may not relay their own *opinions* to other actors, but instead may simply repeat what they have heard from others. This raises an important question, however: what mechanism will determine when a message is relayed and when it is not? Clearly, if all signals are transferred, then the presence of cycles within the network would result in a never-ending transfer of any signals received by any actors; this seems a bit unreasonable, given that we do not observe such behaviors in the real world. A more realistic assumption is that actors pass on received signals with some probability, otherwise failing to transfer the information. For our purposes, we shall assume that the probability of transfer is proportional to the actor's subjective probability of the signal's validity; in other words, the signal is passed on with a probability equal to the actor's current belief, θ (or $1-\theta$, if the signal is a "disbelief" signal) multiplied by some probability range and added to a minimum transfer probability. This assumes that actors tend to share what they regard as reliable information with others, and that they likewise tend to block signals which they regard as unlikely to be "correct," but that actors have in all cases a certain minimum and maximum probability of signal transfer.

Given the original BBFM, the presence of external signals, and our four saliency rules, we are left with a series of possible model variants with which to explore the generation of panic within populations. Due to the obvious complexity of the problem, we turn to simulation as a means of studying model behavior across conditions; in particular, we shall seek to implement a virtual experiment to conduct our exploration.

4. Generating Panic: A Virtual Experiment

Just as it is often useful to study the behavior of human subjects under controlled conditions, so too is it useful to examine the behavior of a computational model across multiple values of its parameters. In the present case, we shall employ such a virtual experiment in order to draw inferences regarding the factors which encourage or discourage the formation of belief panics under the Bayesian Belief Feedback Model. Like any other experiment, this virtual experiment requires attention to issues of measurement, experimental design, and data analysis; we shall proceed, then, to treat these matters in turn.

4.1 Measuring Panic

In order to assess the manner in which the various models under study generate panic within populations, it is first necessary to specify a concrete means of measuring panic. This is a non-trivial task: panic, as defined previously, is a relatively spontaneous, transient phenomenon which occurs in the presence of (but which is distinct from) background noise. Further complications are introduced by the socio-spatial nature of the phenomenon; when panics occur in large, complex social networks, it may be extremely difficult to verify their existence simply due to the amount of computation required to search the entire network for regions of minority opinion (and to track their evolution over time). Because of these difficulties, a simplified approach is here employed which renders measurable certain key aspects of the phenomenon and studies these across a range of models. These results are then compared with qualitative investigations of model behavior under more restricted circumstances, in order to construct an overall picture of model behavior. While this approach is somewhat less than ideal, it nevertheless offers an initial means of studying a (currently) fairly poorly-specified phenomenon in a formal manner.

As noted earlier, panic is here considered to be the result of an influence process which causes the emergence and subsequent dissolution of regions of minority opinion in the midst of an otherwise stable majority. Presumably, this influence process is itself triggered by the introduction of novel information from the environment, such that some number of individuals become persuaded to change their views and then proceed to influence others. While influence from external sources is theoretically important as an initiating condition, then, it is not a part of the panic itself; any measure of panic, then, must "edit out" these environmental influences. Similarly, the presence of minority opinion per se is not evidence of panic as we have defined it. Stable minority regions are quite possible under the Bayesian Belief Feedback Model (among others), but are not transient phenomena and must be omitted from our measurements. Putting all of this together, then, we find that measuring panic means (in some sense, at least) assessing a minority influence process while controlling both on external information and the presence of stable regions of opinion.

The measure suggested here as one means of satisfying the above conditions is the *belief direction change index*, C_{BD}, defined as the population sum

$$[12] \quad C_{BD}(t) = \sum_{i=1}^{N} I(\theta_{i,t-1}, \theta_{i,t})$$

where the directional influence function $I(\theta_{i,t-1}, \theta_{i,t})$ is given by

$$[13] \quad I(\theta_{i,t-1}, \theta_{i,t}) = \begin{cases} 1 & \text{if } \theta_{i,t-1} \le 0.5, \theta_{i,t} > 0.5, \iota^* \\ -1 & \text{if } \theta_{i,t-1} > 0.5, \theta_{i,t} \le 0.5, \iota^* \\ 0 & \text{otherwise} \end{cases}$$

and ι^* represents the logical statement that actor *i*'s belief change is due to a social influence process[11] (as opposed to observation of environmental signals). Thus, C_{BD} indicates the net effect of social influence on the direction of belief in the population between time periods. Because of this, it is insensitive to both environmental influences (which are discounted) and to the presence of stable groups (which do not contribute to change in belief). Panics, on the other hand, should cause a clear change in the population's C_{BD} score, which can easily be measured and compared across experimental conditions.

One means of studying the presence of panic within various models is to examine the evolution of the C_{BD} measure over time; another, by contrast, is to assemble summary statistics on C_{BD} values and to relate these to experimental conditions. It is this approach which will be pursued here. In particular, we shall be concerned with the mean and spread (as measured by the standard deviation) of the C_{BD}, as well as the number of time steps in which the absolute value of the C_{BD} exceeded some critical threshold. The rationale for this last is clear: insofar as they represent episodes of unusually strong social influence, panics will yield C_{BD} values which are much higher than is typical for the population (which is likely to have some background belief change due to the periodic "waffling" of majority nodes who are exposed to unusually high numbers of minority signals from the environment). As a simple criterion for a threshold value, a difference of two standard deviations from the population mean value is here considered to be sufficiently large in magnitude to suggest the presence of panic. This may be overly conservative, but provides us with fairly strong evidence for the phenomenon in question. Given the degree of ambiguity present in our working definition of panic, this would seem to be a prudent fashion in which to proceed.

[11] Due to the nature of the extended Bayesian Belief Feedback Model, it is generally possible to separate social from environmental influences. This becomes somewhat problematic under the signal transfer model; the specific implications are discussed below.

4.2 Experimental Design

Having established a set of measurements which allow for the detection of panic, we now turn to the question of experimental design. As always, we would like to consider as wide a range of possible values across as wide a range of relevant variables as possible; unfortunately, computational constraints limit the number of conditions which we can study, and the number of replicate observations we can perform per condition. With this in mind, we shall here limit ourselves to consideration of six experimental variables: saliency model, minimum degree, network configuration, signal rate, minimum signal transfer probability (signal transfer saliency model only), and maximum signal transfer probability (likewise). Population size and length of model runs (to name two possible variables) will be held constant; as we shall see, this is not a terribly crippling limitation.

The four saliency models which will be used here have been described above; their present implementation follows that description. Signal rate, which controls the number of (randomly allocated) external signals per time step, is allowed to vary from a low value of 50 (an average of one per four actors per round) to 200 (an average of one per actor per round). Obviously, some actors may receive multiple signals in a given round. In the present implementation, external signals are presumed to be "belief" or "disbelief" oriented with equal probability; the results which follow, then, do not emerge from systematic signal bias. Transfer probabilities for the signal transfer model are set such that both a minimum of 0.0 and a maximum of 1.0 are possible, with various limitations in between. This affects both the probability of signal transfer per se and also the degree to which an actors' beliefs can "filter" his or her signals.

As noted previously, numerous researchers (e.g., (Victor, 1993), (Hicks, 1991)) have suggested that social network density plays an important role in the development of panic. In particular, a tendency of panics to occur in tight-knit, highly inbred subgroups has been posited, presumably due to the redundancy of information provided by alters (Granovetter, 1973). In order to examine this hypothesis, two basic network structures are considered here: a "flat" structure in which actors have an equal probability of being tied to any given alter, and a "bridged subgroup" structure in which actors are divided into highly interconnected subgroups with bridging ties distributed randomly among all actors. In each case, the base degree (number of alters an ego may nominate) of each actor is controlled exogenously and is fixed across actors, though the actual degree may be higher due to nomination of ego by multiple alters. Thus, while we here sample across the space of possible networks, we restrict ourselves to certain structural forms which are of theoretical interest.

The combined experimental conditions for the assessment of the BBFM are given in Table 1 below. Note that each condition is replicated five times, and that each simulation is executed for 1100 time steps with a population of 200. This does not change across conditions. After execution, each series of C_{BD} values is collapsed to find the mean C_{BD} and the panic index as described above. Analysis of the data which results, then, should allow us to draw general inferences regarding the behavior of the Bayesian Belief Feedback Model.

order to examine this hypothesis, two basic network structures are considered here: a "flat" structure in which actors have an equal probability of being tied to any given alter, and a "bridged subgroup" structure in which actors are divided into highly interconnected subgroups with bridging ties distributed randomly among all actors. In each case, the base degree (number of alters an ego may nominate) of each actor is controlled exogenously and is fixed across actors, though the actual degree may be higher due to nomination of ego by multiple alters. Thus, while we here sample across the space of possible networks, we restrict ourselves to certain structural forms which are of theoretical interest.

The combined experimental conditions for the assessment of the BBFM are given in Table 1 below. Note that each condition is replicated five times, and that each simulation is executed for 1100 time steps with a population of 200. This does not change across conditions. After execution, each series of C_{BD} values is collapsed to find the mean C_{BD} and the panic index as described above. Analysis of the data which results, then, should allow us to draw general inferences regarding the behavior of the Bayesian Belief Feedback Model.

Table 1: Design for the Virtual Experiment	
Treatment	**Values**
Population Size	200
Simulation Time	1100 Time Steps[12]
Saliency	Constant Excitation, Signal Excitation, Change Excitation, Signal Transfer
Network Configuration	Flat, Bridged Clusters
Minimum Degree	4, 8, 12
Signal Rate	50, 100, 200
Minimum Transfer Probability	0.0, 0.1, 0.2
Maximum Transfer Probability	0.5, 0.8, 1.0
Total Conditions	648
Replications per Condition	5
Total Observations	3240
Limitations	Due to computational constraints, Signal Transfer model experiments included only degrees 4 and 8 of the flat network condition, thus reducing the actual number of observations to 2590. For the same reason, depth of recursion in the signal transfer sub-model was restricted to 5 hops.

4.3 Data Analysis

As indicated above, we are interested in assessing, for the present model, the degree to which various experimental treatments influence the belief direction change index, C_{BD}. In particular, we would like to assess the effects of these parameters on the number of simulation time steps in which the C_{BD} is more than two standard deviations away from the mean, divided by the total number of time steps (the "panic index," or rate of presumed panic activity). A number of methodological questions remain, however, which must be resolved prior to undertaking such an analysis.

The first of these questions pertains to the "window" over which C_{BD} values are to be sampled. This amounts to the issue of stationarity of the present model[13]: if the model is in fact

[12] Of which the first 100 are removed in order to avoid transient effects; see below.
[13] It should be noted that we are *not* attempting a time series analysis of the C_{BD} index, as our goal is to predict rates of occurrence across time rather than particular values of a time-dependent variable at particular times. The methods utilized here would be inappropriate for such a task.

stationary, then out particular choice of window will be fairly unimportant. If the model is not stationary, however, additional complications arise. Consider, for instance, the possibility that C_{BD} (or its variance) changes systematically over time; if this is in fact the case, then any prediction of the panic index will be biased by the particular interval over which the sample is taken. Is this in fact a problem for the present experiment? The answer, as it turns out, is no. Figure 1 shows the C_{BD} values produced by a single run of the Bayesian Belief Feedback Model; as can be seen, these values become quite stationary after a brief initial period. Systematic examination of multiple runs across multiple conditions[14] reveals this pattern to be a stable feature of the model's behavior. As a result, it was decided to remove the first 100 time steps from the sampling window[15], leaving a sample of 1000 time steps for each replication.

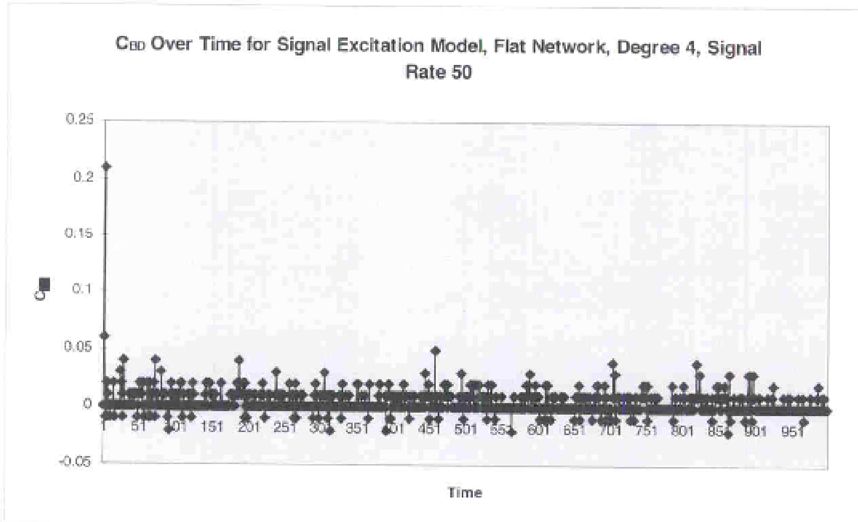

Figure 1: A Sample Experimental Run

[14] All combinations of saliency model, network type, and degree treatments employed in this experiment were examined; no deviation from the indicated pattern was found.
[15] No obvious transient behavior was observed beyond 70 time steps under any experimental conditions, with most conditions stabilizing in under 15 steps.

Having dispensed with the question of stationarity, we are now in a position to examine the C_{BD} index, and likewise our index of panic behavior. This raises our second methodological question: how are these variables distributed? Figures 2 and 3, below, give the distributions of the mean C_{BD} and the panic index for our experimental data. As can be seen, neither variable is normally distributed; both, in fact, are heavily positively skewed. This positive skew suggests that panic events' frequency is inversely related to their size: a very large number of small panics occur for every massive event[16]. This general pattern is a very common finding across a wide range of systems, measures, and phenomena (including the incidence of earthquakes (Sornette and Sornette, 1989), evolution of cellular automata (Langton, 1990), size distributions of firms (Simon, 1995), and frequency of word use (Zipf, 1949)), and is often considered to be one of several properties of systems in the complex regime. While the importance of this finding here should *not* be overstated, its consistency with previous work on qualitatively similar phenomena bears mention. Regardless of its importance or irrelevance, however, this distribution of values is not well-suited to standard regression techniques. In order to avoid the difficulties associated with fitting linear models to such data, therefore, we shall here employ a standard natural log transformation to remove skew from the index values. (While several alternatives were tried, including root transformations, the natural logarithm was found to result both in the best distribution and in the most reasonable overall fit.) All treatment of these variables henceforth will consider this transformed data[17].

[16] It should be emphasized that this is an *interpretation* of the result. An alternative possibility is that a heavily skewed relationship exists between the frequency of panic events within a condition and the frequency of that occurrence across conditions. Due to the nature of the panic index, this cannot be decided by the index alone.

[17] Similar analyses carried out on the raw data showed an overall extremely poor model fit, with patterned residuals and $R^2 << 0.1$.

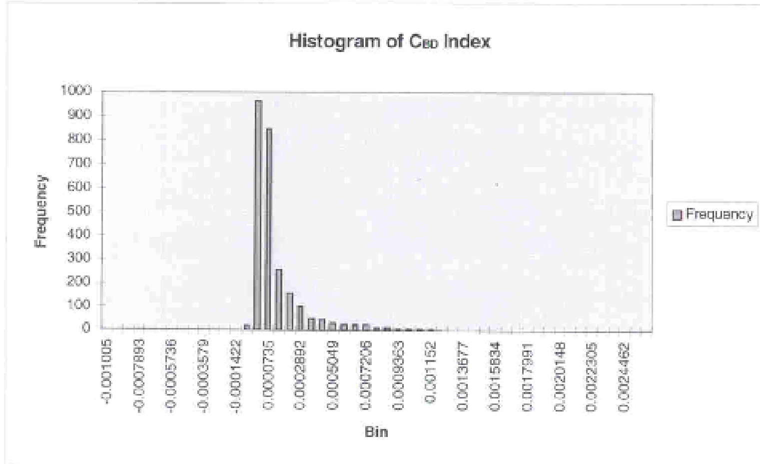

Figure 2: Distribution of Mean C_{BD} Values

Figure 3: Distribution of the Panic Index

After verifying the reasonability of the panic index, and after adjusting the index to remove positive skew, we may now attempt to fit a linear model to the transformed data. While it is naturally quite possible (perhaps even probable) that the underlying effects being studied are not in fact (log) linear, linear regression techniques are nevertheless effective at revealing general trends in data such as that considered here. In the absence of a well-specified theory of nonlinear effects, likewise, the fitting of a linear model can serve as a reasonable (and conservative) initial approach to data modeling. Qualitative results from these analyses, then, can be used to guide further investigation.

The first model we shall use to fit the panic index data is also the largest: here, we regress the logarithm of the panic index values on the natural logs of the experimental treatment variables, and on the full set of interactions between saliency sub-models and other conditions[18]. By attempting to fit such a comprehensive model, we are able to control for effects both at the sub-model (interaction term) and full-model level; the results of this regression are given in Table 2. As can be inferred from the modest R^2, the fit for this model is less than excellent. While it is able to explain approximately one third of the variance in the panic index, its residual standard error is reasonably high and patterning was observed in the model's residuals. This is not unexpected, given the nature of the data in question: even the aggregate behavior of a highly nonlinear system is not always amenable to simple linear prediction. Nevertheless, the model below has a great deal to tell us about the qualitative effects of the various experimental treatments. Saliency sub-models, for instance, had significant overall effects on the panic rating. Constant excitation and change excitation, all other things being equal, were found to have a negative effect on the panic index, while signal based excitation increased it. Signal rate had an overall positive effect on the panic index, while, interestingly enough, network configuration and minimum degree appeared not to; this picture becomes more complicated, however, when one considers the interactions between saliency model and other treatment variables. Minimum degree, for instance, while insignificant alone, has a strong negative effect in conjunction with the signal excitation model (suggesting that this model is more sensitive to changes in network density than the others). Signal rate, similarly, produced a weak negative effect with signal excitation but reversed itself to cause a significant *positive* effect under constant excitation! Signal transfer maximum and minimum probabilities (which, because they were set to 0 for all other saliency models, also acted as dummy variables for the saliency model itself) were also multidirectional in effect, suggesting that probability thresholds which limited actors' ability to invoke their own beliefs in passing on information increased the panic index[19]. Clearly, our

[18] The statistical analyses described here were performed using S-Plus.
[19] An alternative interpretation of this finding is that certain probabilities of signal transfer (near 0.5) are more productive of panic than others. While this is possible, informal experimentation suggests that the other interpretation is more likely.

intuition that different saliency models behave in very different ways is borne out even by these simple results.

Table 2: Regression of log(Panic Index + 1) on Experimental Treatments[20]

	Value	Std. Error	t value	Pr(>|t|)
(Intercept)	0.0053	0.0367	0.1438	0.8857
log(ConExcite + 1)	-0.1119	0.0548	-2.0417	0.0413
log(SigExcite + 1)	0.2266	0.0548	4.1350	0.0000
log(ChgExcite + 1)	-0.1006	0.0548	-1.8353	0.0666
log(NetType + 1)	0.0017	0.0029	0.5958	0.5514
log(MinDegree + 1)	-0.0023	0.0122	-0.1896	0.8496
log(SigRate + 1)	0.0107	0.0043	2.4496	0.0144
log(TransPMin + 1)	0.1387	0.0317	4.3818	0.0000
log(TransPMax + 1)	-0.0879	0.0190	-4.6334	0.0000
log(ConExcite*NetType + 1)	-0.0039	0.0041	-0.9650	0.3346
log(SigExcite*NetType + 1)	-0.0020	0.0041	-0.4898	0.6243
log(ConExcite*MinDegree + 1)	-0.0094	0.0125	-0.7523	0.4519
log(SigExcite*MinDegree + 1)	-0.0583	0.0125	-4.6666	0.0000
log(ChgExcite*MinDegree + 1)	0.0190	0.0125	1.5227	0.1280
log(ConExcite*SigRate + 1)	0.0143	0.0047	3.0518	0.0023
log(SigExcite*SigRate + 1)	-0.0095	0.0047	-2.0205	0.0434
log(ChgExcite*SigRate + 1)	-0.0012	0.0047	-0.2579	0.7965

Residual standard error: 0.02829 on 2572 degrees of freedom
Multiple R-Squared: 0.3345
F-statistic: 80.8 on 16 and 2572 degrees of freedom, the p-value is 0

While the full model presented in Table 2 highlights the most robust effects of experimental treatments on the panic index, it does not represent an optimal fit to the data. As a result, some variance inflation on coefficient estimates may have been present, undermining the significance of weaker predictors[21]. To compensate for this, three alternative methods were used to select best-fit linear models from the above set of predictive variables. These methods - stepwise model selection (F entry/exit threshold of 2.0), leaps and bounds using Mallow's C_p, and leaps and bounds using adjusted R^2 - produced somewhat reduced models with more reliable estimates

[20] Because of singularities, redundant variables and interactions (such as SigTrans) were removed from this and the following regression models.
[21] One source of such variance inflation is multicollinearity. While our experimental design systematically eliminates the possibility of multicollinearity for most variables, this is not the case for interaction effects; the presence of a large number of interactions in our initial model may have adversely affected the estimation of their components (and vice versa).

of regressor effects and better-behaved residuals, although no models were identified which resulted in a truly optimal fit.

The first of the reduced models to be considered was produced by the stepwise selection procedure, and is given in Table 3. This model differs from that given above in several respects. First, and perhaps most importantly, the model chosen by the stepwise procedure omits independent saliency effects, with the exception of the signal excitation model which continues to have a strong positive effect on the panic index. Minimum degree, by contrast, is not only present in this model but also highly significant, acting overall to increase panic while reducing it under the constant and signal excitation models. A fairly similar effect may be observed for the external signal rate, which as before produces different effects for different models, this time significantly so for the belief change excitation model.

Table 3: Stepwise-Selected Model of log(Panic Index)

	Value	Std. Error	t value	Pr(>|t|)
(Intercept)	-0.0658	0.0068	-9.7082	0.0000
log(SigExcite + 1)	0.3290	0.0172	19.0785	0.0000
log(MinDegree + 1)	0.0170	0.0023	7.4978	0.0000
log(SigRate + 1)	0.0172	0.0024	7.2923	0.0000
log(TransPMin + 1)	0.1522	0.0306	4.9683	0.0000
log(TransPMax + 1)	-0.0751	0.0168	-4.4774	0.0000
log(ConExcite * MinDegree + 1)	-0.0297	0.0029	-10.1532	0.0000
log(SigExcite * MinDegree + 1)	-0.0776	0.0034	-22.8658	0.0000
log(ConExcite * SigRate + 1)	0.0067	0.0024	2.7482	0.0060
log(SigExcite * SigRate + 1)	-0.0161	0.0030	-5.4380	0.0000
log(ChgExcite * SigRate + 1)	-0.0075	0.0022	-3.4169	0.0006

Residual standard error: 0.02828 on 2578 degrees of freedom
Multiple R-Squared: 0.3331
F-statistic: 128.8 on 10 and 2578 degrees of freedom, the p-value is 0

The second reduced model to be considered was identified by a leaps and bounds procedure which attempted to optimize on various goodness of fit statistics. In particular, two different measures - Mallow's C_p and the adjusted R^2 - were employed in an attempt to find an optimal model. As it happened, both searches identified the same model, which had both the minimum value of C_p (8.120117) and the maximum adjusted R^2 (33.137%). This model is given in Table

4, below. One notable difference between this and the preceding model lies in the selection of variables for inclusion: this model explicitly includes both signal based excitation and change based excitation, and drops the main effect for minimum degree in exchange for an extra interaction term (losing the interaction between signal rate and the change based excitation saliency model in the process). Interpretively, this allows us to see the strong negative effect of the continuous and change based excitation models on the panic index (effects which were noted in the first model) and permits us to note that the positive effect of minimum degree seen in the previous model appears constrained in practice to the change excitation model. Effects present in both previous models are present here and consistent (such as the already noted signal transfer variables), and those which did not attain significance elsewhere (such as network type) continue to be left out; the overall picture, then, is consistent with previous models, despite the superficial differences.

Table 4: Optimal Model By Leaps and Bounds Procedure

	Value	Std. Error	t value	Pr(>|t|)
(Intercept)	0.0051	0.0131	0.3891	0.6972
log(ConExcite + 1)	-0.1127	0.0237	-4.7657	0.0000
log(SigExcite + 1)	0.2267	0.0237	9.5827	0.0000
log(ChgExcite + 1)	-0.1012	0.0178	-5.6761	0.0000
log(SigRate + 1)	0.0097	0.0016	5.9935	0.0000
log(TransPMin + 1)	0.1414	0.0304	4.6472	0.0000
log(TransPMax + 1)	-0.0874	0.0189	-4.6274	0.0000
log(ConExcite * MinDegree + 1)	-0.0117	0.0025	-4.6432	0.0000
log(SigExcite * MinDegree + 1)	-0.0606	0.0025	-24.0187	0.0000
log(ChgExcite * MinDegree + 1)	0.0167	0.0025	6.6180	0.0000
log(ConExcite * SigRate + 1)	0.0153	0.0024	6.3652	0.0000
log(SigExcite * SigRate + 1)	-0.0085	0.0024	-3.5558	0.0004

Residual standard error: 0.02827 on 2577 degrees of freedom
Multiple R-Squared: 0.3342
F-statistic: 117.6 on 11 and 2577 degrees of freedom, the p-value is 0

5. Discussion

By comparing and contrasting the results obtained by the three linear models fitted above, it is possible to form some inferences regarding the overall behavior of the BBFM in producing panic within populations. First, and in accordance with our initial intuition on the subject, it would seem that the model of saliency which is used to govern actor communication significantly affects the development of panic. Overall, continuous and change based excitation models appear to reduce propensities towards panic, while the signal based excitation model encourages it. These general tendencies are somewhat misleading, however: as different saliency models respond differently to variables such as minimum degree (here, a proxy for density) and the rate at which novel signals are introduced from the external environment, the actual rate of panic will greatly depend on the particular condition in question. On the other hand, some (albeit puzzling) constants appear clearly from the data. Network structure, for instance, is uniformly non-predictive both as a main effect and within particular models. Of course, many aspects of the social structure are determined by the minimum degree (which *is* predictive) and theoretically by size (which is fixed in this case), and it could easily be argued that insofar as the various saliency models reflect different types of social relations that they too are structural...yet, the failure of bridged clusters to produce higher levels of panic (as implied by Granovetter's information redundancy arguments regarding weak ties, and by ethnographic accounts such as those of Victor (1993) and Hicks (1991)) is troubling. Is it actually the case that subgroup configuration is irrelevant to the generation of panics in populations? Or is there simply some aspect of the way in which panics are measured here which hides the true relationship? It may be the case that such effects are only pertinent to the signal transfer model (for which sufficient data was not available to permit their estimation), but this seems rather counterintuitive in light of the communication which goes on under all saliency models. This is a matter for future research, but the question would definitely seem to be an important one.

Another issue regarding the present results concerns their general inability to account for the majority of the variance in the model's behavior. Given the degree of structure observed in the residuals of the above regressions, it would seem that *something* is going on; we are not simply facing a high-variance phenomenon. Whatever is happening, evidently, is either too subtle to be captured via our rather blunt instruments, or else depends upon some other aspect of the experimental condition which was not controlled for. Considering the controls which were in place, such an aspect could well be structural...although it is not obvious what structural variable could govern the results independently of both subgroup configuration and minimum degree. In either case, further systematic study will be needed to explain the behavior of the panic index under the Bayesian Belief Feedback Model.

6. Conclusions

Beginning with a simple model of influence and communication, we have here constructed a general framework for the simulation and measurement of the emergence of panic phenomena within structured populations of actors. The panics which develop within this framework exhibit distributional properties which are typical of complex phenomena, and their incidence can be predicted to some degree using simple linear models. The particular assumptions made regarding actor communication and issue saliency have strong effects on the emergence of panic, and interact with situational elements such as network density and the introduction of environmental signals in non-trivial ways. By contrast, features such as the presence or absence of bridged subgroups within the population's social network appear to have no effect on panic, and do not interact with other variables to produce significant effects.

While the present work utilizes a fairly general measure of panic, this measure is limited in its ability to fully represent our sociological intuitions regarding the term. Further research is needed to uncover better ways of operationalizing the definition of panic itself, and of discerning the properties of such measures. Likewise, the notion of "panic" considered here is very limited - only panics in belief are included - leaving open the possibility of a wide array of complementary research on panics of other sorts. Though relatively amenable to formal study and treated extensively by ethnographers, this area is currently underinvestigated by computational and mathematical sociologists; it is hoped that this work will help rectify this unfortunate situation.

References

Anderson, N. H. (1959), "Test of a Model for Opinion Change," *Journal of Abnormal and Social Psychology*, 59, 371-81.

-- (1964), "Linear Models for Responses Measured on a Continuous Scale," *Journal of Mathematical Psychology*, 1, 121-42.

-- (1971), "Integration Theory and Attitude Change," *Psychological Review*, 79, 171-206.

Asch, Solomon (1952), *Social Psychology*, NJ: Prentice Hall.

Baron, Robert, Hoppe, Sieg, Kao, Chuan Feng, Brunsman, Bethany, Linneweh, Barbara and Rogers, Diane (1996), "Social Corroboration and Opinion Extremity," *Journal of Experimental Social Psychology*, 32, 537-60.

Baron, Robert, Vandello, Joseph and Brunsman, Bethany (1996), "The Forgotten Variable in Conformity Research: Impact of Task Performance on Social Influence," *Journal of Personality and Social Psychology*, 71(5), 915-27.

Banerjee, Abhijit V. (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107(3), 797-817.

Bikhchandani, Sushil, Hirshleifer, David and Welch, Ivo (1992), "A Theory of Fashion, Custom, and Cultural Change," *Journal of Political Economy*, 100(5), 992-1026.

Brass, Daniel (1984), "Being in the Right Place: A Structural Analysis of Individual Influence in an Organization," *Administrative Science Quarterly*, 29.

Bulgatz, Joseph (1992), *Ponzi Schemes, Invaders from Mars and More Extraordinary Popular Delusions*, New York: Harmony Books.

Burt, Ronald (1987), "Social Contagion and Innovation: Cohesion Versus Structural Equivalence," *American Journal of Sociology*, 92, 1287-1335.

-- (1992), *Structural Holes: The Social Structure of Competition*, Cambridge: Harvard University Press.

Butts, Carter (1997), "A Bayesian Model of Panic in Belief," Working Paper.

Carley, Kathleen (1990a), "Group Stability: A Socio-Cognitive Approach," *Advances in Group Processes*, 7.

-- (1990b), "On the Persistence of Beliefs," Working Paper, Department of Social and Decision Sciences, Carnegie Mellon University.

-- (1991), "A Theory of Group Stability," *American Sociological Review*, 56(3), 331-354.

-- (1995), "Communication Technologies and Their Effect on Cultural Homogeneity, Concensus, and the Diffusion of New Ideas," *Sociological Perspectives*, 38(4), 547-571.

Coleman, James (1990), *Foundations of Social Theory*, Cambridge: Harvard University Press.

Degroot, Morris (1989), *Probability and Statistics*, 2nd Edition. Mass.: Addison-Wesley.

Deutsch, Morton and Gerard, Harold B (1955), "A Study of Normative and Informational Social Influences upon Individual Judgment," *Journal of Abnormal and Social Psychology*, 51(3), 629-36.

Erikson, Kai (1966), *Wayward Puritans: A Study in the Sociology of Deviance*. New York: Wiley.

Friedkin, Noah, and Cook, Karen (1991), "Peer Group Influence," *Sociological Methods and Research*, 19, 122-143.

Goel, Prem and Iyengar, N. Sreenivas, (Eds) (1992), *Bayesian Analysis in Statistics and Econometrics*, Lecture Notes in Statistics, 75, New York: Springer-Verlag.

Goode, Erich and Ben-Yehuda, Nachman (1994), "Moral Panics: Culture, Politics, and Social Construction," *Annual Review of Sociology*, 20, 149-71.

Granovetter, Mark (1973), "The Strength of Weak Ties." *American Journal of Sociology*, 78, 1360-80.

-- (1978). "Threshold Models of Collective Behavior." *American Journal of Sociology*, 83, 1420-43.

-- (1982), "The Strength of Weak Ties: A Network Theory Revisited," in Marsden and Lin (Eds), 1982.

-- (1985), "Economic Action and Social Structure: the Problem of Embeddedness," *American Journal of Sociology*, 91, 481-515.

Hicks, Richard (1991), *In Pursuit of Satan: The Police and the Occult*, Buffalo: Prometheus Books.

Jeffrey, Richard (1983), *The Logic of Decision*, Chicago: University of Chicago Press.

Krackhardt, David (1996), "Organizational Viscosity and the Diffusion of Controversial Innovations," *Journal of Mathematical Sociology*, in press.

Langton, Chris (1990), "Computation at the Edge of Chaos: Phase Transitions and Emergent Computation," *Physica D*, 42, 12-37.

Latané, Bibb (1981), "The Psychology of Social Impact," *American Psychologist*, 56(4), 343-56.

Latané, Bibb and L'Herrou, Todd (1996), "Spatial Clustering in the Conformity Game: Dynamic Social Impact in Electronic Groups," *Journal of Personality and Social Psychology*, 70(6), 1218-30.

Lux, Thomas (1995), "Herd Behaviour, Bubbles and Crashes," *The Economic Journal*, 105(431), 881-96.

Macy, Michael (1991), "Chains of Cooperation: Threshold Effects in Collective Action," *American Sociological Review*, 56, 730-47.

March, James and Simon, Herbert (1958), *Organizations*, New York: Wiley.

Mehan, Hugh, and Wood, Houston (1975), *The Reality of Ethnomethodology*, John Wiley and Sons.

Morgan, D.L., Neal, Margaret B., and Carder, Paula. (1997), "The Stability of Core and Peripheral Networks Over Time," *Social Networks*, 19, 9-25.

Orléan, André (1989), "Mimetic Contagion and Speculative Bubbles," *Theory and Decision*, 27, 63-92.

Schelling, Thomas C. (1969), "Models of Segregation," *American Economic Review*, 59(2), 483-93.

Scott, John (1991), *Social Network Analysis: A Handbook*, London: Sage Publications.

Sherif, M. (1936), *The Psychology of Social Norms*, New York: Harper.

Shiller, Robert J. (1995), "Conversation, Information, and Herd Behavior," *American Economic Review*, 85(2), 181-85.

Simon, Hebert (1955), "On a Class of Skew Distribution Functions," *Biometricka*, 42, 425-440.

Sornette, A. and Sornette, D. (1989), "Self-Organized Criticality and Earthquakes," *Europhysics Letters*, 9(3), 197-202.

Suitor, Jill, Keeton, Shirley (1997), "Once a Friend, Always a Friend? Effects of Homophily on Women's Support Networks Across a Decade," *Social Networks*, 19, 51-62.

Victor, Jeffrey S. (1993), *Satanic Panic: The Creation of a Contemporary Legend*, Open Court.
-- (1995), "The Dangers of Moral Panics," *Skeptic*, 3(3), 44-51.

Wasserman, Stanley and Faust, Katherine (1994), *Social Network Analysis: Methods and Applications*, Cambridge: Cambridge University Press.

Wellman, Barry, Wong, Renita Yuk-lin, Tindal, David, and Nazer, Nancy (1997) "A Decade of Network Change: Turnover, Persistence and Stability in Personal Communities," *Social Networks*, 19, 27-50.

Zipf, G. K. (1949), *Human Behavior and the Principle of Least Effort*, Reading: Addison-Wesley.