Title
The burden of social proof: Shared thresholds and social influence.

Journal
Psychological Review, 119(2)

ISSN
1939-1471 0033-295X

Author
MacCoun, Robert J

Publication Date
2012

DOI
10.1037/a0027121

Peer reviewed

Psychological Review
2012, Vol. 119, No. 2, 345–372

0033-295X/12/$12.00 DOI: 10.1037/a0027121

The Burden of Social Proof: Shared Thresholds and Social Influence

Robert J. MacCoun
University of California at Berkeley

Social influence rises with the number of influence sources, but the proposed relationship varies across theories, situations, and research paradigms. To clarify this relationship, I argue that people share some sense of where the "burden of social proof" lies in situations where opinions or choices are in conflict. This suggests a family of models sharing 2 key parameters, one corresponding to the location of the influence threshold, and the other reflecting its clarity—a factor that explains why discrete "tipping points" are not observed more frequently. The plausibility and implications of this account are examined using Monte Carlo and cellular automata simulations and the relative fit of competing models across classic data sets in the conformity, group deliberation, and social diffusion literatures.

Keywords: conformity, deliberation, jury, influence, threshold

Supplemental materials: http://dx.doi.org/10.1037/a0027121.supp

Regimes topple, neighborhoods gentrify, financial bubbles collapse, and fads burst onto the scene. These stark discontinuities of social life galvanize our attention, marked by many labels, including critical mass, information cascades, bandwagons, domino effects, the "hundredth monkey phenomenon," and most famously, tipping points—a phrase attributed to Grodzins (1958), formalized by Schelling (1969, 1978), and popularized by Gladwell (2000).

Tipping points surely exist, but they are far from ubiquitous in social life. Change often happens slowly, linearly, and gradually. One reason might be that we are often situated near a critical point of inflection. But there are many domains in which we observe considerable shifts in social movement—crowds assembling at an event, new products gaining market share, politicians mustering popular support—without seeing anything as starkly discontinuous as a tipping point. Thus, a Google search (March 9, 2011) turned up about 3.8 million references to a "tipping point" and 5 million to a "critical mass" (some of which may involve nonsocial phenomena)—but there are over a billion entries matching the search string "increasing* popular".

The theoretical account offered here attempts to clarify several unresolved issues in the literature. One puzzle is how to reconcile seemingly knife-edged "tipping points" with the more gradual change that is the norm. Ideally, a single theory should explain both types of response patterns and clarify when each should be observed. Another puzzle is why the impact of the first few influence sources qualitatively differs across settings and research paradigms: a concave "r-curve" (Gerard, Wilhelmy, & Conolley, 1968; Milgram, Bickman, & Berkowitz, 1969), an "s-curve" (observed in the classic Asch study and in research on small group decision processes), or even a checkmark pattern (Cialdini, Reno, & Kallgren, 1990). While many constructs and processes in the influence literature can be deployed to verbally explain such qualitative discrepancies, our existing formal models fail to account for them. And a third puzzle is why models that aptly describe behavior in one influence paradigm (e.g., conformity, group deliberation, or diffusion of responsibility) fare more poorly in others. These paradigms differ in psychologically meaningful ways, and they were developed with the intention of capturing different social influence situations. But none of the current formal models successfully reconciles these paradigms or clarifies their differences. The models I propose here help to resolve all three puzzles using a common integrative framework. The models serve a predictive and explanatory purpose but also provide an empirical tool for estimating psychologically meaningful parameters across social settings and research paradigms.

Threshold Concepts in Psychology

Kimble (1996) identified the concept of thresholds as one of a handful of bedrock principles belonging at the core of a unified scientific psychology. In some areas of psychology, we make routine use of thresholds to interpret data; examples include statistical significance testing (e.g., Cohen, 1994), signal detection theory (e.g., Swets, 1988), dose-response functions in pharmacology (Berkson, 1944), and item-response theory in testing (e.g., Reise, Ainsworth, & Haviland, 2005). But psychologists have made more use of threshold concepts to explain perception and judgment than to explain social behavior.

Many social behaviors involve continuously varying responses—how fast one walks down a hallway, how much one contributes to a public good, how generously one tips a waiter. And for methodological reasons, psychologists often prefer continuous or at least interval-level dependent measures. But much social action takes a more binary, "digital" form, and this type of action is often rich in

This article was published Online First February 20, 2012.

The arguments developed here owe much to my long collaboration with Norbert Kerr. I am grateful for his generosity, patience, and enthusiasm as a mentor and friend.

Correspondence concerning this article should be addressed to Robert J. MacCoun, Goldman School of Public Policy, University of California at Berkeley, 2607 Hearst Avenue, Berkeley, CA 94720. E-mail: maccoun@berkeley.edu

social meaning (see Harré, Clarke, & De Carlo, 1985; Watzlawick, Bavelas, & Jackson, 1967)—voting for a candidate, joining an organization, choosing a university, deciding to drink at a party, deciding to use a condom, proposing marriage, quitting one's job. Nonlinear threshold concepts are an important tool for understanding how continuous latent preferences get converted into overt categorical choices and actions.

Threshold models conventionally assume that the probability of a response varies as a nonlinear, sigmoidal function of some latent propensity. Examples include models of neuronal firing (Rumelhart, Hinton, & McClelland, 1986), declarative and procedural memory activation (J. R. Anderson et al., 2004), and the behavior of social insects (Camazine et al., 2001). With this kind of sigmoid function, the linear assumption may be a reasonable approximation, but only for part of the range of the latent propensity, and only when the slope of the function is fairly shallow.

The logistic threshold models I propose in this article share features with item-response models in psychometrics (Birnbaum, 1968; De Boeck, Wilson, & Acton, 2005; Reise et al., 2005), as well as discrete choice models in psychology (Augustin, 2005; Luce, 1959), economics (McFadden, 1974, 2001), and sociology (Macy, 1991). But not all logistical models are threshold models, and the threshold models presented here do not require the logistic function. I show that another threshold model, derived by Norbert Kerr (Kerr, MacCoun, & Kramer, 1996, Appendix B) bears little resemblance to a logistic equation but produces similar behavior. I began this project by trying to extend that model beyond its original domain (group deliberation), but I concluded that a logistic threshold model was more suitable as a bridge across paradigms and disciplines. Still, it is important to emphasize that the Kerr model can be successfully applied well beyond its original application.

Previous Threshold Accounts of Social Phenomena

The two most famous threshold models of collective behavior differ in how they treat the issue of variability. In the Schelling (1969, 1971, 1978) "tipping point" model of racial segregation in housing, it is assumed for convenience that all members of the same group share the same discrete threshold—that is, that "everybody wants at least half his neighbors to be like himself" (Schelling, 1971, p. 150). Granovetter's (e.g., Granovetter, 1978, p. 1421) models, in contrast, "take as the most important causal influence on outcomes the *variation* of norms and preferences within the interacting group." In his simplest example, thresholds are distributed uniformly across 100 people, so that the first person ("the instigator") will riot even if no one else does so, the second will riot if at least one other person does so, the third if two others do so, and so on. In his more complex examples, Granovetter assumed for convenience that thresholds have a normal distribution rather than a uniform distribution (also see Granovetter & Soong, 1983).

The models developed in this article capture the essential features of both the Schelling and Granovetter accounts by treating the position of the threshold and the degree of consensus ("clarity") about the threshold as separate parameters. When fit to data from a wide variety of important experimental and real-world social situations, the models imply that social thresholds do vary but that there is nevertheless some social consensus about where they lie in a given domain, even though the location of the threshold varies considerably from domain to domain.[1]

The BOP Framework

This article examines and compares several alternative threshold and nonthreshold models of social influence, but the focus is on a family of closely related logistic threshold models; for convenience, I refer to these as BOP models, where BOP can stand for "burden of (social) proof" or "balance of pressures."

People are acutely sensitive to social consensus information (Asch, 1951, 1955, 1956; Chudek & Henrich, 2011; Festinger, 1954; Fiedler & Juslin, 2006; Leary & Baumeister, 2000), a trait we appear to share with many other organisms (Couzin, 2009). For a dichotomous issue, social consensus judgments require us to track at least two frequencies, those who share our position and those who favor an alternative position. Following Latané's (1981) notation, I label these frequencies T (for targets) and S (for sources). I operationalize outside influence in two different ways—as a ratio (S/T) or as a proportion (S/N, where $N = S + T$). The social pressure may be either active and intended by its sources or passive and unintended by the sources. Consensus information, alone, is sometimes sufficient to bring about opinion change (or at least *stated* opinion change) even in the absence of any contact with or arguments provided by other people (Asch, 1956; Kerr, MacCoun, Hansen, & Hymes, 1987; Mutz, 1998).

Whether or not an actor changes positions in response to social influence will of course be influenced by a great many factors, including informational versus normative influence (Deutsch & Gerard, 1955); compliance, identification, and internalization (Kelman, 1958); or coercive power, reward power, reference power, legitimate power, and expert power (French & Raven, 1960). The motivation to maintain one's position in the face of opposition is also influenced by attitude importance, knowledge, elaboration, certainty, extremity, and accessibility (Visser, Bizer, & Krosnick, 2006), and these in turn are shaped by cognitive style, personality traits, and motivational goals like accuracy, ego protection, and identity expression (Eagly & Chaiken, 1993; Petty & Cacioppo, 1986). The BOP models abstract away these important details and thus form more of a skeletal structure than a complete anatomy of the influence process.

Deriving the Basic Model

In the BOP framework, I treat this complexity in black-box fashion by defining a parameter b as the net effect of these factors on one's resistance to social pressure. Appendix A explains how the model can be derived from either Luce's (1959) choice model or a random utility model; here I offer a more mechanistic account.

The b parameter functions like an internal threshold. Thus the mere existence of opposition to one's position may attract attention, but it will have minimal influence on one's behavior if it falls below the threshold. I model this matching-to-threshold process using a simple linear difference operation:

[1] Signal detection models have also been applied to persuasion and group judgment (e.g., Kriss, Kinchla, & Darley, 1977; Sorkin, Hays, & West, 2001), but these applications are quite different than the social thresholds proposed here.

$$c(\theta - b),$$

where θ represents the degree of opposition (later operationalized as either *S*/*T* or *S*/*N*), *b* is the threshold parameter, and *c* functions as a slope parameter and is hypothesized to represent "norm clarity," as explained below. This function is reminiscent of linear difference operators used in mathematical learning theory (see Vogel, Castro, & Saavedra, 2004) or in cybernetic models (e.g., Powers, 1973), where *c* would be a dampening or gain parameter. I propose that

$$\ln\Omega = c(\theta - b),$$

where Ω represents the odds of changing positions:

$$\Omega = \frac{p_\Delta}{1 - p_\Delta}.$$

Solving for the odds of change,

$$\Omega = e^{c(\theta - b)} = \exp[c(\theta - b)],$$

where $\exp(x) = e^x$, and *e* is the base of the natural logarithm (approximately 2.718). Converting the odds back to a probability, we get

$$p_\Delta = \frac{\exp[c(\theta - b)]}{1 + \exp[c(\theta - b)]},$$

where *b* is a location parameter and *c* is a slope parameter.

Of course, people differ in their susceptibility to social influence, and each person is more susceptible in some settings than in others. Thus, I weight θ by a "max" parameter *m*, which represents the maximum possible external influence in the situation. Dispositionally, *m* might reflect individual differences in traits like agreeableness, need for approval, or public self-consciousness (see Trafimow & Finlay, 1996), but in this article, it will be estimated in the aggregate and can be viewed as a feature of the situation. In this respect, it might reflect "dependence" in the sense that term is used in game theory and interdependence theory (Kelley et al., 2003)—that is, the extent to which desired outcomes are under the partial control of others.

Thus, the core BOP model is

$$p_\Delta = m\frac{\exp[c(\theta - b)]}{1 + \exp[c(\theta - b)]} = \frac{m}{1 + \exp[-c(\theta - b)]}.$$

Table 1 presents several variants of this model adapted for particular research paradigms, as discussed below. Note that in the table, the θ term is replaced with more specific *S*, *T*, and *N* relationships.

Interpreting the Model Parameters

In the terminology of classical persuasion theory (e.g., McGuire, 1968), *b* can be conceived of as an index of resistance to persuasion. But in the present account, I argue that *b* is usefully conceived as an index of the *burden of social proof*. This is only partially analogous to standards of proof or decision criteria as derived in legal theory, statistical decision theory, or signal detection theory. The notion of proof here is not epistemological so much as it is social—what Cialdini (2001) called "social proof." Specifically, the question is not "how much evidence must I have before I make this change?" but rather "how unpopular must my position be before I'm willing to change it and adopt the majority view?" (see Stasser, Kerr, & Davis, 1989). As we will see, there is

Table 1
The BOP Social Influence Models

Situation	Model
uBOP: Unidirectional influence. Fixed number of targets responding to the influence of 1 to $S_{\max}$ sources.	$p_\Delta = \dfrac{m}{1 + \exp\left[-c\left(\dfrac{S}{T} - b\right)\right]}$
bBOP: Bidirectional influence. 1 to $S_{\max}$ sources and $T = S_{\max} - S$ targets jointly influencing each other in a fixed group of size *N*.	$p_\Delta = \dfrac{m}{1 + \exp\left[-c\left(\dfrac{S}{N} - b\right)\right]}$
cBOP: Contexts in which the normative response is ambiguous but the clarity of the local norm rises with the number of sources.	$p_\Delta = \dfrac{m}{1 + \exp\left[-cS\left(\dfrac{S}{N} - b\right)\right]}$
iBOP: Social imitation. Adopters (*S*) at time *t* become new sources at time $t + 1$.	$\text{Adoption}_t = S_{t-1} + \dfrac{1 - S_{t-1}}{1 + \exp[-c(S_{t-1} - b)]}$
gBOP: Single-peaked goal pursuit (city block metric or one-dimensional Euclidean metric).	$p(\text{Act}) = \dfrac{2m}{1 + \exp\left[c\left\lvert\dfrac{S}{N} - b\right\rvert\right]}$
g2BOP: gBOP with squared Euclidean metric or loss function.	$p(\text{Act}) = \dfrac{2m}{1 + \exp\left[c\left(\dfrac{S}{N} - b\right)^2\right]}$

Note. Across models, *m* is an optional ceiling parameter, setting the maximum influence in the setting; *b* is a threshold parameter expressed as a ratio (uBOP) or a proportion (the other models); and *c* is a "norm clarity" parameter that varies inversely with the variance in *b*. BOP = burden of proof or balance of pressures.

good evidence that even legal standards of proof operate this way in juries (MacCoun & Kerr, 1988).

Although the models in Table 1 can be characterized as discrete choice models (see Appendix A), they do not require a rational choice foundation, and I do not assume the process requires either conscious deliberation or utility maximization. There is good evidence that norms are often primed by situational cues in a fairly automatic, associative fashion (Aarts & Dijksterhuis, 2003; Cialdini et al., 1990; Harvey & Enzle, 1981; Hertel & Kerr, 2001; Kahneman & Miller, 1986). This process may be more like the matching-to-standard mechanisms of self-regulatory theories (e.g., Carver & Scheier, 1998; Powers, 1973) than the deliberative error tradeoffs described in traditional decision theories. But while normative standards may be consistent with error tradeoffs, they are often learned directly from others. This is clearly true for the practice of statistical significance testing, which might seem like a canonical example of deliberative rationality. How many of us would reach the exact .05 norm for statistical significance by independently reasoning about the relative costs of Type I versus Type II statistical errors? Presumably, some would adopt a .20 threshold, some a .01 threshold, and many of us might fail to select a threshold at all, simply treating it as a continuous index (albeit a problematic one) of the reliability of our evidence. In reality, most of us uphold the .05 alpha level as a social convention—like any convention, one that imposes social costs when thwarted.

I suggest that the c parameter can be interpreted as an index of *norm clarity*—the extent to which the operative threshold is indeed a shared convention. The c parameter has the effect of shifting the threshold from a very shallow slope (as c approaches 0) to a step function (as c gets very large). Evocative terms like "tipping points" and "critical mass" aptly describe collective behavior when c is high (Gladwell, 2000) but misleading when it is low.

There are two overlapping ways of interpreting clarity. One is internal and individual; one's sense of $(\theta - b)$ might be noisy, uncertain, and unstable. Another is external and social; there may be variability across people due to a lack of shared understanding of the applicable burden of social proof in a situation. The shared consensus might reflect an explicit voting rule (e.g., "majority rules"; see Stasser et al., 1989), a legal standard of proof ("beyond a reasonable doubt"; MacCoun & Kerr, 1988), the strict enforcement of a law or sanctioning principle ("zero tolerance"), or a shared conceptual system for recognizing a "correct" answer (e.g., arithmetic, symbolic logic, or core disciplinary assumptions; Lakatos, 1970; Laughlin, 1999; MacCoun, 1998).

In Appendix B, I present Monte Carlo simulations that support the notion that the c parameter captures the extent to which thresholds are shared. In these simulations, threshold and clarity levels were sampled under a variety of distributions, and then the implied choice behavior was simulated. The best fitting BOP model parameters match the simulated conditions with considerable accuracy. Appendix B also shows that when fit to random monotone data (i.e., data not generated by any threshold process), the average parameter estimates are what one would expect for linear responses without an abrupt threshold inflection.

The Behavior of the Core Model

Figures 1 and 2 illustrate the effects of the threshold (b), clarity (c), and max (m) parameters on $p(\Delta)$ as pressure increases from zero to five sources (with a single target) in bBOP.[2]

The left panel of Figure 1 shows that as the burden-of-persuasion threshold (b) approaches the maximum number of sources, the function changes shape, shifting from an r-curve to an s-curve to an inverted j-curve ("hockey stick"). Also note a signature characteristic of the model as it is currently formulated. If clarity (c) is low, then when the threshold (b) is near zero, the target has a nonzero chance of changing views, even given minimal external pressure. A threshold near zero implies that the actor has very little confidence and/or places very little value on holding the personal position. With a threshold at zero, the actor has a 50:50 chance of changing views, and even the slightest pressure will be sufficient to tip the balance. (The actor, in effect, says "sure . . . whatever you say.") When clarity is low, the actor may infer that the threshold has already been met. The right panel shows that when clarity is high ($c = 25$ rather than 5), the threshold approaches a step function at the threshold, much more akin to the notion of a "tipping point."

The left panel of Figure 2 shows the effect of the clarity parameter (c) at $b = .5$ (the s-shaped function near the middle of the previous two figures). When clarity is very low, the function flattens out to a linear pattern. But as clarity rises, the function becomes sigmoid and then approaches a step function.[3] At zero clarity, the model predicts that the decision to change positions becomes random, even with a threshold of .5. It is not clear whether any social situations approximate zero clarity, and even in situations of low clarity, such extreme reaction may be damped by the max parameter. Finally, the right panel of Figure 2 shows how the max parameter (m) influences response probabilities. At the back wall of the plot, the actor is highly dependent on the influence sources, and hence highly responsive. As the actor becomes more independent, it shows less response; the completely independent actor is basically impervious to the sources' social pressure.

Comparison to Previous Social Influence Models

In the social psychology literature, there are three major models of the functional form predicting influence as a function of relative faction size (see Table 2): social impact theory (SIT; Latané, 1981; Latané & Wolf, 1981), the other-total ratio (OTR; Mullen, 1983), and the social influence model (Tanford & Penrod, 1984).

[2] The uBOP model produces all the same qualitative behavior illustrated here. I had expected uBOP to provide superior fit for conformity data sets, but analyses presented below suggest otherwise.

[3] Interestingly, the left panel of Figure 2 captures some of the qualitative features of the "cusp catastrophe" of Thom's (1972/1975) catastrophe theory, and it does so with a parametric model closely related to theories of choice and of psychometrics, without some of the heroic assumptions one must make to apply Thom's topological account to psychological phenomena (see Liu & Latané, 1998; Nowak & Vallacher, 1994; van der Maas, Kolstein, & van der Pligt, 2003).

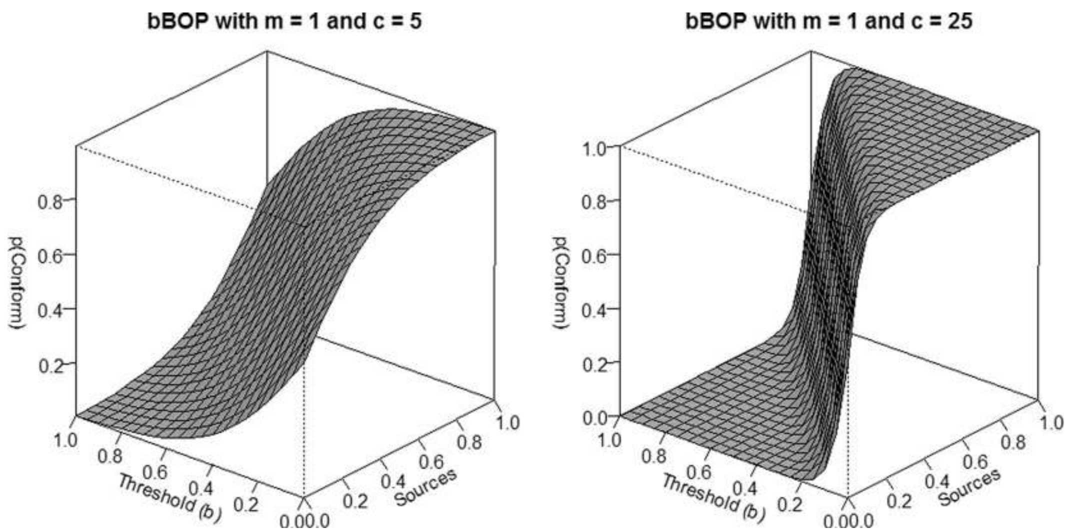

Figure 1. Probability of conformity as a function of the threshold (b) and proportion of sources (S/N) in bBOP. The left panel shows a moderate level of norm clarity ($c = 5$); the right panel shows a high level of norm clarity ($c = 25$). bBOP = bidirectional influence burden of proof model; m = ceiling parameter.

Social impact theory (SIT; Latané, 1981; Latané & Wolf, 1981) is a comprehensive and remarkably successful analysis of collective social influence. Latané and his colleagues (Latané, 1981; Latané & Wolf, 1981; Nowak, Szamrej, & Latané, 1990) have demonstrated the predictive power of SIT across a wide variety of domains, including social imitation, the diffusion of responsibility in helping behavior, social loafing, and anxiety during public speaking. SIT is motivated by a core principle in psychology: diminishing marginal sensitivity to stimuli (see Kahneman, 2003; Stevens, 1957). SIT describes the multiplicative impact (I) of sources (S) on each target (T), as well as the division of impact as any given source's influence is spread across multiple targets. For simplicity, I ignore the role of source strength and immediacy, two features of the original SIT that are not included in the alternative account presented here. If not for the ceiling parameter (here, m for "max"), I could be interpretable as an odds ratio, where $I < 1$ if one's supporters are more influential than one's opponents, $I > 1$ if one's opponents are more influential than one's supporters, and $I = 1$ when the two forces are equivalent. As such, it could be converted into a probability ($p = I/[1 + I]$), although SIT theorists have relied on the ceiling parameter rather than this transformation in applications involving proportional outcomes (e.g., percentage conforming).

Mullen (1983) proposed an alternative social influence function, which he called the *other-total ratio* (OTR); it is simply the ratio of sources to sources + targets, weighted by a ceiling parameter (see Table 2, second row). Mullen demonstrated that this extremely simple model accounts for much of the same data as SIT, and he also applied the model to lynchings and other "mob" phenomena. Mullen motivated the model using theories of self-attention and self-regulation (Carver & Scheier, 1998), although other interpretations are possible.[4] Unlike SIT's I, OTR's I is readily interpretable as the probability of changing one's position. Because of its simplicity, OTR lacks flexibility—one can "squash" the function using the ceiling, but one can't change the basic shape in any psychologically meaningful way. Comparisons of the SIT and OTR models (e.g., Bond, 2005; Latané & Wolf, 1981; Mullen, 1983; Tanford & Penrod, 1984; Tindale, Davis, Vollrath, Nagao, & Hinsz, 1990) have largely ignored the fact that one model implies an odds format while the other can be interpreted as a probability. If the SIT function is converted to a probability, OTR is a special case with an exponent of 1.

As we shall see, SIT and OTR preclude the possibility that a second source has greater impact than the first (Asch, 1956; see Campbell & Fairey, 1989). A third model, Tanford and Penrod's (1984) *social influence model* (SIM), is able to handle this using a more complex functional form. SIM originated in computer simulations of jury behavior (Penrod & Hastie, 1980; Tanford & Penrod, 1983). The computer model's predictions for a 12-person jury setting were then fit to a Gompertz-type growth model (see Table 2, row 3).[5] Tanford and Penrod (1984) held the Gompertz max parameter (here labeled m) constant at 1, but in analyses reported below I found that the model is more accurate when m is treated as a free parameter.[6] SIM is appealing because it provides

[4] As noted by Bond (2005), Stasser and Davis (1981) independently proposed the same function in an analysis of social interaction in groups, which they called the "inform influence function." (A squared version of this function was called the "norm influence function.") I follow convention in referring to the model as OTR.

[5] Tanford and Penrod reported the value 1.75 for the second constant, but Coultas (2004) showed that the value 1.075 is necessary to fit the theoretical curves in their article. I found that the latter version outperforms the original for nearly every data set I examined, so I use the corrected version in this article.

[6] The constants 4 and 1.075 in the formula could also be converted to free parameters to vary across situations, which would make this a full Gompertz model, but Tanford and Penrod (1984) did not take this route. In analyses not reported here, I found that a three-parameter Gompertz model fits data sets nearly as well as uBOP and bBOP, but with parameter values that have less psychological meaning.

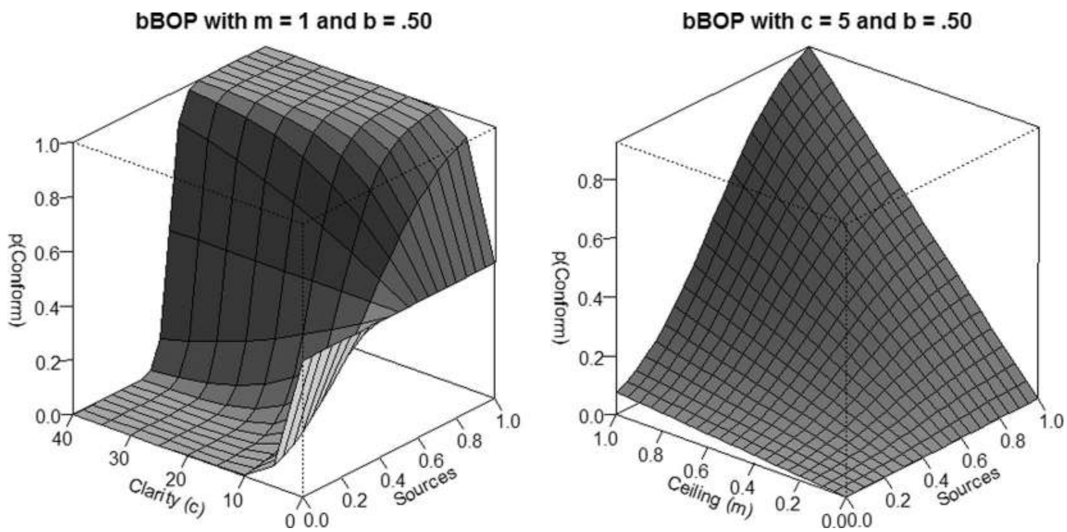

Figure 2. Probability of conformity as a function of the norm clarity parameter (c, left panel) and ceiling parameter (m, right panel) for a varying proportion of sources (S/N). bBOP = bidirectional influence burden of proof model; b = threshold parameter.

a potential link between social psychological experiments and literatures on aggregate growth phenomena (epidemiological incidence models, diffusion models, etc.), which are frequently analyzed using Gompertz models (Allison, 1984). But the SIM model seems more difficult to motivate using psychological principles of psychophysics or choice behavior.

A threshold model of group deliberation, derived by Norbert Kerr, was presented in Appendix B of Kerr et al. (1996), where we noted its ability to accommodate a wide range of social influence patterns. The model (which I refer to as the *Kerr influence model*, or KIM) was proposed for the deliberation paradigm (discussed below) but is readily adapted for use in the unidirectional conformity situation. The model has two parameters, α and β, where $0 \leq \alpha \leq 1$, and $0 < \beta \leq 1$. The model can be expressed with respect to either faction's position, but for comparison purposes, Table 1 expresses it in terms of the probability of adopting the source's position (given the ith distribution of opinion).[7] When $\alpha = .5$, the function is symmetrical; when $\alpha = .5$ and $\beta \rightarrow 0$, the model produces a simple majority rule. When $\beta = 1$, the model reduces to a strict proportionality rule (i.e., OTR). To adapt the KIM model to the unidirectional situation, I simply multiply the above formulas by the same sort of ceiling parameter (m) used in SIT, OTR, and SIM.

Fitting the Models Across Three Experimental Paradigms

The Conformity Paradigm

To examine how well threshold models explain classic influence findings, I examine data from several major research paradigms. In the *conformity paradigm*, multiple sources impinge upon a single target. In the cases I examine here, the sources are confederates of the experimenter, the key independent variable is the number of sources the target encounters, and the dependent variable is the percentage of targets who endorse the source position. This endorsement is public and may reflect only overt public compliance rather than internalization (Kelman, 1958). Because arguments in these cases are held constant, the influence is more likely to be normative than informational (Deutsch & Gerard, 1955; French & Raven, 1960) and will more directly reflect descriptive norms (what I perceive others doing) than injunctive norms (what I believe others think I should be doing; see Cialdini et al., 1990).

Several investigators (Bond, 2005; Campbell & Fairey, 1989; Mullen, 1983; Tanford & Penrod, 1984; Tindale et al., 1990) have compared the relative performance of these models, and the result has been something of a stalemate; the models have similar predictive power, and no consistently superior candidate has emerged. SIM probably suffers most from this stalemate, because it is less parsimonious and more complex to implement. Yet SIM remains the best account of the most famous conformity data—the classic Asch (1951, 1955, 1956) experiments.

Figure 3 shows the best fitting predictions from three nonthreshold models (left panel: SIT, OTR, and SIM) and three threshold models (right panel: bBOP, uBOP, and KIM) for three key conformity studies that have served as a "test bed" in previous discussions of the nature of conformity (Bond, 2005; Latané, 1981; Tanford & Penrod, 1984). The Milgram et al. (1969) data are from a classic field experiment on the streets of New York. Milgram varied how many confederates on the street stared into an empty portion of the sky and recorded whether passersby looked up in the same direction. As seen in the figure, Milgram found the r-shaped concave function that is most typically observed in conformity studies. The Gerard et al. (1968) study is a more rigorously controlled laboratory experiment, with a similar concave pattern. The Asch (1956) data

[7] Table 1 corrects a small but important typographical error in the original presentation. (Using different notation, the original showed $S/N - \alpha$ rather than $\alpha - S/N$ in the numerator of the first expression.)

Table 2
Five Other Social Influence Models

Situation	Model
Social impact theory (SIT; see Latané, 1981)	$I = m\left(\frac{S}{T}\right)^x$
Other-total ratio (OTR; see Mullen, 1983)	$I = m\left(\frac{S}{N}\right)$
Social influence model (SIM; see Tanford & Penrod, 1984)	$I = m * \exp(-4 * \exp(-(S^{1.075/T})))$
Kerr influence model (KIM; see Kerr et al., 1996)	$p_\Delta = \begin{cases} \alpha - \dfrac{\alpha\left(\alpha - \frac{S}{N}\right)^\beta}{\alpha^\beta} & \text{when } S/N \leq \alpha \\ \alpha + \dfrac{(1-\alpha)\left(\frac{S}{N} - \alpha\right)^\beta}{(1-\alpha)^\beta} & \text{when } S/N > \alpha \end{cases}$
Bass diffusion model (Bass, 1969; Mahajan et al., 1995)	$S_t = S_{t-1} + p(m - S_{t-1}) + qS_{t-1}(m - S_{t-1})$, where S_t is the proportion of the population who have adopted the innovation by time t.

come from the most famous demonstration of conformity, with an s-shaped pattern of responses that has posed challenges for past modeling efforts. In Asch's paradigm, an unwitting participant is confronted with zero to S confederates posing as participants, each of whom announces the same judgment that clearly conflicts with the judgment of the true participant (and with the judgments of large samples of nonparticipant judges in pilot testing).

I fit the models by minimizing the mean squared error using the *optim* procedure in the R programming language (see http://www.r-project.org/); very similar solutions were found using Microsoft Excel's Solver function. To compare the relative fit of different models, I used several indices, including the familiar root-mean-squared error (RMSE) and an R^2 statistic computed using least-squares regression with no intercept term. Because the models differ in the number of free parameters (viz., three for bBOP, uBOP, and KIM; two for SIT; and one for OTR and SIM), I also used two fit criteria that put them on a more equal footing, The adjusted R^2 penalizes the raw R^2 for the number of free parameters in each model.[8] The Akaike information criterion (AIC; Akaike, 1974) is a log-likelihood criterion with a very conservative penalty factor ($2k$, where k is the number of free parameters).[9] Table 3 presents the best fitting parameter estimates, and Table 4 presents the fit statistics.

As seen in Table 4, the threshold models (especially bBOP) provide the best fit to these data sets with respect to the RMSE and unadjusted R^2s, and this advantage holds even for the adjusted R^2, which penalizes their extra free parameters. The more parsimonious one-parameter OTR and SIM models fared better with respect to the more conservative index (AIC). But we will see that what the threshold models lose in parsimony (in the conformity paradigm), they gain in generality across paradigms. Arguably, how well these models fit the data seems less interesting and useful than the question of whether they shed light (through the parameter values) on how social influence varies across tasks and settings (see Kerr, Stasser, & Davis, 1979). The uBOP and bBOP parameters suggest that the Asch task provided a higher threshold than did the Milgram and Gerard experiments, arguably because Asch used a simple perceptual task with a seemingly transparent right answer, and there was little influence until two or more confederates (viz., $b = 2.096$ in uBOP) supported the nonobvious response option.

The Cialdini checkmark. Asch's s-shaped conformity function is not the only notable exception to the typical r-shaped function. A more recent example is the checkmark-shaped pattern reported by Cialdini et al. (1990, Studies 2 and 3). Their focus theory of normative conduct has many facets, but of present interest is their claim (Cialdini et al., 1990, p. 1015) that "norms should motivate behavior primarily when they are activated (i.e., made salient or otherwise focused on)." In their Study 2, Cialdini et al. removed all litter from a parking structure and then systematically arranged 0, 1, 2, 4, 8, or 16 pieces of litter and examined how people returning to their cars chose to dispose of a handbill the authors had placed on their windshields. They predicted that "a single piece of litter lying in an otherwise clean environment" should actually reduce littering "by drawing attention to an environment whose descriptive norm (except for one aberrant litterer) was clearly antilitter" (Cialdini et al., 1990, p. 1017). The predicted

[8] The adjusted R^2 statistic usually adjusts for the number of participants in each condition in the experiment; I instead adjusted for the number of experimental conditions, an approach less forgiving toward the three-parameter models.

[9] In this context, $\text{AIC} = 2k + \sum_i -2[p_i \log \hat{p}_i + (1 - p_i)\log(1 - \hat{p}_i)]$, where p_i was the proportion conforming in the ith experimental condition.

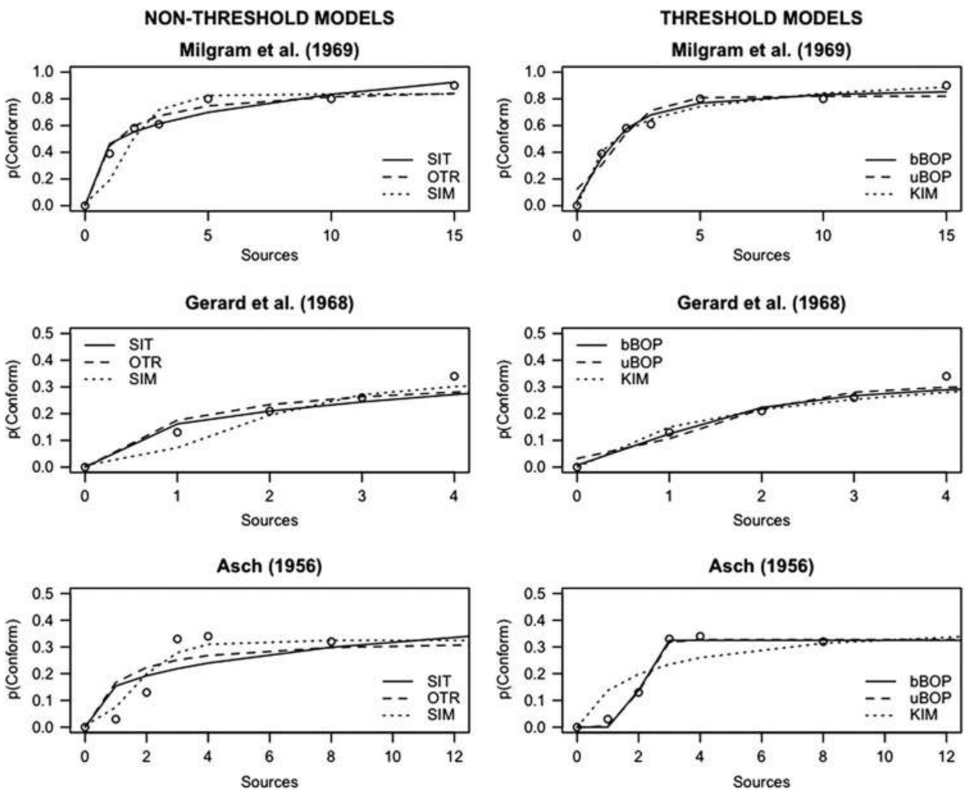

Figure 3. The relative fit of nonthreshold models (left panels) and threshold models (right panels) for three experiments in the conformity paradigm. SIT = social impact theory; OTR = other-total ratio; SIM = social influence model; bBOP = bidirectional influence burden of proof model; uBOP = unidirectional influence burden of proof model; KIM = Kerr influence model.

checkmarked pattern was obtained (Figure 4) and later replicated (with 0, 1, or "many" pieces of litter) in a different setting in their Study 3.

None of the formal models examined so far can produce this checkmarked pattern. But the BOP clarity parameter has the potential to capture the logic of Cialdini's argument. In applications so far, I have made the simplifying assumption that clarity is constant at all levels of source influence. This seems reasonable, but it is not strictly required by the logic of the model. To see whether the model could capture Cialdini's prediction, I first fit the standard bBOP model to the data in Figure 4. I then fit a variant that weights c by S (see the cBOP model in Table 1); this is equivalent to assuming that norm clarity rises with the amount of source information in the environment. This cBOP model is able to capture the "Cialdini checkmark" fairly well.

It would be unreasonable to count this as strong evidence for BOP; clearly, the BOP framework does not offer any a priori reasons for this adjustment. Rather, cBOP demonstrates that the BOP framework can accommodate both Cialdini's nonmonotonic checkmark and his arguments about why it occurs. Nevertheless, in analyses not shown, I found that cBOP did a poor job of fitting the Milgram, Gerard, and Asch data sets. It could be that there is more genuine indecision in Cialdini's littering context; also, one's interpretation of the presence or absence of litter will be influenced by beliefs about whether someone recently cleaned the area.

The Helping Paradigm

A major contribution of Latané's SIT is its distinction between the multiplication of impact (of multiple sources on a given target) versus the division of impact (of a given source across multiple targets). Latané and Darley (1970) are justly famous for their demonstrations that the probability that a bystander will help in an emergency falls off with the number of bystanders, and Latané (1981) later showed that this "diffusion of responsibility" effect is directly predicted by the source-to-target ratio in SIT. I compared the fit of various models to two key data sets in this literature that are especially useful because they parametrically varied the number of bystanders. Experimenters working for Latané and Dabbs (1975) "accidentally" dropped handfuls of coins or pencils in the presence of one to six bystanders in elevators in three American cities. Wiesenthal, Austrom, and Silverman (1983) solicited charitable donations from groups of one to five people in pubs and in classrooms. Figure 5 and Table 5 show that SIT, uBOP, and bBOP provide nearly identical fit to data showing diminishing likelihood that a bystander will help as a function of group size. OTR and SIM (not plotted) fared more poorly, as did the KIM

Table 3
Parameter Estimates for Each Model

Study	SIT		OTR	SIM	KIM			uBOP			bBOP		
	m	*x*	*m*	*m*	*m*	β	α	*m*	*c*	*b*	*m*	*c*	*b*
Conformity													
Milgram et al. (1969)	0.464	0.254	0.896	0.836	1.000	0.746	0.980	0.819	1.208	1.431	1.000	5.329	0.610
Gerard et al. (1968)	0.162	0.375	0.351	0.316	0.489	0.526	1.000	0.305	1.528	1.406	0.368	6.580	0.603
Asch (1956)	0.154	0.318	0.335	0.325	0.475	0.492	1.000	0.327	3.963	2.096	0.325	107.253	0.670
Helping													
Latané & Dabbs (1975)	0.498	0.386	1.000	1.000	1.000	1.000	0.727	0.482	5.772	0.180	0.502	8.978	0.163
Wiesenthal et al. (1983)	0.776	0.757	1.000	1.000	1.000	0.621	0.138	1.000	2.996	0.573	1.000	7.179	0.343
Deliberation													
Hinsz (1990)	0.226	0.615	0.600	0.618	N/A	0.378	0.878	N/A	0.515	3.931	N/A	4.732	0.758
Tindale (1993)	0.647	0.106	1.000	1.000	N/A	0.516	0.000	N/A	0.436	0.254	N/A	1.971	0.195
Kerr & MacCoun (1985)													
p(guilty verdict)	0.305	0.584	0.912	0.983	N/A	0.235	0.667	N/A	2.610	1.745	N/A	18.415	0.624
p(not guilty verdict)	0.589	0.468	1.000	1.000	N/A	0.234	0.333	N/A	7.800	0.615	N/A	18.416	0.376

Note. SIT = social impact theory; OTR = other-total ratio; SIM = social influence model; KIM = Kerr influence model; uBOP = unidirectional influence burden of proof model; bBOP = bidirectional influence burden of proof model; *m* = ceiling parameter; *x* = SIT exponent; β = KIM shape parameter; α = KIM threshold parameter; *c* = norm clarity parameter; *b* = threshold parameter; N/A = not applicable.

threshold model, though freeing SIM's −4 constant significantly improves its fit. The uBOP and bBOP models appear to match the fit of SIT by incorporating a similar "division of impact"; it is not clear that the threshold itself plays a role here.

Thresholds may matter more in other types of helping situations, for example, when a discrete public good will be provided only if a "minimal contributing set" of contributors is reached (van de Kragt, Orbell, & Dawes, 1983). A model for predicting behavior in such situations may require a single-peaked goal function, a point that is addressed in the Discussion section.

The Deliberation Paradigm

A third major domain of social influence research can be called the *deliberation paradigm*, in which proponents of multiple viewpoints interact and attempt to influence each other.[10] This paradigm includes analyses of jury behavior, committee behavior, legislative voting, and so forth and has produced a host of prescriptive and descriptive formal models. Many of these models (especially in the public choice literature) introduce various complications—institutional constraints, differences in power and resources, and decisions involving multidimensional outcomes or options that can be bundled or traded off. In contrast, social psychological models have tended to focus on more minimal group contexts, generally involving actors without prior history and with little or no differences in formal authority or resources.

Much of this work involves the large family of models called social combination schemes (Lorge & Solomon, 1955; Smoke & Zajonc, 1962) or social decision schemes (Davis, 1973; Kerr et al., 1996; Stasser et al., 1989). In the social decision scheme tradition, it has been quite useful to use a family of theoretical decision schemes as benchmarks against which to compare actual group performance data. Table 6 shows that bBOP reproduces three of the five most common benchmarks perfectly and closely approximates the other two. This is appealing, because bBOP helps to situate those benchmarks in a two-dimensional ($c \times b$) parameter space. Notice that the asymmetric "truth wins" and "truth supported wins" schemes require higher c values. This supports the interpretation of c as a "norm clarity" parameter, because these schemes come closest to matching actual group behavior in situations where the group has some shared conceptual system for identifying a "correct" response (Kerr et al., 1996; Laughlin, 1999).

Interestingly, bBOP only approximates the two-thirds majority decision scheme. The logistic function is slightly regressive relative to strictly linear models like OTR or the proportionality decision scheme; hence, bBOP allows some small probability that an initially unanimous group will reverse itself. This idea seems more plausible for "evidence-driven" groups (which delay voting until after discussing the evidence) than for "verdict-driven" groups (which take an immediate vote before deliberating; see Hastie, Penrod, & Pennington, 1983), but in any case, the Proportionality D is more of a baseline than a good predictive model. Note that the two-thirds scheme requires a low clarity level (c = 10 vs. 100 for a simple majority) because the symmetrical threshold (b = .5) falls halfway between the .33 or .66 points that would constitute a victory for either side.

Figure 6 and Table 7 provide details on the fit of all five models to three data sets from the group deliberation paradigm. The first study (Hinsz, 1990) involves the accuracy of group performance on a recognition memory task. There is a pronounced asymmetry due to the shared recognition of a correctly recalled answer once it gets articulated by a group member in a memory task; thus, advocates of the wrong answer consistently show less than proportional influence. Because the inflection pattern is more gradual, all the models fit it easily, though none exceed bBOP's perfect fit.

The second study (Tindale, 1993) involves performance on "Linda" style conjunction problems (Tversky & Kahneman, 1983),

[10] Intermediate between these paradigms are cases where factions exchange information without deliberating, or "bandwagon" cases where the mere knowledge of relative faction size influences opinions (see the agent-based simulations below).

Table 4
Fit Statistics for the Conformity Paradigm

Study	Model	RMSE	R^2	Adj. R^2	AIC
Milgram et al. (1969)	SIT	.051	.994	.993	10.779
	OTR	.045	.995	.995	**8.767**
	SIM	.094	.979	.979	9.045
	KIM	**.032**	**.998**	**.996**	12.729
	uBOP	.077	.986	.979	13.090
	bBOP	.041	.996	.994	12.823
Gerard et al. (1968)	SIT	.038	.978	.974	11.906
	OTR	.038	.978	.978	**9.908**
	SIM	.038	.977	.977	9.941
	KIM	.035	.981	.974	13.895
	uBOP	.035	.981	.973	13.958
	bBOP	**.033**	**.983**	**.977**	13.903
Asch (1956)	SIT	.081	.895	.874	10.416
	OTR	.075	.910	.910	8.393
	SIM	.039	.976	.976	**8.185**
	KIM	.070	.923	.885	12.341
	uBOP	**.014**	**.997**	**.995**	12.155
	bBOP	**.014**	**.997**	**.995**	12.976
Average fit	SIT	.057	.956	.947	11.033
	OTR	.053	.961	.961	**9.023**
	SIM	.057	.977	.977	9.057
	KIM	.045	.967	.952	12.988
	uBOP	.042	.988	.983	13.068
	bBOP	**.029**	**.992**	**.989**	13.234

Note. Best fit statistics are in bold font. SIT = social influence theory; OTR = other-total ratio; SIM = social influence model; KIM = Kerr influence model; uBOP = unidirectional influence burden of proof model; bBOP = bidirectional influence burden of proof model; RMSE = root-mean-squared error; AIC = Akaike information criterion.

a domain where people often respond as if they believed p(A & B) $>$ p(B)—a logical impossibility. Here, there is a pronounced asymmetry, such that advocates of the error have greater than proportional influence—a phenomenon Tindale (1993) called “bias supported wins” (also see Kerr et al., 1996). This data set is unique (among those examined here) in that groups often made the conjunction error even when none of their members had done so individually. The BOP models are able to capture this feature; in bBOP, nonzero influence at $S = 0$ implies low clarity.

The third data set is the pattern of unanimous verdicts for 12-person criminal mock juries in Kerr and MacCoun (1985). It shows a clear asymmetry, such that advocates of acquittal are more influential than their proportionally sized counterparts favoring a guilty verdict. Norm clarity (c) is higher here than in the nonlegal cases, perhaps because jurors were given an explicit standard of proof, decision rule, and other instructions (see below). In Table 3, I also show the estimated coefficients when the models are fit to the acquittal rate rather than the conviction rate. Note that the parameter values for the nonthreshold models (SIT, OTR, and SIM) change in seemingly arbitrary ways. In contrast, the parameter changes for the threshold models (KIM, uBOP, and bBOP) are readily interpretable; for KIM, $.33 = 1 - .67$; for bBOP, $.38 = 1 - .62$; and for uBOP, $.38 \approx 1 - .64$ after translating from odds to probabilities.

The asymmetry in mock juries is not an artifact of leniency in student populations; MacCoun and Kerr (1988, Experiment 1) replicated the effect with jury-eligible Michigan adults. Nor is it an inherent property of jury behavior. Table 8 shows the results of three different experiments showing that the asymmetry is a direct result of the asymmetrical “beyond a reasonable doubt” standard used in criminal trials. When mock juries are instructed to use the symmetrical “preponderance of evidence” standard, the asymmetry disappears. Interestingly, traditional “reasonable doubt” instructions have little effect on individual jurors’ predeliberation judgments (Kagehiro & Stanton, 1985; MacCoun & Kerr, 1988; cf. Kerr et al., 1976). This indicates that the standard of proof might work more like a *social* mechanism than a decision theoretic one. MacCoun and Kerr (1988; Kerr & MacCoun, in press) argued that the reasonable doubt criterion gives a rhetorical advantage to advocates of acquittal. If jurors favoring conviction discover that well-meaning peers feel strongly that the defendant is innocent, this might serve as “social proof” that there is, in fact, a reasonable doubt.

This is of more general relevance to the present theory. The data in Table 8 provide direct evidence in support of the “burden-of-proof” interpretation of the BOP models. In three different experiments, mock criminal juries have been randomly assigned to evaluate cases under either a strict “beyond a reasonable doubt” standard of proof (which is asymmetric and favors the defendant) or a less stringent alternative—either lax or undefined (Kerr et al., 1976) or the symmetrical “mere preponderance of the evidence” standard more typically used in civil cases (MacCoun, 1984; MacCoun & Kerr, 1988). Table 8 shows the probability of acquittal given a unanimous verdict for mock jurors that were evenly split (50% for guilty, 50% for not guilty) at the onset of deliberation. The cell sizes are largest for the third study, where evenly split groups were assembled by design. In all three cases, there is a significant asymmetry when the assigned standard of proof is asymmetrical, and the outcomes are symmetrical when the assigned standard is more symmetrical. (The asymmetries are more pronounced in Table 8 than in the deliberation threshold estimates in Table 3, because the latter includes groups that initially favored guilty.)

Interestingly, MacCoun (1984) found that assigned standards of proof did not have their expected effect on *individual* predeliberation verdict preferences. Similarly, Kagehiro and Stanton (1985) failed to find any effect of traditional verbal standard-of-proof instructions on individual mock jurors' verdicts, although they did find an effect of quantified instructions (e.g., greater than 50% for preponderance of evidence). Thus, the asymmetries shown in Table 8 illustrate that even legal standards of proof operate more like social norms than like epistemological inferences. We may sometimes resolve error tradeoffs through social comparison rather than through private cogitation.

Influence in the Real World and Two Virtual Worlds

The Social Imitation Paradigm

The BOP models were developed to model social psychological phenomena in small group interpersonal settings. But there is value in extending the model toward larger scale collective phenomena, to see whether there is some continuity between the micro and macro levels. Here I examine whether the BOP model might be useful in the study of the diffusion of innovations,[11] a paradigm that is itself closely related to models of fashion (Bikhchandani, Hirshleifer, & Welch, 1992; Miller, McIntyre, & Mantrala, 1993), the spread of contagious diseases (Daley & Gani, 1999), and the evolution of culture (Boyd & Richerson, 1985; Henrich & Boyd, 1998).

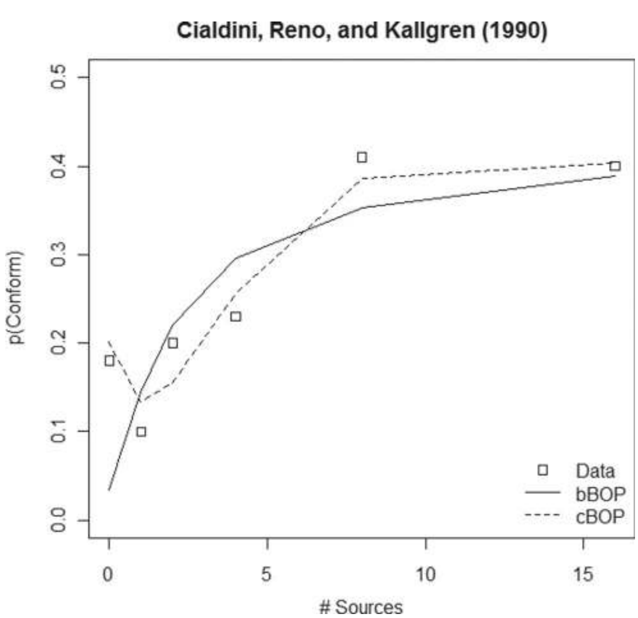

Figure 4. Fitting bBOP and cBOP to the Cialdini et al. (1990) checkmark pattern. bBOP = bidirectional influence burden of proof model; cBOP = contextual clarity burden of proof model.

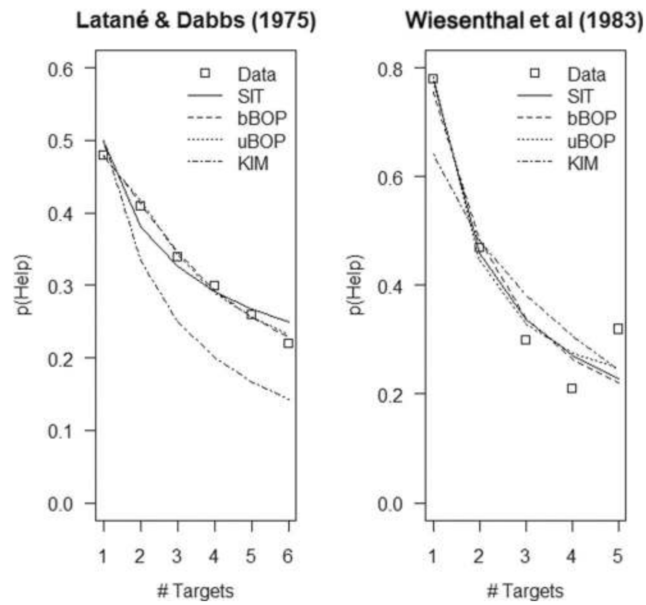

Figure 5. The relative fit of various models for two experiments in the helping paradigm. SIT = social impact theory; bBOP = bidirectional influence burden of proof model; uBOP = unidirectional influence burden of proof model; KIM = Kerr influence model.

Note that in the literature on contagion and diffusion-of-innovation processes, social influence is plotted using time rather than number of sources on the horizontal axis. The iBOP model appears in Table 1; it generates a predicted adoption rate at time t by applying the bBOP model to the targets (nonadopters), with the $t - 1$ adoption rate serving as the source influence.[12] As a point of comparison, I use the influential Bass diffusion model (Bass, 1969; Mahajan, Muller, & Bass, 1995), which appears in the final row of Table 2. It has two key parameters: a coefficient of innovation (p; generally around .03) and a coefficient of imitation (q; generally around .3 to .5 in most marketing applications; Mahajan et al., 1995). Bass interprets p as the probability that an individual would learn about the innovation on his or her own, without social contact; many authors refer to this (somewhat confusingly) as "external influence" (as in "external to the social system") and interpret it as an effect of media and advertising. It seems likely, however, that media attention may increase as a response to the growth in adoption, so the q parameter may be at least as likely to pick up media effects. (This may be why the p parameter is generally so close to zero in most applications of the Bass model.)

I examined four qualitatively distinct case studies. The first was the classic Ryan and Gross (1943) study of the adoption of hybrid seed corn among Iowa farmers. The second series shows the

[11] The diffusion of innovations concept should not be confused with the "diffusion of responsibility" label often used for the helping paradigm; indeed, these phrases use the word diffusion in nearly opposite ways.

[12] A forward-recursion version (Feller, 1957), which used the model's predicted S_{t-1} rather than the observed S_{t-1} as the source index, fit the data sets nearly as well.

Table 5
Fit Statistics for the Helping Paradigm

Study	Model	RMSE	R^2	Adj. R^2	AIC
Latané and Dabbs (1975)	SIT	.020	.997	.996	11.453
	OTR	.081	.965	.965	**9.664**
	SIM	.257	.784	.784	13.045
	KIM	.081	.965	.942	13.664
	uBOP	.007	1.000	.999	13.444
	bBOP	**.006**	**1.000**	**1.000**	13.443
Wiesenthal et al. (1983)	SIT	.0521	.9873	.9830	10.0101
	OTR	.1570	.9816	.9816	**8.5067**
	SIM	.3506	.9380	.9380	12.1194
	KIM	.0909	.9630	.9260	12.1372
	uBOP	**.0470**	**.9896**	**.9793**	11.9960
	bBOP	.0551	.9858	.9715	12.0196
Average fit	SIT	.0359	.9920	.9895	10.7317
	OTR	.1188	.9734	.9734	**9.0855**
	SIM	.3037	.8610	.8610	12.5823
	KIM	.0858	.9641	.9340	12.9008
	uBOP	**.0272**	**.9946**	**.9892**	12.7199
	bBOP	.0303	.9928	.9856	12.7313

Note. Best fit statistics are in bold font. SIT = social impact theory; OTR = other-total ratio; SIM = social influence model; KIM = Kerr influence model; uBOP = unidirectional influence burden of proof model; bBOP = bidirectional influence burden of proof model; RMSE = root-mean-squared error; AIC = Akaike information criterion.

diffusion of television into American households between 1946 and 1951 (De Fleur, 1966). The third series shows the participation of Leipzig citizens in the famous Monday demonstrations prior to the fall of the communist regime in East Germany, which allegedly encompassed most of the city's population by the end of a month (Banta, 1989; Braun, 1995). And the final series depicts the rapid growth in MDMA (Ecstasy) use among American young adults over 12 during the 1990s (Substance Abuse and Mental Health Services Administration, 2002). A challenge in applying influence models in field settings is that it is not always clear what constitutes the relevant reference population, which is necessary to operationalize the source-to-target ratio. For the seed corn case, I indexed the adoption rate to its maximum level in the final year. The television case shows adopters as a percentage of the total number of households. The Leipzig case used weekly counts as a fraction of the estimated 1989 population for the Leipzig metropolitan area. And the Ecstasy case shows cumulative initiation among Americans who were between 18 and 24 years old at some point during the 1984–2000 period.

As seen in Figure 7 and Table 9, the iBOP and Bass models provide a quite strong fit to these data sets. It is clear that most of the explanatory work of the Bass model is being done by its coefficient of imitation (q), but that parameter is open to many different interpretations, and neither of the Bass parameters varies much across cases. And indeed, applications of the Bass model have tended to emphasize forecasting rather than explanation. In contrast, iBOP's parameter estimates vary more across cases; with a larger set of cases, it might be possible to use these parameters to identify situational variations in social imitation dynamics.

That the iBOP model fits these data so well is encouraging, but the results must be interpreted cautiously. First, the number of people involved is far more heterogeneous, and orders of magnitude greater, than would be observed in any psychology experiment. In none of these cases were polling data widely available to all participants, and no participant was likely to have a firm estimate of the level of consensus in the relevant population. At best, people would have to infer consensus information by sampling their own local social network of acquaintances, combined perhaps with mass media images (see Dawes & Mulford, 1996; Krueger & Clement, 1994; Prentice & Miller, 1996). In principle, one might fit iBOP to people's *perceived* prevalence estimates, if they were measured by survey or else modeled (e.g., using Bayesian updating). Second, the iBOP model does not explicitly incorporate nonsocial factors (changes in opportunity, price, perceived risks and rewards) that presumably influence adoption decisions, and these will potentially bias the parameter estimates if not incorporated into the analysis. And third, as with any aggregate estimates, iBOP's threshold and clarity parameters might accurately capture the population at large without distinguishing the characteristics of subgroups within it. Recent work in diffusion theory and epidemiology has sought to incorporate heterogeneity and social structure into traditional models (Caulkins, Behrens, Knoll, Tragler, & Zuba, 2004; Strang & Soule, 1998). When data will permit it, fitting iBOP separately for different subgroups should yield greater insight into the individual and social dynamics of the adoption process.[13]

To illustrate the latter point, the Ecstasy (MDMA) initiation rates plotted in Figure 7 describe the successive waves of 18- to 24-year-olds during the 1984–2000 period (collectively, all those between the ages of 18 and 41 by 2001), with parameter estimates of b = 0.673 and c = 6.780). But if we expanded the relevant population to all Americans aged 12 and over in those years, the iBOP estimates would suggest an extremely steep threshold (b =

[13] Alternatively, one might adapt a mixed Rasch type of model (see Meiser, Hein-Eggers, Rompe, & Rudinger, 1995) for use in the BOP framework.

Table 6
Fitting bBOP to Five Major Social Decision Schemes

Initial split		Proportionality		Simple majority		Two-thirds majority		Truth wins[a]		Truth supported wins[a]	
S	*T*	D	BOP	D	BOP	D	BOP	D	BOP	D	BOP
0	12	0.00	0.08	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1	11	0.08	0.11	0.00	0.00	0.00	0.01	1.00	1.00	0.00	0.00
2	10	0.17	0.16	0.00	0.00	0.00	0.02	1.00	1.00	1.00	1.00
3	9	0.25	0.22	0.00	0.00	0.00	0.04	1.00	1.00	1.00	1.00
4	8	0.33	0.30	0.00	0.00	0.00	0.11	1.00	1.00	1.00	1.00
5	7	0.42	0.40	0.00	0.00	0.42	0.26	1.00	1.00	1.00	1.00
6	6	0.50	0.50	0.50	0.50	0.50	0.50	1.00	1.00	1.00	1.00
7	5	0.58	0.60	1.00	1.00	0.58	0.74	1.00	1.00	1.00	1.00
8	4	0.67	0.70	1.00	1.00	1.00	0.89	1.00	1.00	1.00	1.00
9	3	0.75	0.78	1.00	1.00	1.00	0.96	1.00	1.00	1.00	1.00
10	2	0.83	0.84	1.00	1.00	1.00	0.98	1.00	1.00	1.00	1.00
11	1	0.92	0.89	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.00
12	0	1.00	0.92	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Pearson's *r*			.99		1.00		.95		1.00		1.00
Parameters		*b*	0.5	*b*	0.5	*b*	0.5	*b*	0.05	*b*	0.125
		c	5	*c*	100	*c*	10	*c*	200	*c*	200
		m	1	*m*	1	*m*	1	*m*	1	*m*	1

Note. bBOP = bidirectional influence burden of proof model; *S* = source; *T* = target; D = predicted decision probability; BOP = burden of proof; *b* = threshold parameter; *c* = norm clarity parameter; *m* = ceiling parameter.
[a] Assumes that sources are advocating "truth" relative to shared conceptual scheme.

0.990, with c = 6.287), because so few Americans other than young adults tried Ecstasy. At the other extreme, if we were to identify the most highly susceptible groups of young adults on the basis of relevant risk factors, we would expect an even shallower threshold than 0.673. Another example of the effects of mixing populations is explored below (for the simulation "Scenario 10:90").

Agent-Based Simulations of the Implications of the Model

Social influence processes are distributed across space as well as time, but it is impractical to examine these dimensions in any broad way using social psychology experiments. A meta-analysis by Bond (2005) identified 125 Asch-type conformity studies, which works out to roughly 2.5 studies a year since Asch. But very few had examined more than three different faction sizes, and almost all the experiments examined responses to a small number of influence sources encountered for a brief period in a single location.

I therefore explored some of the implications of the bBOP model using cellular automata computer simulations. In recent years, many investigators have used such models to explore aggregate social phenomena that are difficult to analyze in the field or simulate in the lab (Axelrod, 1997; Epstein & Axtell, 1996; Nowak et al., 1990). These models represent social systems as a grid of cells occupied by simple "agents" who respond to their neighbors and their environment according to simple behavioral rules. This grid can be construed as a model of geographical space (e.g., adjacent houses in neighborhoods) or as a more abstract model of social space (in terms of close vs. remote social contacts). Each agent's behavior is updated iteratively in response to changes in this environment to identify aggregate consequences of individual behavior.

My simulations were run in NetLogo (Wilensky, 1999), adapting and significantly modifying code from the NetLogo segregation model (Wilensky, 1998). In the simulations, agents are either blue or red. The program first randomly seeds up to 2,500 agents in a 50 × 50 grid. In the simulations presented here, I compare results at *high density* (2,500 agents, or 100% occupation) versus *low density* (1,250 agents, or 50% occupation). I also compare two different topologies. In the *torus* topology, the grid wraps around itself so that agents at the right border of the grid are immediately adjacent to those at the left border; similarly, those at the bottom border are immediately adjacent to those at the top border. In the *box* topology, the four borders of the grid are firm barriers so that agents near the borders are more isolated than those near the center.

For each iteration, the program calculated for each agent how many of its neighbors are similar in color. In the simulations presented here, a key parameter was *vision*, the radius within which neighbors can be observed. In the "Moore" neighborhoods simulated here, each agent can observe neighbors along eight directions (corresponding to the compass directions N, NE, E, SE, S, SW, W, and NW) at a radius of from 1 to 10. A vision of 1 corresponds to $(2 * \text{radius} + 1)^2 - 1 = 8$ neighbors; a vision of 2 corresponds to 24 neighbors; and a vision of 10 corresponds to 440 neighbors. This vision parameter allows an examination of the effects of local versus global majorities. When vision is at higher levels, individual behavior change becomes more correlated across agents.

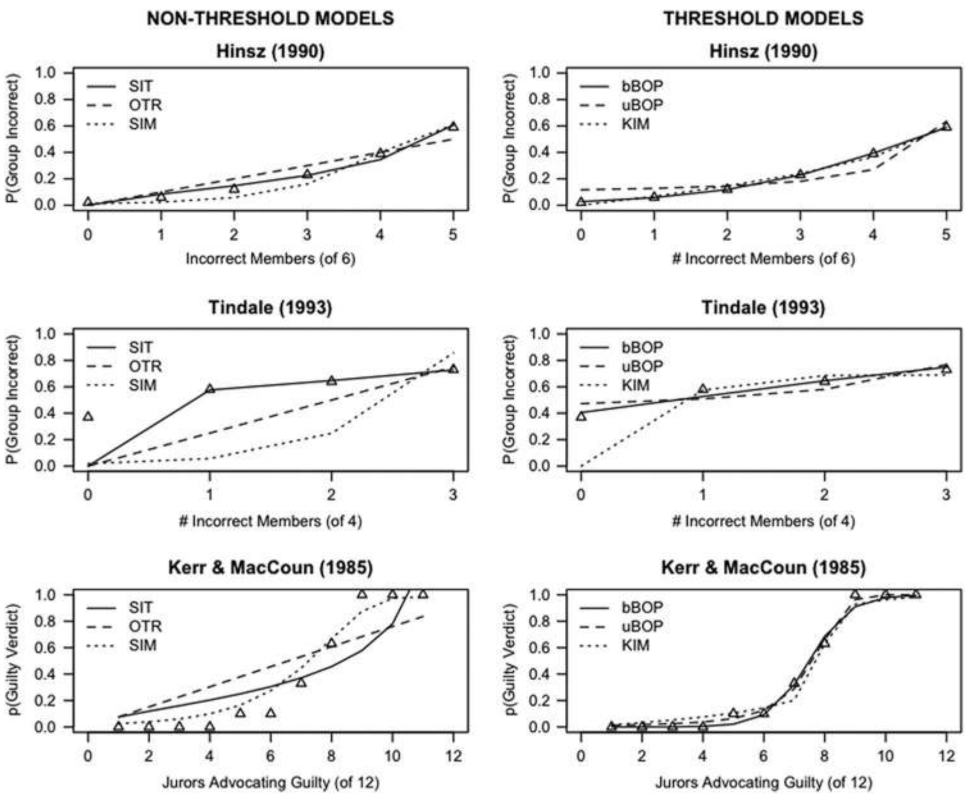

Figure 6. The relative fit of nonthreshold models (left panels) and threshold models (right panels) for three experiments in the deliberation paradigm. SIT = social influence theory; OTR = other-total ratio; SIM = social influence model; bBOP = bidirectional influence burden of proof model; uBOP = unidirectional influence burden of proof model; KIM = Kerr influence model.

The program then computed the probability that the agent will change colors, using one of two processes. In the *constant threshold* process, all agents of the same color are assigned a common threshold value (b) and clarity value (c). The probability of change was then computed using the bBOP equation. For comparison purposes, the *sampled threshold* process is intended to correspond more closely to a Granovetter-type threshold model, in which each agent's threshold value is sampled from a normal distribution, where the mean and standard deviation are fixed for each agent color. The mean threshold value matched the value used for the constant threshold process. The standard deviation for the sampled threshold process was calibrated to the fixed clarity value of the constant threshold process: $SD = 1.8/c$.[14]

The agent's probability of change was compared to a random number below 1 to determine whether the agent would change colors. (In the version used here, this change occurs only if the agent is prompted to change by having at least one dissimilar neighbor; if all neighbors are similar then no change occurs. Experiments allowing unprompted change despite completely similar neighbors produced nearly identical results.) The updating across agents was done in a synchronous manner, so that updated values were not applied to neighbors until a round of updating was complete.[15]

I describe two different "experiments" using this modeling approach, but it is important to emphasize that these simulation "results" are implications of the model rather than empirical findings that might validate the model. With that caveat in mind, my simulation data confirm (sometimes dramatically) a core implication of earlier cellular automata models of social behavior: Even when members of different opinion factions are quite dispersed across social space at the beginning of a simulation, after relatively few iterations there is a great deal of clustering, such that most agents are surrounded by "like-minded" others—indeed, some become quite isolated from members of the other factions. Because this phenomenon is well described elsewhere (Axelrod, 1997; Epstein & Axtell, 1996; Nowak et al., 1990), it is not documented here, but it was readily apparent in these simulations.

[14] This calibration was determined inductively after simulations testing a range of conversion values from $1/c$ to $1.8/c$. It is close to the $1.68/c$ conversion suggested by the simulations in Appendix A, and both are within Long's (1997, p. 48) suggested 1.6 to 1.8 range for converting a probit to the corresponding logit.

[15] The agent-based simulations examined in this article allowed agents to change opinions upon contact with other agents but did not allow them to move to another location. A model variant in which agents cannot change their opinion but can move to another location can produce strikingly similar clustering patterns (MacCoun, Cook, Muschkin, & Vigdor, 2008).

Scenario 50:50. In the first set of simulations, I examined a scenario involving equal numbers of red and blue agents (50:50 ratio). For both groups, the threshold value (b) was set at 0.50, with clarity (c) set at 50. This clarity level is intermediate between that in the simple majority decision scheme ($c = 10$) and that in the two-thirds majority decision scheme ($c = 100$) presented earlier.

This set of simulations had a 2 (torus vs. box topology) × 2 (50% vs. 100% density) × 2 (constant vs. sampled thresholds) × 3 (vision: 1, 5, or 10) factorial design, with 20 runs per cell, producing 480 outcomes, where each outcome was the final percentage of blue agents after 50 iterations of the entire network. Despite this large sample size, an analysis of variance (ANOVA) revealed no significant effects (Fs < 2.9). Methodologically, this provides evidence that the BOP model's constant threshold assumption closely mimics a Granovetter-type sampled threshold process. This is unsurprising, as it is known (e.g., Long, 1997) that a logistic model with appropriate parameters can closely approximate cumulative normal probit models that are based on the normal distribution. I also found no differences between a torus and box topology, though it is possible that topology matters more for other combinations of parameter values. Substantively, the lack of mean differences for density and vision is unsurprising because there is little reason to expect that either parameter would produce an asymmetry when the simulations started with a symmetrical 50:50 initial ratio of agents with identical threshold and clarity levels.

Table 7
Fit Statistics for the Deliberation Paradigm

Study	Model	RMSE	R^2	Adj. R^2	AIC
Hinsz (1990)	SIT	.026	.993	.991	9.196
	OTR	.060	.962	.962	7.306
	SIM	.042	.981	.981	**7.235**
	KIM	.017	.997	.996	9.180
	uBOP	.075	.942	.928	9.457
	bBOP	**.004**	**1.000**	**1.000**	9.167
Tindale (1993)	SIT	.185	.903	.855	11.256
	OTR	.258	.828	.828	9.783
	SIM	.377	.600	.600	11.117
	KIM	.204	.889	.833	11.447
	uBOP	.072	.986	.985	9.241
	bBOP	**.033**	**.997**	**.995**	**9.173**
Kerr & MacCoun (1985)	SIT	.206	.868	.854	11.275
	OTR	.234	.829	.829	10.548
	SIM	.086	.977	.977	**7.051**
	KIM	.072	.984	.982	8.924
	uBOP	**.028**	**.998**	**.997**	8.243
	bBOP	.041	.995	.995	8.284
Average fit	SIT	.139	.921	.900	10.576
	OTR	.184	.873	.873	9.212
	SIM	.168	.853	.853	**8.468**
	KIM	.098	.957	.937	9.850
	uBOP	.058	.975	.968	8.980
	bBOP	**.026**	**.997**	**.997**	8.875

Note. Best fit statistics are in bold font. SIT = social influence theory; OTR = other-total ratio; SIM = social influence model; KIM = Kerr influence model; uBOP = unidirectional influence burden of proof model; bBOP = bidirectional influence burden of proof model; RMSE = root-mean-squared error; AIC = Akaike information criterion.

Table 8
Effects of Assigned Standards of Proof on Probability of Acquittal in Evenly Split Groups

	Outcomes for groups starting with equal faction sizes			
Study	Convict	Acquit	Hung	$P(\mathrm{A}|\mathrm{V})$
Kerr et al. (1976)				
Stringent reasonable doubt	0	3	5	1
Lax or undefined reasonable doubt	5	5	4	.5
MacCoun (1984)				
Reasonable doubt	1	5	3	.83
Preponderance of evidence	2	2	1	.5
MacCoun & Kerr (1988)				
Reasonable doubt	10	28	9	.74
Preponderance of evidence	15	21	6	.58

Note. $P(\mathrm{A}|\mathrm{V})$ = probability of acquittal given unanimous verdict.

But an examination of the mean results obscures the most interesting result of the experiment, which is the effect that the vision manipulation had on the variance in outcomes. As seen in Figure 8, at low vision (a simple Moore neighborhood of 8 neighbors), almost every run ended with the initial 50:50 distribution of blue and red agents unchanged. But as vision increased and the number of visible agents increased, nearly half of the runs ended with a unanimous state in which all agents were the same color.[16] This occurs because at high vision, behavior change becomes highly correlated across agents (a form of spatial and temporal autocorrelation), and slight differences in the initially random allocation of agents across space allow one color to prevail. This is reminiscent of the sort of "extreme sensitivity to initial conditions" invoked in "chaos" models, but it was produced here using a model that is stochastic rather than deterministic. The result is also analogous to the emergence of social coordination norms as documented in the game theory literature (Arthur, 1994; Kelley et al., 2003; Luce & Raiffa, 1957).

Scenario 10:90. In the second set of simulations, I examined a scenario involving an initial minority of 10% blue agents confronting a majority of 90% red agents. For red agents, the threshold value is set at 50, with clarity set at a relatively low 5. The minority agents had threshold values of either 50 or 75, with clarity of either 5 or 10. In this scenario, the topology was always a torus with 50% density (1,250 agents). This set of simulations had a 2 (blue threshold) × 2 (blue clarity) × 2 (constant vs. sampled thresh-

[16] Intuitively, one might expect bimodal distributions (bottom) to indicate a highly polarized society and the single-peaked distributions (top) to represent a more conformist society. But these are distributions of societies (i.e., simulation runs) rather than of actors within a society, and so the reverse is true. In the unimodal distribution, a scenario's population remains divided 50:50, whereas in the bimodal distributions, most simulation runs end with everyone red or everyone blue.

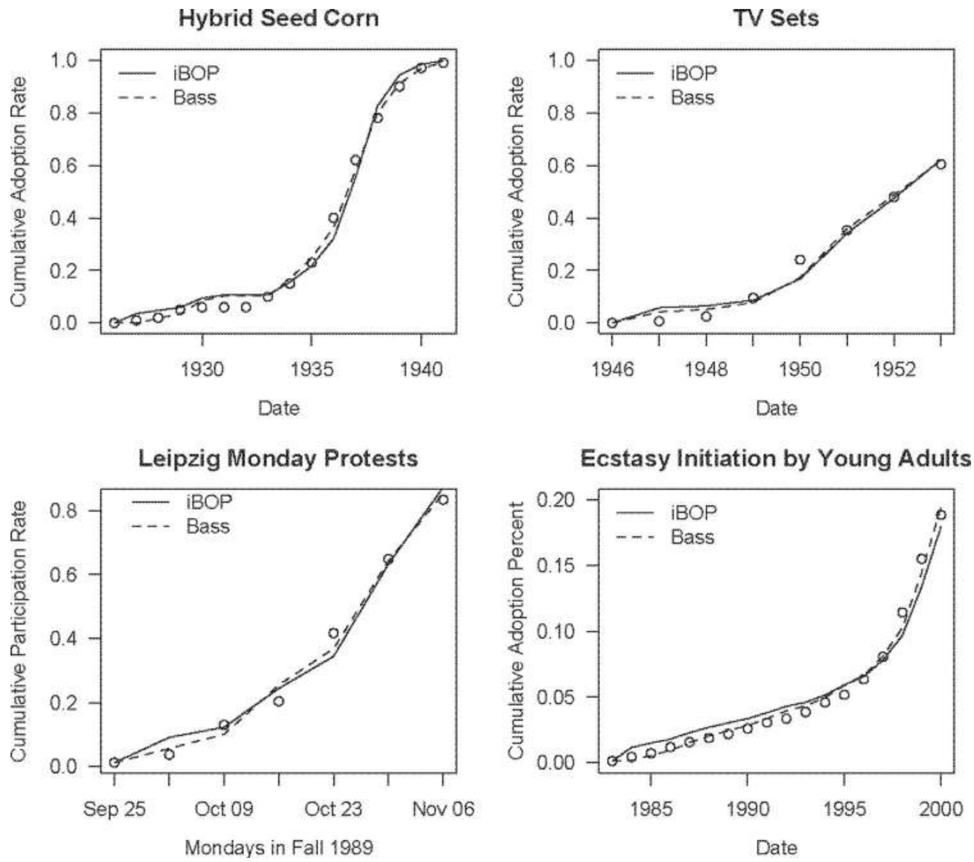

Figure 7. The relative fit of two models for four case studies in the social imitation paradigm. iBOP = social imitation burden of proof model; Bass = Bass diffusion model.

olds) × 2 (vision: 1 vs. 10) factorial design, with 20 runs per cell, producing 320 outcomes based on 50 iterations of the entire network for each run.

In contrast to Scenario 50:50, almost every effect in the 2 × 2 × 2 × 2 ANOVA was significant; $F(1, 303) = 20.92$, $p < .0001$, for the four-way interaction. As seen in Figure 9, the constant and sampled thresholds again produced extremely similar effects. When vision was low (1, corresponding to an average of 4 neighbors), the minority faction was most successful when they shared a high threshold with high clarity. When this happened, the blue faction was able to grow from an initial 10% to over 60% after 50 iterations. This is reminiscent of the empirical finding that minority factions in small group experiments are significantly more influential when they behave in a very consistent and persistent manner (Moscovici, Lage, & Naffrechoux, 1969; Nemeth, 1986; Wood, Lundgren, Oullette, Busceme, & Blackstone, 1994).

But when vision was high (10, corresponding to an average of 220 neighbors), even the lower clarity level was sufficient for the minority faction to completely prevail, so long as the blue threshold (75) exceeded the red threshold (50). In essence, a faction's steep threshold is advantageous only if minority members can find each other and if majority members are aware of the minority viewpoints (see Sunstein, 2003).

Table 9
Fit Statistics for the Social Imitation Paradigm

	Bass diffusion model			iBOP		
Case study	p	q	Adj. R^2	c	b	Adj. R^2
Diffusion of hybrid seed corn	0.000	0.751	.996	5.549	0.592	.990
Adoption of TV sets	0.042	0.446	.978	3.779	0.739	.969
Participation in 1989 Leipzig Monday Protests	0.035	0.844	.987	4.728	0.529	.977
Initiation to MDMA (Ecstasy) use	0.000	0.300	.990	6.780	0.673	.987

Note. In the Bass model, p is the coefficient of innovation and q is the coefficient of imitation. See text for data sources. iBOP = social imitation burden of proof model; c = norm clarity parameter; b = threshold parameter.

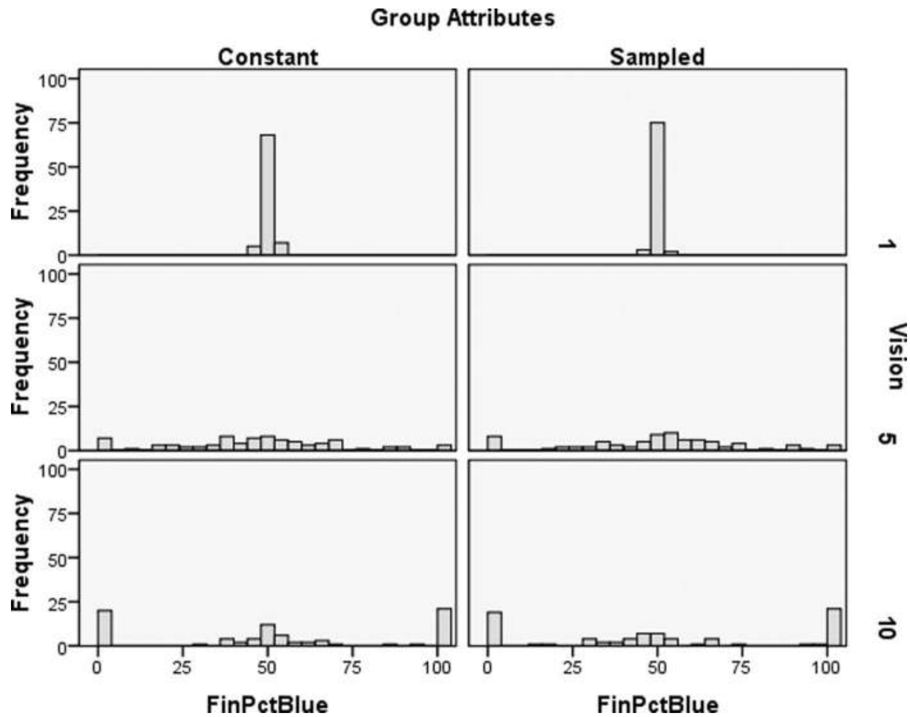

Figure 8. Average final distribution in the 50:50 scenario simulations. Data points describe the distribution of "societies" across simulation runs, rather than the distribution of agents within a society. FinPctBlue = percentage of agents whose final choice was blue (vs. red).

Discussion

What This Theory Does and Doesn't Accomplish

The idea that influence grows with relative prevalence is well established in the work of Solomon Asch, Bibb Latané, James H. Davis, Robert Cialdini, and many other scholars. The idea that such phenomena can be modeled as a threshold process or "tipping point" has been fruitfully explored by Thomas Schelling, Mark Granovetter, and others. What is distinctive about the present account is that I show how a set of closely related formal models can unify much of this past work. First, the BOP models allow a more explicit mapping of social influence processes across several distinct research paradigms: conformity, deliberation, helping behavior, and the diffusion of innovations. I show that logistical threshold models with a threshold parameter and a clarity parameter can successfully fit major data sets from each of these paradigms, often better than alternative models in the literature. And core features of earlier models can be seen as special cases of the BOP model: the knife-edged tipping points of the Schelling segregation model, the distributed social thresholds of Granovetter's model, the concave diminishing marginal influence of Latané's SIT, and the s-shaped curves of the Asch data plotted by faction size and of diffusion-of-innovation data plotted cumulatively over time.

These models also offer a methodological tool for trying to infer something about social situations. Given data on how opinions change with relative faction size, the models allow one to estimate two interesting social parameters: the degree to which one side or the other seems to bear the "burden of social proof" and the degree to which there is consensus about that point. In support of this particular interpretation of the parameters, I have offered evidence from mock jury experiments in which the operative legal standard of proof was manipulated; these experiments suggest that standards of proof sometimes operate in a manner that is more social (how many people disagree with me?) than epistemic (how much evidence do I have?).

It is also important to clarify what this theoretical account does not do. This is hardly a "grand unified theory" of social influence. It is an account, and only a partial account, of dichotomous choices between positions or actions; it says nothing about the content of the attitudes underlying choices, or about the passion (or passive aggressiveness) with which actions get implemented. Further, these models describe the functional relationship between social consensus and individual behavior but say little about the details of the processes that give rise to that functional relationship. As a result, this analysis resists easy classification in terms of classic distinctions like informational versus normative influence (Deutsch & Gerard, 1955), compliance versus identification versus internalization (Kelman, 1958), deliberative versus associative or heuristic processing (e.g., Petty & Cacioppo, 1986), or injunctive versus descriptive social norms (Cialdini et al., 1990). Although the models more readily represent normative influence ("strength in numbers"), outside the laboratory informational influence ("strength in arguments") is often correlated with consensus information. And there are many ways in which these various processes might affect the parameters of the BOP model. In this sense, the

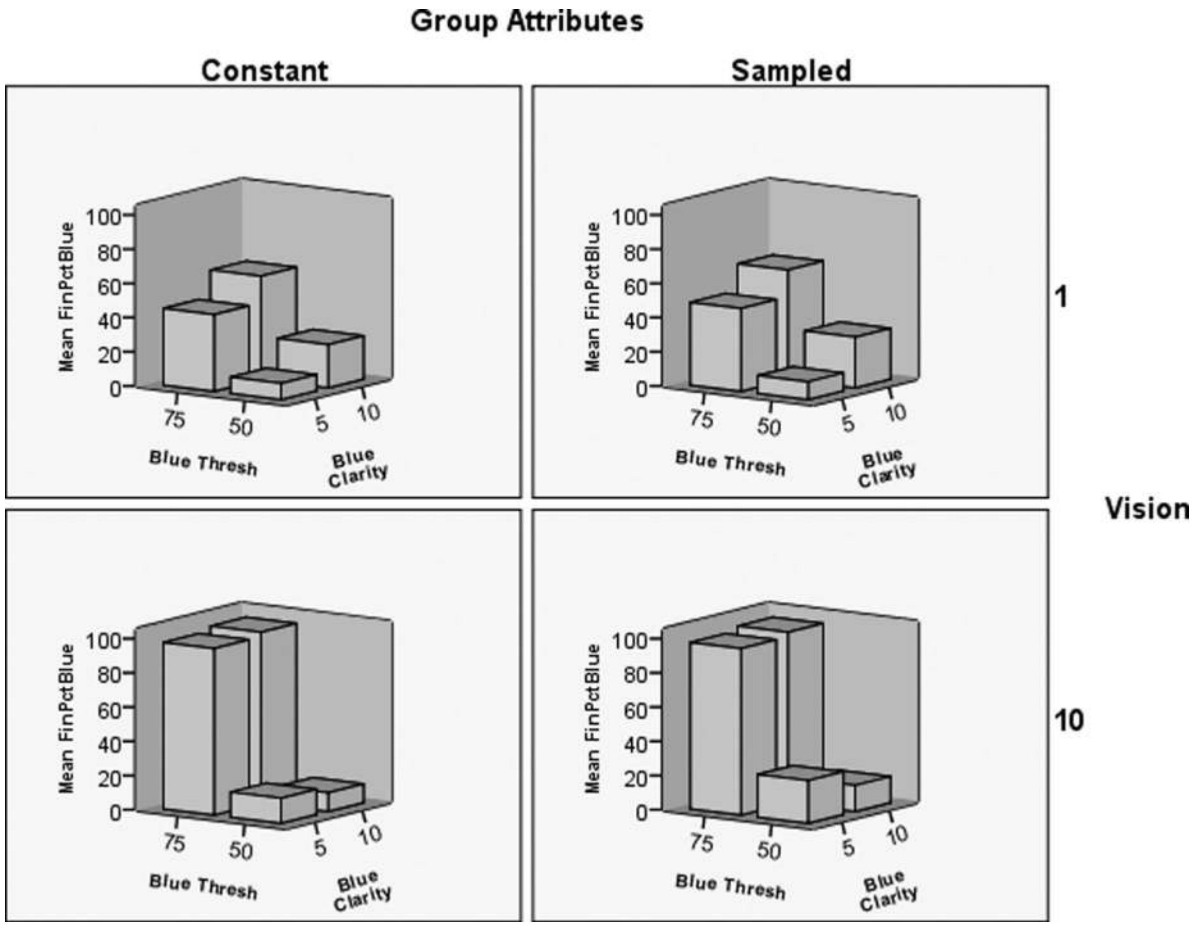

Figure 9. Average final distribution in the 10:90 scenario simulations. Data points describe the distribution of "societies" across simulation runs, rather than the distribution of agents within a society. Thresh = threshold; FinPctBlue = percentage of agents whose final choice was blue (vs. red).

BOP models might play the same kind of "skeletal" role as other formal models in social psychology, like Davis's (1973) social decision scheme theory for small groups or any of the recent connectionist models for social cognition (Smith, 2009). Like these approaches, the BOP model is as much an analytic tool as it is a substantive theory, and to fully understand the particular model parameters, analysts should draw on the rich body of social psychological theory on influence processes, a point discussed in more detail below.

This article provides considerable evidence that is consistent with the hypothesized "matching-to-threshold" process but little direct evidence of that process. Bailenson and Rips (1996) identified appeals to standards and burdens of proof in ordinary argumentation. And as shown above, three different experiments suggest that experimentally manipulated standard-of-proof instructions influence group behavior in the manner predicted by this account. Two of these studies (MacCoun, 1984; MacCoun & Kerr, 1988) also presented evidence that questionnaire measures of standards of proof varied with the manipulations in the expected way. MacCoun (1984) compared a wide variety of measurement approaches to directly assessing standards of proof—including probability measures, odds ratios, and aversion to different decision errors—and showed that the most explicit self-report measures had only partial validity. This may indicate that the matching-to-threshold process is not in fact occurring, but in light of the success of the BOP models, it seems more likely that the thresholds are implicit rather than explicit, with the matching process largely occurring unconsciously. Moreover, the models suggest that standards of proof often operate as social rather than epistemic processes—as standards of *social* proof.

The analyses presented here clearly show the value of incorporating a slope parameter as an index of the degree to which thresholds are shared, but its interpretation as an index of "norm clarity" is still somewhat speculative. It is important to emphasize that it is clarity about the norm and not about the task or the evidence that is being modeled. The Monte Carlo simulations suggest that the c parameter can capture lack of social consensus about the operative threshold *across* people, and additional simulations not reported here show that it can also model noisiness in perceptual and choice processes *within* each person. But further empirical research is needed before the parameter's proper psychological interpretation can be established.

Linking the BOP Parameters to Other Psychological Constructs

With relatively few parameters (ceiling, threshold, clarity, and vision) and variables (*S*, *T*, and *N*), the bBOP model clearly abstracts away much of the nuance and complexity of social influence. The agent-based simulations presented above sampled just two small regions of this parameter space, yet both produced novel behavioral predictions with potential real-world relevance. Unfortunately, testing such predictions in the laboratory is daunting because of the cost and logistical complexity of examining large groups of interacting individuals in real time. One way to manage this challenge is to manipulate social consensus information via experimentally provided polling feedback (e.g., Kerr et al., 1987). This approach maintains a high level of experimental control, but it involves "mere consensus information" rather than actual verbal and nonverbal interaction; whether this is advantageous or disadvantageous will depend on one's research questions. Alternatively, it is now possible to obtain vast quantities of group interaction data via social networking media like Facebook, Twitter, YouTube, and other websites. But in nonexperimental settings it is more difficult to parse "strength in numbers" from "strength in arguments," and the BOP parameters estimated from such data may be biased by spurious factors that covary with consensus information. With these caveats in mind, I offer some hypotheses about how the model parameters may relate to various processes and constructs in the influence literature.

The threshold parameter (*b*). The threshold parameter should covary in predictable ways with any variable known to influence resistance or yielding to persuasion (see generally, Eagly & Chaiken, 1993), and manipulating these factors (or constructing social groupings based on them) should produce corresponding changes in the observed threshold. For example, thresholds should vary directly with stable personal traits like private self-consciousness, dogmatism, and dominance but should vary inversely with traits like public self-consciousness or need for affiliation. Thresholds should be increase with attitude extremity and certainty but decrease with attitude ambivalence. Higher thresholds seem likely when arguments supporting one's position are strong but also when they are deeply embedded in a network of other beliefs, especially those that directly implicate self-identity. Thresholds should fall in response to the "social powers" of the sources, including expertise, legitimacy, prestige, and control over rewards and punishments. Thresholds should also be lower in collectivist cultures, in groups with opportunities for compromise and logrolling, and when tasks involve a normative conceptual scheme establishing a "correct" answer (e.g., arithmetic, engineering) or a shared threat requiring a coordinated response.

The clarity parameter (*c*). The clarity parameter reflects the extent to which a particular threshold level is shared among targets of social influence. Under the random utility interpretation (Appendix A), it also reflects stochastic uncertainty within a single respondent, as might be expected when one's attitudes toward an issue are to some extent assembled "on the fly" (Tourangeau, Rips, & Raskinski, 2000; Zaller, 1992). Thus, while the threshold parameter is expected to vary with the mean level of any of the above influence factors, the clarity parameter can be expected to vary with their *variance* within a target population. Clarity might also be higher when topics and situations are familiar and/or simple than when they are novel and/or complex. Tipping points (very high clarity) might be especially likely when there is a "monoculture" of very few or very redundant arguments favoring one position.

The max parameter (*m*). None of the models examined here can adequately fit the classic conformity sets without invoking a ceiling parameter, yet none of the models (BOP included) offers a compelling theoretical account. What is clear is that *m* indexes a form of resistance persuasion that is empirically distinct from the threshold parameter. It is not possible to mimic the effects of a low ceiling on the conformity data by substituting any combination of the clarity and threshold parameters, even when these parameters are allowed to take any value. My conjecture is that a high ceiling reflects *dependence* as conceptualized by interdependence theorists (Kelley et al., 2003) and that conformity experiments minimize dependence by studying complete strangers in one-shot interactions.

Relative social support (*S/N*, *S/T*). In the analyses reported here, relative faction size was an objectively measured or manipulated feature of the data sets. But of course, the target's subjective perception of social support might not be accurate, and perceived support might be more predictive than actual support. This has two implications, one statistical and the other psychological. Statistically, if the target's perceived consensus is quite discrepant from the objective consensus as measured by the researcher, then threshold estimates fit to the objective *S/N* ratio will be biased indicators of the subjective threshold. Psychologically, if the discrepancy is noticed, it may create a sense of surprise, and expectancy violations are well known to trigger or amplify psychological responses. For example, in cases of "pluralistic ignorance," the target may incorrectly presume that others will hold a different opinion (Prentice & Miller, 1996), so that a relatively low *S/N* ratio will be surprising. This may have played a role in producing Cialdini's "checkmark" pattern (see the discussion of cBOP, above). In cases of social projection (Robbins & Krueger, 2005), or when the manipulated level of opposition is a distortion of the true base rate (Asch, 1956), the target may encounter less support than expected. In either case, changes in opinion or position may be more common than in situations where there is little surprise.

Extending and Generalizing the BOP Framework

Previous work on logistic models suggests useful ways in which the BOP models might be extended. On the dependent variable side, McFadden (2001) and others have extended the logit choice model to the case of three or more response alternatives (see Cramer, 2003; Liao, 1994). On the independent variable side, the logistic framework can readily accommodate other influence factors in addition to the source-to-target ratio. For example, Zaller (1992) operationalized McGuire's (1968) persuasion framework using a logit framework, applying it to an impressive array of political opinion phenomena.

The current article applies the BOP model to variations in social faction size. But the underlying logic applies more generally to threshold processes for other dichotomous choice situations, and it should be possible to substitute other variables in place of *S/T* or *S/N*, for example, argument strength or monetary incentives. By varying two such factors parametrically, it should be possible to directly compare threshold and clarity values for, say, strength in numbers versus strength in arguments.

The models examined in this article (BOP, SIT, SIM, etc.) involve what Coombs (1964) called *dominance* processes, in which responses are a monotonically increasing (or decreasing) function of some input variable. Although the threshold models posit a point of inflection, the function itself does not change direction. But as Coombs argued, many response functions involve a *proximity* process, with a single-pointed peak or "ideal point." Although goal definitions in the literature are not always clear on this point, arguably single-pointed peaks help to distinguish *goals* from ordinary preferences—while we may welcome surpassing a goal, we often diminish our efforts once it is reached, if only because of competing goals in the queue (see Aarts & Elliot, 2012).

The basic BOP model can be adapted to proximity goals; two such functions are offered at the bottom of Table 1. The gBOP model involves a Euclidean distance metric in one dimension and a city block metric if extended to two or more dimensions (see Gärdenfors, 2000). The g2BOP model is a Euclidean metric in two or more dimensions but a squared Euclidean loss function (like mean squared error) in one dimension. As seen in Figure 10, gBOP produces a sharply pointed function reminiscent of goal gradients proposed by Hull (1938) and Heath, Larrick, and Wu (1999). In the group setting, I conjecture that the gBOP function might appropriately characterize responding in step-level public goods, where there is a specific target or "minimal contributing set" (van de Kragt et al., 1983). In contrast, g2BOP suggests a less focused goal, where small discrepancies are more tolerable; g2BOP also provides a good summary characterization of more complex opponent-process situations, where it approximates the geometric mean of a positive and a negative dominance function. Thus, the plotted g2BOP function closely resembles Brewer's (1991, Figure 2) presentation of her optimal distinctiveness model and might characterize a situation in which members welcome some support but prefer not to join the most popular faction. An example might be Arthur's (1994) "El Farol problem," in which patrons want to go to a bar so long as no more than 60 of the available 100 seats are taken.

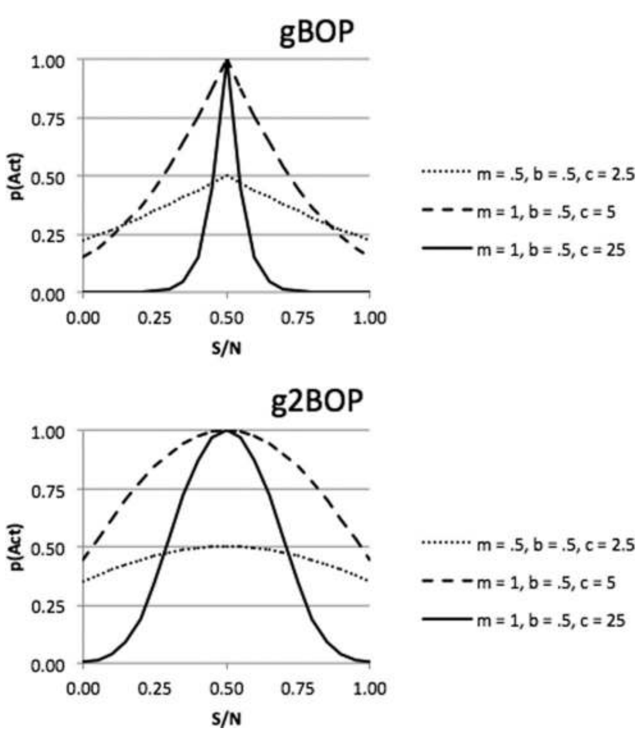

Figure 10. From thresholds to goals: Two burden of proof (BOP) models for single-peaked goals and preferences. gBOP = single-peaked goal pursuit burden of proof model; g2BOP = gBOP with squared Euclidean metric or loss function; m = ceiling parameter; b = threshold parameter; c = norm clarity parameter; S = number of sources; N = number of sources plus number of targets.

Interpreting the BOP Parameter Space

In previous investigations of social influence, formal modelers have sometimes sought to describe general or typical parameter values across a range of situations. In contrast, I suggest that the models examined here are useful in large part because they help us to infer and quantify *differences* across social situations. To illustrate, various threshold and clarity estimates from this article are plotted alongside each other in Figure 11. (Clarity is depicted in natural logs because the parameter has no ceiling and the values were very skewed.) The threshold is relative to the way the relevant topic was framed (e.g., litter vs. don't litter, error vs. no error), which is somewhat arbitrary. Thus Figure 11 transforms b into an asymmetric index ($|b - .5|/.5$) by "folding" the $b \times c$ space at the $b = .5$ axis. Four of the social decision scheme benchmarks—"equiprobability," "proportionality," "simple majority," and "truth (bias) wins"—are plotted for convenience.

There are too few data points (studies and cases) in the figure for a comprehensive meta-analysis of social thresholds. And the diagram must be interpreted cautiously, because some of the estimates (the diamonds) are inferred from uncontrolled archival data sources and may be biased. Still, it seems clear that there is no emergent "Planck's constant" here; these values are all over the map. It appears that social thresholds vary considerably across settings and tasks.

Variation on the horizontal axis suggests that an implicit "simple majority rule" is rare; instead influence appears fairly asymmetric, with one faction implicitly getting "the benefit of the doubt" (also see MacCoun & Kerr, 1988). This is curious because many theoretical accounts have suggested that simple majority schemes are a powerful cultural adaptation (see Hastie & Kameda, 2005; Henrich & Boyd, 1998; Kurzban & Leary, 2001; Rosenwein & Campbell, 1992), but few of the estimates in Figure 11 align near the 50–50 knife edge (viz., 0 on the asymmetry index) that a simple majority heuristic would predict. (Poignantly, the closest match is provided by the Leipzig protests that brought democracy to East Germany.) This distribution is hard to reconcile with the view that simple majorities hold a special psychological status, though it could be an artifact of the selection of data sets.

The clustering in the middle of the vertical axis suggests social thresholds are shared to a degree that makes it reasonable to refer to them as "norms." If people do indeed share a sense of the burden of social proof, it is natural to wonder where that burden might fall in the case of highly contentious and chron-

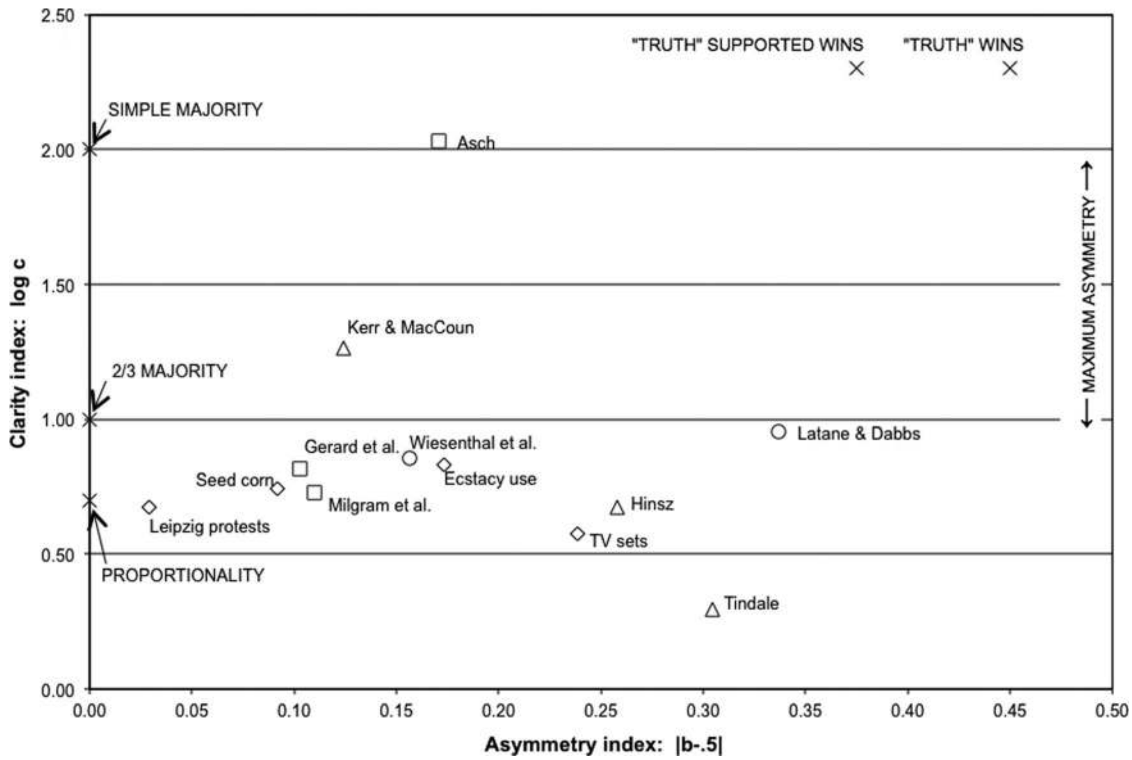

Figure 11. Locating the data in the burden of proof (BOP) parameter space (□ = conformity; ○ = helping; Δ = deliberation; ◇ = social imitation; × = social decision scheme benchmark). b = threshold parameter; c = norm clarity parameter.

ically polarized topics, like abortion, the death penalty, or drug legalization (see MacCoun & Paletz, 2009; MacCoun & Reuter, 2001). Unfortunately, these "culture war" topics were not among those represented in the suitable data sets for testing the BOP models; while hundreds of studies examine attitudes toward these topics, none appear to have parametrically varied "local" consensus and opposition in the manner required for testing the model. It is conceivable that such hot-button topics differ from those represented in Figure 11. Between 1972 and 2006, the General Social Surveys found that 60.6% felt their views toward abortion were "very unlikely to change"; 47.6% said the same thing about the death penalty.[17] Thus, one might expect the distribution of thresholds to be bipolar, with each side holding the other to a nearly impossible burden of proof. This would create high within-group clarity but low aggregate clarity. But this image is a bit of a caricature; even on these extreme issues, most Americans' views are more nuanced than those of the most visible opinion advocates (see Evans, 2003; Fiorina, Abrams, & Pope, 2005; cf. Abramowitz & Saunders, 2005). Also, abstract questions may not capture what happens when people are actually in a room, face to face with others who disagree with them (Asch, 1955; Tetlock, 1992). Through some combination of political, legal, and social psychological processes, it is possible that all sides of some contentious social debates perceive that one side bears a greater burden of social proof (see MacCoun & Reuter, 2001, pp. 323–325).

Vision and the Changing Nature of Social Networks

The model-fitting analyses presented earlier show that BOP's threshold and norm clarity parameters work together to describe susceptibility to social influence. But the agent-based simulations in the final part of the article demonstrate that these factors interact with a third parameter that is less obvious in controlled laboratory experiments: "vision," or the number of others whose views an agent can monitor. In those simulations, as vision broadened, the social influence process changed qualitatively, dramatically so.

It is important to clarify that this is a conditional argument: The simulations show what the model implies *if* actors have this kind of broad "vision." Whether people do, or are capable of doing so, is an open question. It is likely that in ordinary circumstances, distant neighbors receive less weight than immediate neighbors, a claim that is modeled explicitly in SIT (Latané, 1981; Nowak et al., 1990).[18] But as information technologies continue to advance, these circumstances may be changing, and the direction of the change isn't certain. On the one hand, information technologies foster the development of small cultural niche communities (Sunstein, 2003), thereby diminishing the impact of the mainstream

[17] See National Opinion Research Center (2009, Appendix T).

[18] Latané and his colleagues modeled this decay by assuming that impact falls off as an inverse square of distance. Evidence on this point is controversial (Knowles, 1999), and other functions (e.g., exponential decay) are also plausible.

society (C. Anderson, 2006). But at the same time, the kind of broad vision modeled here becomes increasingly feasible due to the near ubiquity of national opinion polling results and the increasing popularity of social networking tools like YouTube, Facebook, MySpace, and Twitter.[19] If so, the simulations suggest that this broadened scope of influence can make societies more volatile. Whether this is good, bad, or both is an open question. Hopefully, these information technologies will also make it easier to track and model how the burdens of social proof emerge and evolve in a mass society.

[19] The Roper iPoll national polling database (http://roperweb.ropercenter.uconn.edu, accessed on July 20, 2009) lists 17,171 polling questions asked between January 1, 2007 and December 31, 2007. Facebook and MySpace each received over 60 million unique U.S. visitors in March 2009 (see Stelter & Arango, 2009).

References

Aarts, H., & Dijksterhuis, A. (2003). The silence of the library: Environment, situational norm, and social behavior. *Journal of Personality and Social Psychology, 84,* 18–28. doi:10.1037/0022-3514.84.1.18

Aarts, H., & Elliot, A. J. (2012). *Goal-directed behavior*. New York, NY: Psychology Press.

Abramowitz, A., & Saunders, K. (2005). Why can't we all just get along? The reality of a polarized America. *The Forum, 3*(2), Article 1. Retrieved from http://www.bepress.com/forum/vol3/iss2/art1

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, 19,* 716–723. doi:10.1109/TAC.1974.1100705

Allison, P. D. (1984). *Event history analysis: Regression for longitudinal event data* (Quantitative Applications in the Social Sciences Sage University Paper No. 46). Thousand Oaks, CA: Sage.

Anderson, C. (2006). *The long tail: Why the future of business is selling less of more*. New York, NY: Hyperion.

Anderson, J. R., Bothell, D., Bryne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111,* 1036–1060. doi:10.1037/0033-295X.111.4.1036

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American Economic Review, 84,* 406–411.

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership and men: Research in human relations* (pp. 177–190). Oxford, England: Carnegie Press.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American, 193*(5), 31–35. doi:10.1038/scientificamerican1155–31

Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs, 70*(Whole No. 416).

Augustin, T. (2005). An approach to combine the logistic threshold model of psychophysics with the Bradley-Terry-Luce models of choice theory. *Journal of Mathematical Psychology, 49,* 70–79. doi:10.1016/j.jmp.2004.11.002

Axelrod, R. (1997). *The complexity of cooperation*. Princeton, NJ: Princeton University Press.

Bailenson, J. N., & Rips, L. J. (1996). Informal reasoning and burden of proof. *Applied Cognitive Psychology, 10,* 3–16. doi:10.1002/(SICI)1099-0720(199611)10:7<3::AID-ACP434>3.0.CO;2-7

Banta, K. W. (1989, November 27). Leipzig: Hotbed of protest. *Time*. Retrieved from http://www.time.com

Bass, F. (1969). A new product growth model for consumer durables. *Management Science, 15,* 215–227. doi:10.1287/mnsc.15.5.215

Berkson, J. (1944). Application of the logistic function to bio-assay. *Journal of the American Statistical Association, 39,* 357–365. doi:10.2307/2280041

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *The Journal of Political Economy, 100,* 992–1026. doi:10.1086/261849

Birnbaum, A. (1968). Some latent trait models and their use in inferring an examinee's ability. In F. M. Lord & M. R. Novick (Eds.), *Statistical theories of mental test scores* (pp. 397–479). Reading, MA: Addison-Wesley.

Bond, R. (2005). Group size and conformity. *Group Processes & Intergroup Relations, 8,* 331–354. doi:10.1177/1368430205056464

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.

Braun, N. (1995). Individual thresholds and social diffusion. *Rationality and Society, 7,* 167–182. doi:10.1177/1043463195007002005

Brewer, M. B. (1991). The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin, 17,* 475–482. doi:10.1177/0146167291175001

Camazine, S., Deneubourg, J. L., Franks, N., Sneyd, J., Bonabeau, E., & Theraulaz, G. (2001). *Self-organization in biological systems*. Princeton, NJ: Princeton University Press.

Campbell, J. D., & Fairey, P. J. (1989). Informational and normative routes to conformity: The effect of faction size as a function of norm extremity and attention to the stimulus. *Journal of Personality and Social Psychology, 57,* 457–468. doi:10.1037/0022-3514.57.3.457

Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. New York, NY: Cambridge University Press.

Caulkins, J. P., Behrens, D. A., Knoll, C., Tragler, G., & Zuba, D. (2004). Markov chain modeling of initiation and demand: The case of the US cocaine epidemic. *Health Care Management Science, 7,* 319–329. doi:10.1007/s10729-004-7540-4

Chudek, M., & Henrich, J. (2011). Culture–gene coevolution, norm-psychology, and the emergence of human prosociality. *Trends in Cognitive Sciences, 15,* 218–226. doi:10.1016/j.tics.2011.03.003

Cialdini, R. B. (2001). *Influence: Science and practice* (4th ed.). Boston, MA: Allyn & Bacon.

Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology, 58,* 1015–1026. doi:10.1037/0022-3514.58.6.1015

Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist, 49,* 997–1003. doi:10.1037/0003-066X.49.12.997

Coombs, C. (1964). *A theory of data*. New York, NY: Wiley.

Coultas, J. C. (2004). When in Rome . . . An evolutionary perspective on conformity. *Group Processes & Intergroup Relations, 74,* 317–331. doi:10.1177/1368430204046141

Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences, 13,* 36–43. doi:10.1016/j.tics.2008.10.002

Cramer, J. S. (2003). *Logit models*. New York, NY: Cambridge University Press.

Daley, D. J., & Gani, J. (1999). *Epidemic modelling: An introduction*. Cambridge, England: Cambridge University Press.

Davis, J. H. (1973). Group decision and social interaction: A theory of social decision schemes. *Psychological Review, 80,* 97–125. doi:10.1037/h0033951

Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes, 65,* 201–211. doi:10.1006/obhd.1996.0020

De Boeck, P., Wilson, M., & Acton, G. S. (2005). A conceptual and psychometric framework for distinguishing categories and dimensions. *Psychological Review, 112,* 129–158. doi:10.1037/0033-295X.112.1.129

De Fleur, M. L. (1966). Mass communication and social change. *Social Forces, 44,* 314–326.

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *Journal of Abnormal and Social Psychology, 51,* 629–636. doi:10.1037/h0046408

Eagly, A., & Chaiken, S. (1993). *The psychology of attitudes*. Fort Worth, TX: Harcourt Brace Jovanovich.

Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies*. Washington, DC: Brookings Institution Press.

Evans, J. W. (2003). Have Americans' attitudes become more polarized? An update. *Social Science Quarterly, 84,* 71–90. doi:10.1111/1540-6237.8401005

Fechner, G. (1966). *Elemente der psychophysik* [Elements of psychophysics] (Vol. 1, H. E. Adler, Trans.). New York, NY: Holt, Rinehart & Winston. (Original work published 1860)

Feller, W. (1957). *An introduction to probability theory and its applications* (2nd ed., Vol. 1). New York, NY: Wiley.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations, 7,* 117–140. doi:10.1177/001872675400700202

Fiedler, K., & Juslin, P. (Eds.). (2006). *Information sampling and adaptive cognition*. New York, NY: Cambridge University Press.

Fiorina, M. P., Abrams, S. J., & Pope, J. C. (2005). *Culture war? The myth of a polarized America*. New York, NY: Pearson/Longman.

French, J. P. R., Jr., & Raven, B. (1960). The bases of social power. In D. Cartwright & A. Zander (Eds.), *Group dynamics* (pp. 607–623). New York, NY: Harper and Row.

Gärdenfors, P. (2000). *The geometry of thought*. Cambridge, MA: MIT Press.

Gerard, H. B., Wilhelmy, R. A., & Conolley, E. S. (1968). Conformity and group size. *Journal of Personality and Social Psychology, 8,* 79–82. doi:10.1037/h0025325

Gladwell, M. (2000). *The tipping point: How little things can make a big difference*. New York, NY: Little, Brown.

Granovetter, M. (1978). Threshold models for collective behavior. *American Journal of Sociology, 83,* 1420–1443. doi:10.1086/226707

Granovetter, M., & Soong, R. (1983). Threshold models of diffusion and collective behavior. *Journal of Mathematical Sociology, 9,* 165–179. doi:10.1080/0022250X.1983.9989941

Grodzins, M. M. (1958). *The metropolitan area as a racial problem*. Pittsburgh, PA: University of Pittsburgh Press.

Harré, R., Clarke, D. D., & De Carlo, N. (1985). *Motives and mechanisms: An introduction to the psychology of action*. London, England: Methuen.

Harvey, M. D., & Enzle, M. E. (1981). A cognitive model of social norms for understanding the transgression–helping effect. *Journal of Personality and Social Psychology, 41,* 866–875. doi:10.1037/0022-3514.41.5.866

Hastie, R., & Kameda, T. (2005). The robust beauty of majority rules in group decisions. *Psychological Review, 112,* 494–508. doi:10.1037/0033-295X.112.2.494

Hastie, R., Penrod, S. D., & Pennington, N. (1983). *Inside the jury*. Cambridge, MA: Harvard University Press.

Heath, C., Larrick, R. P., & Wu, G. (1999). Goals as reference points. *Cognitive Psychology, 38,* 79–109. doi:10.1006/cogp.1998.0708

Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior, 19,* 215–241. doi:10.1016/S1090-5138(98)00018-X

Hertel, G., & Kerr, N. L. (2001). Priming and in-group favoritism: The impact of normative scripts in the minimal group paradigm. *Journal of Experimental Social Psychology, 37,* 316–324. doi:10.1006/jesp.2000.1447

Hinsz, V. B. (1990). Cognitive and consensus processes in group recognition memory performance. *Journal of Personality and Social Psychology, 59,* 705–718. doi:10.1037/0022-3514.59.4.705

Hull, C. L. (1938). The goal-gradient hypothesis applied to some 'field-force' problems in the behavior of young children. *Psychological Review, 45,* 271–299. doi:10.1037/h0053885

Kagehiro, D. K., & Stanton, W. C. (1985). Legal vs. quantified definitions of standards of proof. *Law and Human Behavior, 9,* 159–178. doi:10.1007/BF01067049

Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist, 58,* 697–720. doi:10.1037/0003-066X.58.9.697

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review, 93,* 136–153. doi:10.1037/0033-295X.93.2.136

Kelley, H. H., Holmes, J. G., Kerr, N. L., Reis, H. T., Rusbult, C. E., & Van Lange, P. A. M. (2003). *An atlas of interpersonal situations*. Cambridge, England: Cambridge University Press.

Kelman, H. C. (1958). Compliance, identification and internalization: Three processes of attitude change. *The Journal of Conflict Resolution, 2,* 51–60. doi:10.1177/002200275800200106

Kerr, N. L., Atkin, R. S., Stasser, G., Meek, D., Holt, R. W., & Davis, J. H. (1976). Guilt beyond a reasonable doubt: Effects of concept definition and assigned decision rule on the judgments of mock jurors. *Journal of Personality and Social Psychology, 34,* 282–294. doi:10.1037/0022-3514.34.2.282

Kerr, N. L., & MacCoun, R. J. (1985). The effects of jury size and polling method on the process and product of jury deliberation. *Journal of Personality and Social Psychology, 48,* 349–363. doi:10.1037/0022-3514.48.2.349

Kerr, N. L., & MacCoun, R. J. (in press). Is the leniency asymmetry really dead? Misinterpreting asymmetry effects in criminal jury deliberation. *Group Processes and Interpersonal Relations*.

Kerr, N. L., MacCoun, R. J., Hansen, C. H., & Hymes, J. A. (1987). Gaining and losing social support: Momentum in decision-making groups. *Journal of Experimental Social Psychology, 23,* 119–145. doi:10.1016/0022-1031(87)90028-X

Kerr, N., MacCoun, R. J., & Kramer, G. (1996). Bias in judgment: Comparing individuals and groups. *Psychological Review, 103,* 687–719. doi:10.1037/0033-295X.103.4.687

Kerr, N. L., Stasser, G., & Davis, J. H. (1979). Model testing, model fitting, and social decision schemes. *Organizational Behavior and Human Performance, 23,* 399–410. doi:10.1016/0030-5073(79)90006-0

Kimble, G. A. (1996). *Psychology: The hope of a science*. Cambridge, MA: MIT Press.

Knowles, E. S. (1999). Distance matters more than you think! *Personality and Social Psychology Bulletin, 25,* 1045–1048. doi:10.1177/01461672992511011

Kriss, M., Kinchla, R. A., & Darley, J. M. (1977). A mathematical model for social influences on perceptual judgments. *Journal of Experimental Social Psychology, 13,* 403–420. doi:10.1016/0022-1031(77)90026-9

Krueger, J., & Clement, R. W. (1994). The truly false consensus effect: An ineradicable and egocentric bias in social perception. *Journal of Personality and Social Psychology, 67,* 596–610. doi:10.1037/0022-3514.67.4.596

Kurzban, R., & Leary, M. R. (2001). Evolutionary origins of stigmatization: The functions of social exclusion. *Psychological Bulletin, 127,* 187–208. doi:10.1037/0033-2909.127.2.187

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). New York, NY: Cambridge University Press.

Latané, B. (1981). The psychology of social impact. *American Psychologist, 36,* 343–356. doi:10.1037/0003-066X.36.4.343

Latané, B., & Dabbs, J. M., Jr. (1975). Sex, group size and helping in three cities. *Sociometry, 38,* 180–194. doi:10.2307/2786599

Latané, B., & Darley, J. (1970). *The unresponsive bystander: Why doesn't he help?* New York, NY: Appleton-Century-Crofts.

Latané, B., & Wolf, S. (1981). The social impact of majorities and minorities. *Psychological Review, 88,* 438–453. doi:10.1037/0033-295X.88.5.438

Laughlin, P. R. (1999). Collective induction: Twelve postulates. *Organizational Behavior and Human Decision Processes, 80,* 50–69. doi:10.1006/obhd.1999.2854

Leary, M. R., & Baumeister, R. F. (2000). The nature and function of self-esteem: Sociometer theory. *Advances in Experimental Social Psychology, 32,* 1–62. doi:10.1016/S0065-2601(00)80003-9

Liao, T. F. (1994). *Interpreting probability models: Logit, probit, and other generalized models* (Quantitative Applications in the Social Sciences Sage University Paper No. 101). Thousand Oaks, CA: Sage.

Liu, J. H., & Latané, B. (1998). The catastrophic link between the importance and extremity of political attitudes. *Political Behavior, 20,* 105–126. doi:10.1023/A:1024828729174

Long, J. S. (1997). *Regression models for categorical and limited dependent variables*. Thousand Oaks, CA: Sage.

Lorge, I., & Solomon, H. (1955). Two models of group behavior in the solution of Eureka-type problems. *Psychometrika, 20,* 139–148. doi:10.1007/BF02288986

Luce, R. D. (1959). *Individual choice behavior*. New York, NY: Wiley.

Luce, R. D., & Raiffa, H. (1957). *Games and decisions: An introduction and critical survey*. New York, NY: Wiley.

MacCoun, R. J. (1984). Modeling the impact of extralegal bias and defined standards of proof on the decisions of mock jurors and juries. *Dissertation Abstracts International, 46,* 700B.

MacCoun, R. J. (1998). Biases in the interpretation and use of research results. *Annual Review of Psychology, 49,* 259–287. doi:10.1146/annurev.psych.49.1.259

MacCoun, R. J., Cook, P., Muschkin, C., & Vigdor, J. (2008). Distinguishing spurious and real norm effects: Evidence from artificial societies, small-group experiments, and real schoolyards. *Review of Law & Economics, 4,* 695–714. doi:10.2202/1555-5879.1226

MacCoun, R. J., & Kerr, N. L. (1988). Asymmetric influence in mock jury deliberation: Jurors' bias for leniency. *Journal of Personality and Social Psychology, 54,* 21–33. doi:10.1037/0022-3514.54.1.21

MacCoun, R. J., & Paletz, S. (2009). Citizens' perceptions of ideological bias in research on public policy controversies. *Political Psychology, 30,* 43–65. doi:10.1111/j.1467-9221.2008.00680.x

MacCoun, R. J., & Reuter, P. (2001). *Drug war heresies: Learning from other vices, times, and places*. New York, NY: Cambridge University Press. doi:10.1017/CBO9780511754272

Macy, M. W. (1991). Chains of cooperation: Threshold effects in collective action. *American Sociological Review, 56,* 730–747. doi:10.2307/2096252

Mahajan, V., Muller, E., & Bass, F. (1995). Diffusion of new products: Empirical generalizations and managerial uses. *Management Science, 14,* G79–G88.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics* (pp. 105–142). New York, NY: Academic Press.

McFadden, D. (2001). Economic choices (2000 Nobel Lecture). *The American Economic Review, 91,* 351–378. doi:10.1257/aer.91.3.351

McGuire, W. J. (1968). The nature of attitudes and attitude change. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (2nd ed., Vol. 3, pp. 136–314). Reading, MA: Addison-Wesley.

Meiser, T., Hein-Eggers, M., Rompe, P., & Rudinger, G. (1995). Analyzing homogeneity and heterogeneity of change using Rasch and latent class models: A comparative and integrative approach. *Applied Psychological Measurement, 19,* 377–391. doi:10.1177/014662169501900407

Milgram, S., Bickman, L., & Berkowitz, L. (1969). Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology, 13,* 79–82. doi:10.1037/h0028070

Miller, C. M., McIntyre, S. H., & Mantrala, M. K. (1993). Toward formalizing fashion theory. *Journal of Marketing Research, 30,* 142–157. doi:10.2307/3172824

Moscovici, S., Lage, E., & Naffrechoux, M. (1969). Influences of a consistent minority on the responses of a majority in a color perception task. *Sociometry, 32,* 365–380. doi:10.2307/2786541

Mullen, B. (1983). Operationalizing the effect of the group on the individual: A self-attention perspective. *Journal of Experimental Social Psychology, 19,* 295–322. doi:10.1016/0022-1031(83)90025-2

Mutz, D. C. (1998). *Impersonal influence: How perceptions of mass collectives affect political attitudes*. Cambridge, England: Cambridge University Press.

National Opinion Research Center. (2009). *GSS methodological report No. 24*. Retrieved from http://www.norc.org/GSS+Website/Data+Analysis/

Nemeth, C. J. (1986). Differential contributions of majority and minority influence. *Psychological Review, 93,* 23–32. doi:10.1037/0033-295X.93.1.23

Nowak, A., Szamrej, J., & Latané, B. (1990). From private attitudes to public opinion: A dynamic theory of social opinion. *Psychological Review, 97,* 362–376. doi:10.1037/0033-295X.97.3.362

Nowak, A., & Vallacher, R. R. (1994). *Dynamical systems in social psychology.* San Diego, CA: Academic Press.

Parker, S., Casey, J., Ziriax, J. M., & Silberberg, A. (1988). Random monotone data fit simple algebraic models: Correlation is not confirmation. *Psychological Bulletin, 104,* 417–423. doi:10.1037/0033-2909.104.3.417

Penrod, S., & Hastie, R. (1980). A computer simulation of jury decision making. *Psychological Review, 87,* 133–159. doi:1037/0033-295X.87.2.133

Petty, R. E., & Cacioppo, J. T. (1986). *Communication and persuasion: Central and peripheral routes to attitude change*. New York, NY: Springer-Verlag.

Powers, W. T. (1973). *Behavior: The control of perception*. Chicago, IL: Aldine.

Prentice, D. A., & Miller, D. T. (1996). Pluralistic ignorance and the perpetuation of norms by unwitting actors. *Advances in Experimental Social Psychology, 28,* 161–209. doi:10.1016/S0065-2601(08)60238-5

Reise, S. P., Ainsworth, A. T., & Haviland, M. G. (2005). Item response theory: Fundamentals, applications, and promise in psychological research. *Current Directions in Psychological Science, 14,* 95–101. doi:10.1111/j.0963-7214.2005.00342.x

Robbins, J. M., & Krueger, J. I. (2005). Social projection to ingroups and outgroups: A review and meta-analysis. *Personality and Social Psychology Review, 9,* 32–47. doi:10.1207/s15327957pspr0901_3

Rosenwein, R. E., & Campbell, D. T. (1992). Mobilization to achieve collective action and democratic majority/plurality amplification. *Journal of Social Issues, 48,* 125–138. doi:10.1111/j.1540-4560.1992.tb00888.x

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing* (Vol. 1, pp. 45–76). Cambridge, MA: MIT Press.

Ryan, B., & Gross, N. C. (1943). The diffusion of hybrid seed corn in two Iowa communities. *Rural Sociology, 8,* 15–24.

Schelling, T. C. (1969). Models of segregation. *The American Economic Review, 59,* 488–493.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology, 1,* 143–186.

Schelling, T. C. (1978). *Micromotives and macrobehavior*. New York, NY: Norton.

Smith, E. R. (2009). Distributed connectionist models in social psychology. *Social and Personality Psychology Compass, 3,* 64–76. doi:10.1111/j.1751-9004.2008.00160.x

Smoke, W. H., & Zajonc, R. B. (1962). On the reliability of group judgments and decisions. In J. H. Crisswell, H. Solomon, & P. Suppes

(Eds.), *Mathematical methods in small group processes* (pp. 322–333). Stanford, CA: Stanford University Press.

Sorkin, R. D., Hays, C. J., & West, R. (2001). Signal-detection analysis of group decision making. *Psychological Review, 108,* 183–203. doi:10.1037/0033-295X.108.1.183

Stasser, G., & Davis, J. H. (1981). Group decision making and social influence: A social interaction sequence model. *Psychological Review, 88,* 523–551. doi:10.1037/0033-295X.88.6.523

Stasser, G., Kerr, N. L., & Davis, J. H. (1989). Influence processes and consensus models in decision-making groups. In P. Paulus (Ed.), *Psychology of group influence* (2nd ed., pp. 327–355). Hillsdale, NJ: Erlbaum.

Stelter, B., & Arango, T. (2009, May 3). Losing popularity contest, MySpace tries a makeover. *The New York Times.* Retrieved from http://www.nytimes.com/2009/05/04/technology/companies/04myspace.html

Stevens, S. S. (1957). On the psychophysical law. *Psychological Review, 64,* 153–181. doi:10.1037/h0046162

Strang, D., & Soule, S. A. (1998). Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology, 24,* 265–290. doi:10.1146/annurev.soc.24.1.265

Substance Abuse and Mental Health Services Administration. (2002). *National Survey on Drug Use & Health: Detailed tables (Table 4.8a)*. Rockville, MD: Office of Applied Studies, Substance Abuse and Mental Health Services Administration, U.S. Department of Health & Human Services.

Sunstein, C. (2003). *Why societies need dissent*. Cambridge, MA: Harvard University Press.

Swets, J. A. (1988, June 3). Measuring the accuracy of diagnostic systems. *Science, 240,* 1285–1293. doi:10.1126/science.3287615

Tanford, S., & Penrod, S. (1983). Computer modeling of influence in the jury: The role of the consistent juror. *Social Psychology Quarterly, 46,* 200–212. doi:10.2307/3033791

Tanford, S., & Penrod, S. (1984). Social influence model: A formal integration of research on majority and minority influence processes. *Psychological Bulletin, 95,* 189–225. doi:10.1037/0033-2909.95.2.189

Tetlock, P. E. (1992). The impact of accountability on judgment and choice: Toward a social contingency model. *Advances in Experimental Social Psychology, 25,* 331–376. doi:10.1016/S0065-2601(08)60287-7

Thom, R. (1975). *Structural stability and morphogenesis: An outline of a theory of models* (D. H. Fowler, Trans.). Reading, MA: Benjamin. (Original work published 1972)

Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review, 34,* 273–286. doi:10.1037/h0070288

Tindale, R. S. (1993). Decision errors made by individuals and groups. In N. Castellan Jr. (Ed.), *Individual and group decision making: Current issues* (pp. 109–124). Hillsdale, NJ: Erlbaum.

Tindale, R. S., Davis, J. H., Vollrath, D. A., Nagao, D. H., & Hinsz, V. B. (1990). Asymmetrical social influence in freely interacting groups: A test of three models. *Journal of Personality and Social Psychology, 58,* 438–449. doi:10.1037/0022-3514.58.3.438

Tourangeau, R., Rips, L. J., & Raskinski, K. (2000). *The psychology of survey response.* New York, NY: Cambridge University Press.

Trafimow, D., & Finlay, K. (1996). The importance of subjective norms for a minority of the people. *Personality and Social Psychology Bulletin, 22,* 820–828. doi:10.1177/0146167296228005

Train, K. (2009). *Discrete choice methods with simulation* (2nd ed.). New York, NY: Cambridge University Press.

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90,* 293–315. doi:10.1037/0033-295X.90.4.293

van de Kragt, A. J. C., Orbell, J. M., & Dawes, R. M. (1983). The minimal contributing set as a solution to public goods problems. *American Political Science Review, 77,* 112–122.

van der Maas, H. L. L., Kolstein, R., & van der Pligt, J. (2003). Sudden transitions in attitudes. *Sociological Methods & Research, 32,* 125–152. doi:10.1177/0049124103253773

Visser, P. S., Bizer, G. Y., & Krosnick, J. A. (2006). Exploring the latent structure of strength-related attitude attributes. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 38, pp. 1–60). New York, NY: Academic Press.

Vogel, E. H., Castro, M. E., & Saavedra, M. A. (2004). Quantitative models of Pavlovian conditioning. *Brain Research Bulletin, 63,* 173–202. doi:10.1016/j.brainresbull.2004.01.005

Watzlawick, P., Bavelas, J. B., & Jackson, D. D. (1967). *Pragmatics of human communication*. New York, NY: Norton.

Wiesenthal, D. L., Austrom, D., & Silverman, I. (1983). Diffusion of responsibility in charitable donations. *Basic and Applied Social Psychology, 4,* 17–27. doi:10.1207/s15324834basp0401_2

Wilensky, U. (1998). *NetLogo segregation model*. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. Retrieved from http://ccl.northwestern.edu/netlogo/models/Segregation

Wilensky, U. (1999). *NetLogo*. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. Retrieved from http://ccl.northwestern.edu/netlogo/

Wood, W., Lundgren, S., Oullette, J. A., Busceme, S., & Blackstone, T. (1994). Minority influence: A meta-analytic review of social influence processes. *Psychological Bulletin, 115,* 323–345. doi:10.1037/0033-2909.115.3.323

Zaller, J. R. (1992). *The nature and origins of mass opinion*. New York, NY: Cambridge University Press.

(*Appendices follow*)

Appendix A

Deriving the BOP Models

The BOP models propose that the probability of changing one's position from one state ($y = 0$) to another ($y = 1$), $p(y = 1)$, is a function of the observed number of targets (T) who hold one's position and sources (S) who hold the opposite position, as well as an unobserved latent propensity to resist the sources' influence. For simplicity, I represent $y = 1$ and $y = 0$ as y_1 and y_0 and use $\Lambda(x)$ as an abbreviation for the inverse logit function:

$$\Lambda(x) = \frac{\exp(x)}{1 + \exp(x)} = \frac{1}{1 + \exp(-x)}.$$

Let $\theta = f(S, T)$ be a function representing the relative support for y_1 versus y_0 (i.e., the popularity of y_1). Plausible candidates for θ include S/T, $S/(S + T)$, $S - T$, and $|S - T|$. In this appendix, I restrict my attention to $S/(S + T)$ (i.e., bBOP).

To derive the BOP model, I first outline the strict and random utility approaches, which reach the same logistic choice model via different assumptions. It is important to emphasize that neither approach requires any notion of utility maximization or strong assumptions about rationality. I then show how BOP can be derived from either approach.

Strict Utility Model

For a simple dichotomous choice, Luce's (1959, Theorem 4, pp. 25, 38) famous choice rule states that there exists a positive ratio scale v such that

$$p(y_1) = \frac{1}{1 + \dfrac{v(y_0)}{v(y_1)}}.$$

More informally, this shows that the ratio of $v(y_1)$ to $v(y_0)$ can be treated as the odds of preferring y_1 to y_0. Following Fechner (1860/1966), Luce argued that $p(y_1) = F[u(x) - u(y)]$ when $u = (1/c)\ln(v) + a$, where $c > 0$ and $0 < p(y_1) < 1$. From this, Luce (p. 40) derived

$$p(y_1) = \Lambda(c[u(y_1) - u(y_0)]).$$

Random Utility Model

McFadden (1974) has shown that a logistic choice model can be motivated by invoking a utility model with a random component (Thurstone, 1927), though McFadden and other economists preferred to view this as a statistical rather than a psychophysical property. Letting ε_i refer to unobservable variables influencing the choice, z_i to observable variables, and β to a parameter to be estimated, then

$$\beta z_1 + \varepsilon_1 > \beta z_0 + \varepsilon_0.$$

McFadden assumed that the ε terms have an extreme value distribution, which results in a logistic distribution for $\varepsilon_0 - \varepsilon_1$, so that

$$p(y_1) = \Lambda(\beta[z_1 - z_0]).$$

Application to the Social Influence Situation

To derive the BOP model from a logistic choice model, assume that the choice as to whether to change position is a function of θ (an index representing the relative popularity of the two options), ψ_1 and ψ_0 (representing all other systematic influences favoring change and no change, respectively), and, in the random utility derivation, ε_1 and ε_0 (representing random error). The ψ terms are alternative-specific constants with nonzero mean; including them in the model makes the ε terms have zero mean (Train, 2009, p. 20). In the situations examined in this article, θ is observed (or experimentally manipulated), but the other terms are inferred from the data. When θ is experimentally manipulated, θ and ψ can be interpreted as representing normative versus informational influence (Deutsch & Gerard, 1955); reference power versus reward, coercive, legitimate, or expert power (French & Raven, 1960); and/or descriptive norms versus injunctive norms (Cialdini et al., 1990). When θ is observed but not manipulated, it may represent "strength in arguments" as well as "strength in numbers," and so without additional information, θ and ψ can be interpreted only as a distinction between those influences that do versus do not vary directly with relative faction size.

The θ term, representing the relative popularity of the two positions, is an attribute of the situation rather than either alternative alone. In the logistic choice model—in either the strict or random utility interpretation—only differences between utilities matter, and we can rearrange the components so that the actor will be indifferent when $c\theta = \psi$, where c is a scaling constant and $\psi = \psi_0 - \psi_1$ (see Cramer, 2003, p. 17; Train, 2009, pp. 21–22). This implies that $p_\Delta = \Lambda(c\theta - \psi)$, which will produce a sigmoid function of θ with an inflection point at $b \equiv \psi/c$. Substituting b into that equation, we get the basic BOP model:

$$p_\Delta = \Lambda[c(\theta - b)].$$

Note that the standard logistic regression model for one predictor, $\Lambda(\beta_0 + \beta_1 x)$, is equivalent to BOP's $\Lambda(c[x - b])$, if $x = \theta$, $c = \beta_1$, and $b = -\beta_0/\beta_1$, where the latter ratio between intercepts is the inflection point of the function (Cramer, 2003, p. 13).

(*Appendices continue*)

Appendix B

Does BOP Find Thresholds Where It Should and Where It Shouldn't?

Finding Known Thresholds

The BOP models posit that the b parameter reveals the mean social threshold in a population and that the c parameter reveals the extent to which there is a shared threshold across individuals, which reveals the extent to which there is a clear norm in the given situation for the population. If this reasoning is correct, the model should correctly infer the mean and variance of a simulated threshold. And ideally, it should do so fairly accurately for a wide range of statistical distributions, which would show that the model does not hinge on the psychophysical assumptions of the strict utility model or the error distribution assumptions of the random utility formulations (Appendix A).

To test this, I conducted 32 Monte Carlo simulations using a 2 (conformity vs. deliberation paradigm) × 2 (threshold as proportion vs. threshold as count), with eight types of statistical distributions crossed with each of the eight combinations. For the proportional thresholds (e.g., "I will change if 25% of the group disagrees with me"), I used a uniform distribution, four different normal distributions, and three different beta distributions. For the count thresholds (e.g., "I will change if two people disagree with me"), I used a uniform distribution, four different normal distributions, and three different Poisson distributions.

In each simulation, I randomly sampled 1,000 thresholds from the distribution. I then trimmed each sample by dropping any case with a threshold outside the 0–1 range for proportions or the 0–10 range for counts. As intended, this produced some skew in each distribution, with skew coefficients ranging from –0.42 to 0.63 in the proportional distributions and –0.61 to 1.52 in the count distributions.

I then simulated a conformity experiment with 1 to 10 sources disagreeing with the sole target, or a deliberation experiment in which the factional splits ranged from 0 sources versus 10 targets to 9 sources versus 1 target. This produces social influence pressure that rises in a concave function (0.5 to 0.91 for conformity/proportional), a linear function (0.1 to 1.0 for deliberation/proportional; 1 to 10 for conformity/odds), or a convex function (0.11 to 9 for deliberation/odds).

Table B1 shows that bBOP does an outstanding job of inferring the true mean for the proportions data ($r^2 = .994$, mean absolute discrepancy = .007), and uBOP performed nearly as well for odds data ($r^2 = .997$, mean absolute discrepancy = .061). There was no detectable difference in performance for the unidirectional versus bidirectional paradigms. Accuracy was influenced somewhat by the shape of the threshold distribution—but the BOP estimates are quite robust across assumptions about the threshold distribution.

The standard deviations from each simulation were regressed onto the inverse of the c estimates (using bBOP for the proportion data and uBOP for the count data) with an intercept of 0. These analyses suggested that a conversion factor of $1.68/c$ provided fairly accurate estimates of the standard deviations (adjusted $R^2 = .89$ for bBOP and .90 for uBOP), as can be seen by comparing the final two columns of Table B1. This is in accord with Long's (1997, p. 48) suggested 1.6 to 1.8 multiplier for converting a normally distributed probit to the corresponding logit.

Finding Thresholds in Random Data?

The BOP models perform sufficiently well across the tests in this article that one may wonder whether logistic threshold models can successfully fit just about anything. More generally, a concern in model fitting is that many algebraic models can do a good job of fitting random monotone data (see Parker, Casey, Ziriax, & Silberberg, 1988).

To examine this concern, I generated 10 simulated data sets, each consisting of 10 random numbers between 0 and 1, which were then arranged from smallest to largest, corresponding to the simulated proportions of participants conforming upon exposure to 1, 2, . . . , 10 sources. To simulate the conformity paradigm, I set targets = 1 for each source condition. To simulate the deliberation paradigm, targets = $10 - S$ for each source condition. The same 10 data series were used for both paradigms.

In fact, bBOP did provide a good fit to these simulated data (average $R^2 = .857$), though this falls well short of the fit to real data sets in this article. But it is important not to overinterpret this finding; this is neither a Turing test to distinguish human and nonhuman data, nor does some degree of fit imply a false positive as in a signal detection test. Rather, the question is whether the model parameters will imply a strong threshold process where there isn't one. In the deliberation paradigm, the mean b and c values for these simulated data were 0.452 and 4.872—quite close to the 0.500 and 5.000 benchmarks corresponding to the proportionality decision scheme model (see Table 6), a rising linear slope with no threshold inflection. Similarly, in the conformity paradigm, the mean b and c values for the simulated data were 0.744 and 16.057; the corresponding values for bBOP's best fit to a threshold-free straight line (from .1 to 1 for $S = 1$ to 10) are 0.813 and 17.052. So while bBOP provides reasonably good fit to random monotone data, it also yields strong clues that it is doing so.

(*Appendices continue*)

Table B1
Fitting the Mean and Variance in Simulated Threshold Distributions

Paradigm	Simulated distribution	*M*	*b*	*SD*	1.68/*c*
Threshold as a proportion			bBOP		bBOP
Conformity (unidirectional)	Uniform (.25 to .75)	0.507	*0.511*	0.142	*0.130*
	Normal (M = .5, SD = .1)	0.500	*0.503*	0.096	*0.089*
	Normal (M = .5, SD = .25)	0.504	*0.508*	0.221	*0.222*
	Normal (M = .66, SD = .1)	0.663	*0.665*	0.104	*0.096*
	Normal (M = .66, SD = .25)	0.616	*0.626*	0.211	*0.209*
	Beta (a = 2, b = 2)	0.501	*0.510*	0.223	*0.221*
	Beta (a = 2, b = 3)	0.393	*0.411*	0.200	*0.183*
	Beta (a = 2, b = 5)	0.285	*0.308*	0.159	*0.146*
Deliberation (bidirectional)	Uniform (.25 to .75)	0.498	*0.497*	0.148	*0.166*
	Normal (M = .5, SD = .1)	0.501	*0.501*	0.097	*0.096*
	Normal (M = .5, SD = .25)	0.506	*0.507*	0.220	*0.229*
	Normal (M = .66, SD = .1)	0.659	*0.660*	0.098	*0.097*
	Normal (M = .66, SD = .25)	0.624	*0.634*	0.209	*0.218*
	Beta (a = 2, b = 2)	0.503	*0.502*	0.219	*0.233*
	Beta (a = 2, b = 3)	0.394	*0.384*	0.208	*0.215*
	Beta (a = 2, b = 5)	0.287	*0.276*	0.162	*0.164*
Threshold as a count (odds)			uBOP		uBOP
Conformity (unidirectional)	Uniform (.33 to 3)	1.704	*1.702*	0.763	*0.851*
	Normal (M = 1, SD = .11)	1.004	*1.004*	0.115	*0.182*
	Normal (M = 1, SD = .33)	1.016	*1.022*	0.337	*0.264*
	Normal (M = 2, SD = .11)	1.996	*1.995*	0.110	*0.178*
	Normal (M = 2, SD = .33)	1.982	*1.995*	0.332	*0.259*
	Poisson (λ = 1)	1.552	*1.452*	0.769	*0.926*
	Poisson (λ = .75)	4.863	*4.829*	1.967	*2.073*
	Poisson (λ = .5)	6.660	*6.773*	1.796	*1.858*
Deliberation (bidirectional)	Uniform (.33 to 3)	1.687	*1.710*	0.777	*0.866*
	Normal (M = 1, SD = .11)	1.001	*1.000*	0.108	*0.091*
	Normal (M = 1, SD = .33)	1.005	*0.997*	0.342	*0.336*
	Normal (M = 2, SD = .11)	2.004	*2.000*	0.110	*0.081*
	Normal (M = 2, SD = .33)	2.005	*2.017*	0.329	*0.303*
	Poisson (λ = 1)	1.609	*1.430*	0.821	*0.549*
	Poisson (λ = .75)	4.699	*4.407*	2.000	*1.677*
	Poisson (λ = .5)	6.651	*6.469*	1.795	*1.792*

Note. Model estimates are in italics. bBOP = bidirectional influence burden of proof model; uBOP = unidimensional influence burden of proof model; b = threshold parameter; c = norm clarity parameter.

Received April 22, 2011
Revision received December 15, 2011
Accepted December 23, 2011 ■