Organizational Behavior and Human Decision Processes 129 (2015) 93–104

Contents lists available at ScienceDirect

Organizational Behavior and Human Decision Processes

journal homepage: www.elsevier.com/locate/obhdp

Balancing evidence and norms in cultural evolution

Robert J. MacCoun *

Goldman School of Public Policy and Berkeley Law, University of California at Berkeley, United States

ARTICLE INFO

Article history:
Received 15 January 2013
Accepted 11 September 2014

Accepted by Michael Morris, Ying-yi Hong and Chi-yue Chiu

Keywords:
Culture
Imitation
Conformity
Norms
Diffusion
Threshold

ABSTRACT

Psychologists have long studied the ways in which individuals draw inferences from evidence in their environment, and the conditions under which individuals forgo or ignore those inferences and instead conform to the choices of their peers. Recently, anthropologists and biologists have given considerable attention to the ways in which these two processes intersect to jointly shape culture. In this paper I extend the BOP ("burden of (social) proof"; MacCoun, 2012) analysis of "strength in numbers" with a parallel account of "strength in arguments," and examine ways the two processes might be linked. I compare these models to some leading accounts of individual learning and social transmission, suggesting opportunities for a closer integration of theory and research on cultural evolution across anthropology, biology, and psychology.

Introduction

Cultures evolve through a balance of individual learning and social transmission (Boyd & Richerson, 1985). This is no less true in academic scholarship than in other cultural communities. Consider the cultural practice of null-hypothesis testing in social science. Most readers of this journal engage in this practice; we were taught the practice, given reasons for the practice, and completed problem sets that allowed us to explore the merits of the practice. But very few of us independently discovered the practice; we adopted it because the community had already adopted it, and we persist in the tradition even when editors try to nudge us into alternative practices (Fidler, Thomason, Cumming, Finch, & Leeman, 2004). And irrespective of one's views of the evidence and logic behind null-hypothesis testing (see Cumming, 2014 for a recent overview), there is one feature we have adopted without any compelling mathematical or empirical reasons – the convention to set the critical rejection region at $p = .05$ (rather than, say, .02 or .20), as was proposed fairly arbitrarily by Fisher (1928, p. 45). Thus, null-hypothesis testing involves two issues: Where to place the threshold, and how strictly and uniformly to place the threshold. But at a meta-level, it illustrates the same issues with respect to two other thresholds – our epistemic and social thresholds for adopting that .05 threshold.

In this paper, I argue that for the advantages of understanding cultural transmission in terms of such shared thresholds on evidence and on norms. These thresholds, which establish our relative responsiveness to evidence and norms, are characterized by two properties (MacCoun, 2012). First, these thresholds have a location, and the asymmetry of that location (i.e., the extent to which it is differs from .5 on a 0–1 metric) reveals whether the assessment has a bias. Second, these thresholds can range from very soft to very hard, a property I call "clarity." I argue that these properties can be estimated from data, and that together, these estimated parameters can indicate the extent to which people have a shared conceptual scheme for assessment.

Beginning with the pioneering works by Cavalli-Sforza and Feldman (1981) and Boyd and Richerson (1985), there is now a large and impressive body of theoretical and empirical work on cultural learning and cultural evolution (see Bentley & O'Brien, 2011; Boyd & Richerson, 2005; Henrich, 2000; Henrich & McElreath, 2003; Hoppitt & Laland, 2013; Rendell et al., 2011). This work demonstrates the value of applying Darwinian concepts like selection, retention, and fitness to the emergence and endurance of cultural practices and beliefs.

In this paper, I take as a starting point the general notion that cultural selection and retention involve the interplay of two forces – "strength in arguments" (reasoning on the basis of evidence and deduction) and "strength in numbers" (imitation and conformity to the behavior of those in one's community). I do so by extending the BOP ("burden of proof") family of logistic threshold models (MacCoun, 2012, 2014), in two ways. First, I present a model of

* Present address: Stanford Law School, Stanford University, 559 Nathan Abbott Way, Stanford, CA 94305, United States.
E-mail address: maccoun@law.stanford.edu

http://dx.doi.org/10.1016/j.obhdp.2014.09.009
0749-5978/

Table 1
Notation used in the models.

Notation	Definition
s	Proportion who have chosen Option 1 at time t
s'	Proportion who have chosen Option 1 at time $t+1$
P_1, P_0	Probability of choosing Option 1 and Option 0, respectively, where $P_1 + P_0 \leqslant 1$
L	Probability that evidence is inconclusive; viz., $L = 1 - P_1 - P_0$
B	Net direct bias favoring Option 1 over Option 0
$1-\gamma$, γ	Weights given to individual and frequency-dependent (conformist) learning, respectively
A_1, A_0	Attractiveness of Option 1 and of Option 0
λ	Influence of differences in attraction scores
n_1, n_0	Numbers who have chosen Option 1 and Option 0 in most recent period
f	Bias toward copying most popular option (when $f > 1$, where $f = 1$ is no bias)
m	Ceiling parameter on bBOP and BEAN; $0 \leqslant m \leqslant 1$
b	bBOP norm threshold; $0 \leqslant b \leqslant 1$
c	Clarity of bBOP norm threshold; $0 \leqslant c \leqslant$
x	Proportion of evidence (excluding consensus information) favoring Option 1
a	aBOP evidence threshold; $0 \leqslant a \leqslant 1$
k	Clarity of aBOP threshold; $0 \leqslant k \leqslant$
α	Exogenous threshold in BEAN model
β	Clarity of BEAN threshold

Table 2
Probability models for choice under individual and social learning.

Source	Label	Model
MacCoun (2012)	bBOP	$p(\text{Option }1) = m/(1+\exp[-c(s-b)])$
This paper	aBOP	$p(\text{Option }1) = 1/(1+\exp[-k(x-a)])$
	BEAN	$p(\text{Option }1) = m/(1+\exp[-\beta(s_1 - x_0 + \alpha)])$
McElreath et al. (2008)	MEA1	$p(Option\ 1)_{IL} = \exp(\lambda A_1)/[\exp(\lambda A_1) + \exp(\lambda A_0)]$
	MEA2	$p(Option\ 1)_{FD} = n_{1,t}^f/(n_{1,t}^f + n_{0,t}^f)$
	MEA3	$p(Option\ 1) = (1-\gamma)\, p(Option\ 1)_{IL} + \gamma p(Option\ 1)_{FD}$

Note: See Table 1 for notation and definitions.

"strength in arguments" (aBOP) that closely parallels the structure of the bBOP model of "strength in numbers". I then link these to models into an integrative model of how people balance evidence and norms (BEAN). Finally, I compare and contrast these models with some major models of cultural evolution developed at the intersection of anthropology and the biological sciences (Boyd & Richerson, 1985, 2005; Henrich, 2000; McElreath et al., 2008). I attempt to show continuities between the two approaches, but also some friendly amendments to illustrate how features of the BOP and BEAN models might link their approach more closely to social psychological data, as well as formal models in psychology, economics, and sociology.

Norms and the burden of social proof

The bBOP model (MacCoun, 2012) describes the probability that an individual will switch positions on a dichotomous issue as a function of "strength in numbers" favoring the opposite position in a local population. The acronym "bBOP" stands for "bidirectional burden of proof" – one of a family of similar models in MacCoun (2012). The notation for this and other models discussed in this paper appears in Table 1 and the equation specifying bBOP appears in Table 2.

MacCoun (2012) shows how the model can be used as a common frame of reference for behavior in studies of conformity, group deliberation, diffusions of innovation, and neighborhood change. Consider a situation where an actor has reached an opinion on some dichotomous issue or choice, adopting a position or behavior or choice we will call Option 0. The actor then encounters a collection of other people, some of whom have made the opposite choice, Option 1. According to bBOP, the probability that the actor will now change from Option 0 to Option 1 is given by a logistic threshold function that compares the proportion (s) of "sources" (S) who favor the position opposite one's own in a population of size N (i.e., $s = S/N$) to a threshold parameter (b) that can be interpreted as the actor's perceived "burden of social proof" – the point at which Option 1's popularity is sufficiently high to begin tipping her toward switching from Option 0 to Option 1. Fig. 1 shows how the probability of influence varies with the location of the threshold and the popularity of the opposing position. When b is near 1 the actor places a steep burden of proof on the other side and is thus quite resistant to change. When b is at .5, the burden is shared by both sides, producing an implicit "majority wins" rule, even in the absence of any formal group procedures for consensus. When b is near 0, the actor is almost completely susceptible to any social influence to change positions.

The c parameter represents the "clarity" of the matching-to-threshold process. Clarity is inversely related to variance at both the individual and aggregate levels. At the individual level, clarity reflects how strictly one enforces the b threshold, and thus low c can reflect uncertainty or fuzziness about whether the level of social consensus exceeds one's personal threshold. At the collective level, c is inversely related to the standard deviation of the distribution of b across actors, so that a high clarity level implies a high degree of consensus about the threshold – a shared sense of where the burden of social proof lies in this situation. When c is very high, the model produces a hard threshold and predicts a step function; when c is very low, the model produces a soft threshold and predicts that choice becomes increasingly random.

Fig. 2a and b shows the effect of clarity under two different threshold levels. When $b = .5$ (Fig. 2a), as clarity increases the function begins to resemble a formal "majority wins" voting rule. But when b is near 0 (Fig. 2b), only one or two endorsers may be sufficient to persuade everyone to adopt their position, and as clarity increases the function suggests an implicit "Truth Wins" norm indicating that the group has some shared conceptual scheme (be it arithmetic, logic, theology, or economic theory) for recognizing a convincing position once it is articulated (see Kerr, MacCoun, & Kramer, 1996; Laughlin, 2011).[1] But note that the winning argument has to evoke a conceptual scheme that strongly favors it, and the conceptual scheme has to be broadly shared for "Truth Wins" to work. "Truth Wins" can also be distinguished from prestige-based influence (French & Raven, 1960; Henrich, 2000). In prestige-based influence, a sole advocate can have a disproportionate impact, but only if he or she has prestigious traits (reputation, maturity, status, a good track record for accuracy). bBOP could be modified to apply prestige weights to each source, but given the model's extremely good fit to data the added complexity and loss of parsimony seem unnecessary.

Finally, the m parameter is a "ceiling" parameter that reflects the maximum predicted opinion change in a given situation. A low ceiling parameter suggests that there are factors producing resistance to persuasion that are independent of relative faction size. MacCoun (2012) found that a ceiling parameter was necessary (for competitor models as well as for bBOP) to fit data in the Asch-type conformity paradigm, where a lone target is confronted with

[1] In 1931, a book entitled *Hundert Autoren Gegen Einstein* ("A Hundred Authors Against Einstein") was published in Germany. According to the sculptor Jacob Epstein (1975, p. 78), Einstein was unimpressed: "Were I wrong," he said, "one Professor would have been enough" – a quote that perfectly exemplifies the logic of the "Truth Wins" rule.

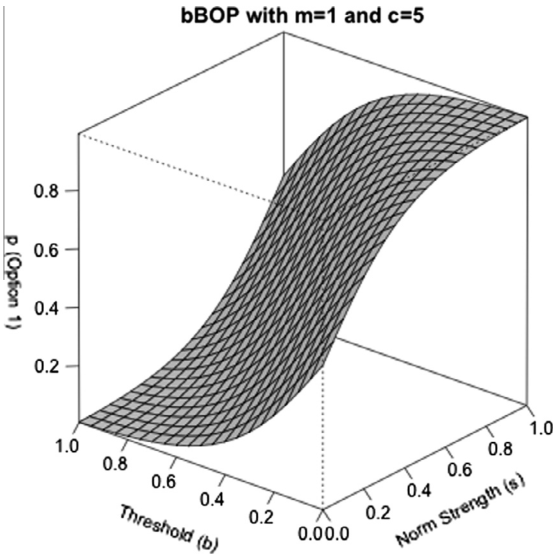

Fig. 1. Probability of Option 1 under bBOP with $m = 1$ and $c = 5$.

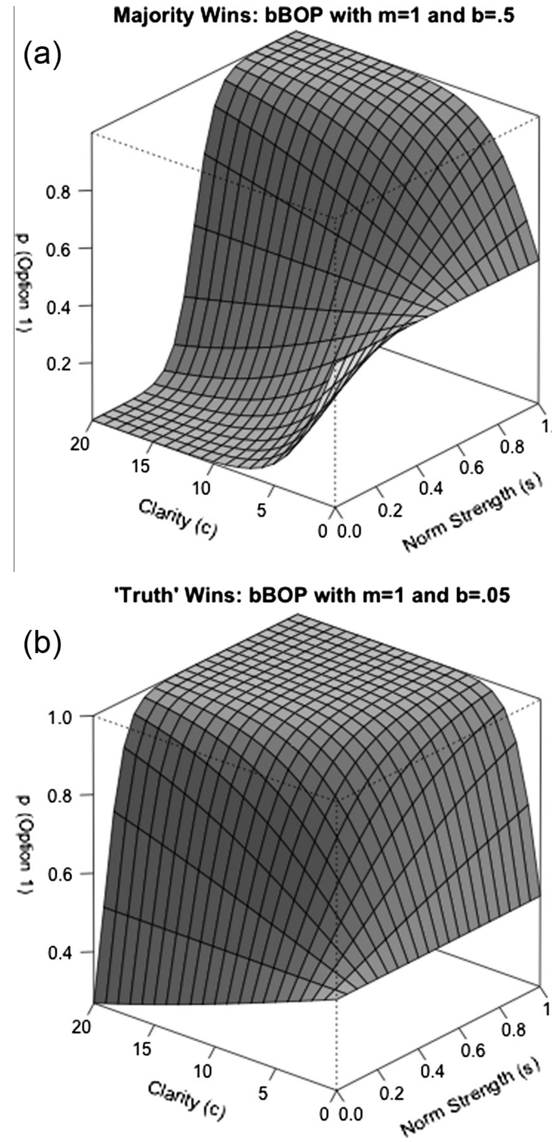

Fig. 2. (a) Majority Wins: bBOP with $m = 1$ and $b = .5$; (b) 'Truth' Wins: bBOP with $m = 1$ and $b = .05$.

some number of sources endorsing the opposite position. No ceiling was required (i.e., $m = 1$) in other experimental paradigms MacCoun (2012) examined.

This matching-to-standard process is a familiar one in psychology; similar processes are featured in control theory (e.g., Carver & Scheier, 1998; Powers, 1973), aspiration-level theories (Lewin, Dembo, Festinger, & Sears, 1944; Simon, 1956), and reference-dependent theories of choice (e.g., Kahneman & Tversky, 1979). It is well known that humans (Festinger, 1954; Leary & Baumeister, 2000) and other species (Couzin, 2009) readily engage in "quorum-sensing" by carefully monitoring the views of their neighbors.[2] Recent research suggests that these social comparison processes may occur in the same neural circuitry (the anterior cingulate cortex) used for error monitoring and cognitive control (Chang, Gariepy, & Platt, 2013). As discussed below, there is increasing evidence that these social consensus judgments are critical to the formation of shared culture (Morris, 2014).

Note that the bBOP model involves the burden of *social* proof (Cialdini, 2001; Festinger, 1954) rather than evidentiary or epistemic proof. In that sense, it reflects normative rather than informational influence (Deutsch & Gerard, 1955), and the influence of *descriptive* norms (what others say and do) rather than *injunctive* norms (what others think I should say and do; Cialdini, Reno, & Kallgren, 1990). There are a great many social influence variables that contribute to attitude and behavior change (see Eagly & Chaiken, 1993; Petty & Cacioppo, 1986) and undoubtedly many of them influence both threshold clarity and level in the model. MacCoun (2012) suggests that, in that sense, bBOP is a "skeletal" model, and is not intended to serve as a complete explanation of an influence situation. Rather, the model is a tool for estimating relevant parameters that provide a concise and useful description of a situation, and for comparing those parameter values across situations, populations, and research paradigms.

MacCoun (2012) examines the fit of the model to a variety of classic data sets in three paradigms. Fig. 3 shows four illustrative examples. In the *conformity* paradigm, a single target is confronted with 1 or more sources; the influence is mostly unidirectional. Panel a shows a typical *r*-shaped pattern of responses, in this case from the Milgram et al. (1969) examination of whether urban pedestrians would look to the sky (at nothing) as a function of the number of observable experimental confederates already doing so. bBOP easily fits these data, though so do other models (e.g., Latané's (1981) social impact theory). Panel b shows the *S*-shaped pattern of results from Asch's (1956) classic study of conformity, which examined whether naïve participants would state a patently incorrect response simply because other people (actually, experimental confederates) had already done so. bBOP is able to capture the fact that the second source has greater impact than the first, a feature that should not occur under the diminishing marginal impact assumption of Latané's (1981) model.[3]

In the *deliberation* paradigm, group members in opposing factions (arbitrarily labeled sources vs. target) influence each other in a bidirectional fashion. bBOP fits these data better than several other influence models examined by MacCoun (2012) – in

[2] "Consensus sensing" might be a more accurate term, but the label has become widespread – a new cultural artifact.

[3] The Asch (1956) research also found that conformity dropped dramatically when the target of social influence had a single supporter at the table. The bBOP model predicts this, because the *s* term, the proportion of opponents in the situation, drops as its denominator shifts from $S + 1$ to $S + 2$, an effect that is particularly pronounced for small *S*.

particular, it captures the profound asymmetry that is consistently seen in mock criminal jury data (Kerr & MacCoun, 2012; MacCoun & Kerr, 1988; for example the data for 12-person mock juries presented in Panel c; also see Fig. 4 discussed below). Specifically, advocates of conviction appear to face a steeper burden than advocates of acquittal, and hence a "Not Guilty" final verdict requires less initial support.

Finally, the *social imitation* paradigm involves the growth in popularity of a trait over time. Examples include drug epidemics, the diffusion of new technologies, and viral YouTube videos. Studies in this paradigm have tended to differ from the other two paradigms in two important ways. First, they examine much larger social aggregates – communities or society as a whole rather than dyads or small groups. Second, the data are usually observational rather than experimental, so the s index of popularity is potentially confounded with changes in information in the environment. The temporal element requires a different formulation of the model: MacCoun (2012) shows that a recursive form of bBOP (iBOP; see Table 3) fits classic datasets on the diffusion of innovations about as well as the classic Bass diffusion model (Bass, 1969; Mahajan, Muller, & Bass, 1995), as with the famous Ryan and Gross (1943) study of the adoption of hybrid seed corn among Iowa farmers (Panel d).

The bBOP model provides a unifying account for a wide variety of other important social influence phenomena. Schelling's (1978) tipping point model is a special case of bBOP where c is very large and b = .50; for that reason, bBOP shows why tipping points are rare and most change is more gradual. bBOP also closely approximates Granovetter's (1978) threshold model, in which thresholds vary across actors. MacCoun (2012) shows that many of the prototypical social decision schemes (Davis, 1973; Kerr et al., 1996; also see the special symposium in this journal edited by Parks & Kerr, 1999) – e.g., simple majority, truth wins, truth-supported wins – can be reproduced using bBOP, which circumvents the need for matrix computations and locates the data in a 2-dimensional parameter space that facilitates comparisons across schemes and studies. For example, Fig. 4 (adapted from MacCoun, 2014) shows data from 9 different mock and real criminal jury studies, as well as 7 different studies using "intellective" group tasks with a demonstrable correct answer, plotted in the bBOP parameter space. Note that the two regions do not overlap; groups working on intellective tasks share an asymmetric threshold that enables good minority solutions to prevail; criminal juries operate under some asymmetry (favoring acquittal) but with greater clarity, as might be expected given that they are instructed to apply a "beyond a reasonable doubt" standard (Kerr & MacCoun, 2012; MacCoun & Kerr, 1988).

The best way to fully explore the implications of the bBOP model is to conduct agent-based modeling (ABM) simulations (see Macy & Willer, 2002), in which dozens or hundreds of programmed "agents" simultaneously influence each other. Using this approach, MacCoun (2012) reproduced many of the qualitative patterns reported in Nowak, Szamrej, and Latané (1990) and other ABM studies, including the formation of remarkably stable spatial attitude clusters or groupings. The ABM approach also makes it possible to explore the effects of varying the number of neighbors whose positions one is able to monitor. Fig. 5a and b illustrate this for many different runs of the model, each extending for 50 updating periods. When this "vision" parameter was set to a very local level (1 visible neighbor in each of the 8 cardinal directions), societies that started out with an even split and a neutral "majority wins" threshold tended to remain evenly split over many generations; for each local neighborhood that tipped one direction, some other neighborhood tipped in the other, offsetting it (Fig. 5a). On the other hand, when vision extended 10 neighbors deep in each of the 8 cardinal directions, there was a global increase in correlated behavior change. As a result, by chance, some entire societies tipped one way while others tipped the other. MacCoun (2012) suggested that we are experiencing large increases in this vision parameter through the proliferation of opinion polls, Facebook

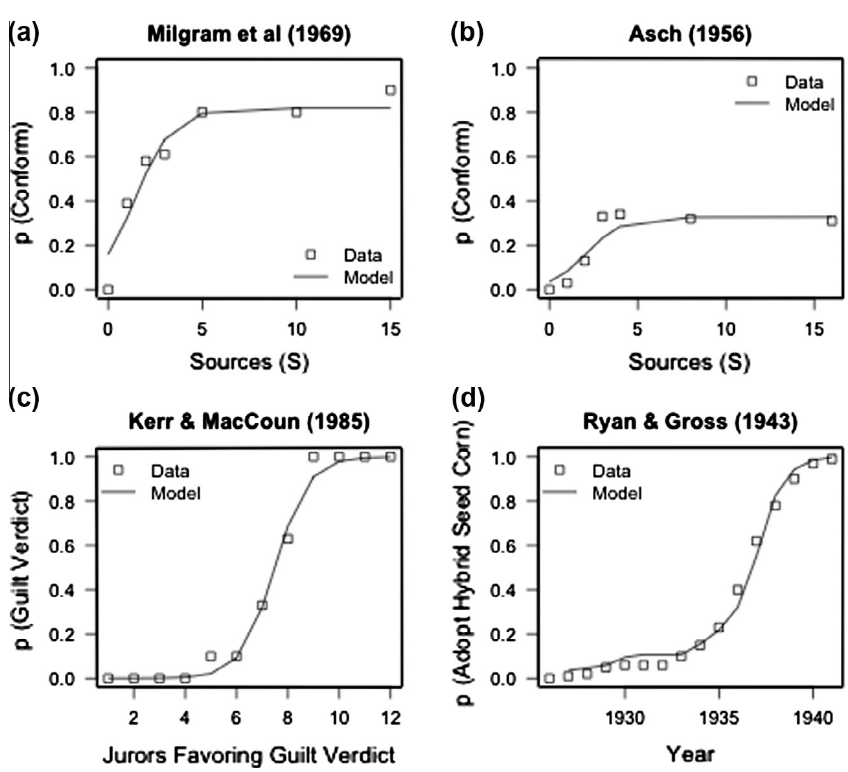

Fig. 3. bBOP best fit to four empirical data sets: (a) Milgram et al. (1969); (b) Asch (1956); (c) Kerr and MacCoun (1985; size 12); and (d) Ryan and Gross (1943).

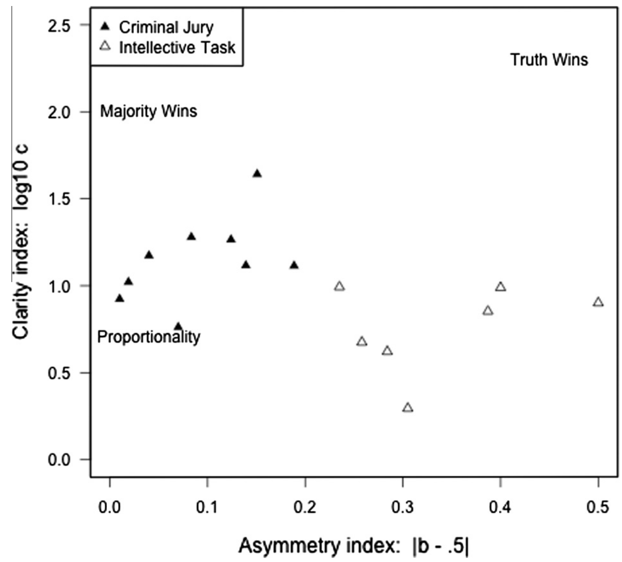

Fig. 4. Nine criminal jury studies and seven intellective task studies located in the bBOP parameter space (adapted from MacCoun, 2014, Figure 2).

Table 3
Growth models for individual and social learning.

Source	Label	Model
MacCoun (2012)	iBOP	$s' = s + (1-s)/(1+\exp[-c(s-b)])$
Mahajan et al. (1995)	BASS	$s' = s + (1-\gamma)(1-s) + \gamma s(1-s)$
Boyd and Richerson (1985), Boyd and Richerson (2005)	BR1	$s' = s + (1-s)P_1 - sP_0 = P_1 + Ls$
	BR2	$s' = s + s(1-s)B$
	BR3	$s' = s + s(1-s)[(1-\gamma)B + \gamma(2s-1)]$

Note: See Table 1 for notation and definitions.

"likes", Amazon reader ratings, and other forms of social media. At the level of national cultures, clusters of distinguishing tastes and customs that have been stable for centuries (e.g., in France vs. neighboring Spain) may decrease as a function of an increase in the vision parameter, reflecting the increasing ease of long-distance travel and communications.

The bBOP model is explicitly an account of social influence, and only represents non-social influence (e.g., "strength in arguments") implicitly, via the b threshold. A symmetrical b value near .50 is consistent with the kind of "majority-rule" influence that features prominently in recent evolutionary accounts of culture (e.g., Boyd & Richerson, 1985, 2000; Hastie & Kameda, 2005; Henrich, 2000). But MacCoun (2012), MacCoun (2014) finds that most data sets yield b estimates that are much more asymmetric. Drawing on earlier work (Kerr et al., 1996; Laughlin, 2011; Lorge & Solomon, 1955), I propose that an asymmetric b, combined with reasonably high c values, are an indication that there is a *shared conceptual scheme* in the population under study that makes arguments or evidence for one option more compelling. In such cases, the other side bears more of the implicit burden of social proof.

Evidence and the burden of epistemic proof

Whatever its merits as a model of social transmission, the bBOP model clearly lacks any account of learning from non-social evidence. But it is straightforward to propose that the same kind of threshold process that describes norm-based choice might describe evidence-based choice. In essence, epistemic proof is the other side of the bBOP coin; whereas the $c(s-b)$ formulation emphasizes the threshold against the pressure of social sources, we can flip things around and estimate the threshold we used to scrutinize evidence. Where bBOP reflects the strength in numbers ("normative influence"; Deutsch & Gerard, 1955), I propose aBOP as an account of the strength in arguments ("informational influence"). The model appears in Table 2, where x indexes the relative evidentiary strength for a trait or behavior or position, a is the evidentiary threshold, and k is an index of clarity. As with c in bBOP, k is inversely related to the standard deviation of the a distribution.

It is clear that this is far from a detailed process model of evidence assessment and integration. There are many different attributes that might get integrated in such a process, and they go by many different labels: observations (in everyday life), data (in science), evidence (in law), cues (in the lens model tradition; Cooksey, 1996), signals (in the signal detection theory tradition; Swets, 1988), arguments (in traditional theories of rhetoric and logic and in much of the persuasion literature in psychology), and payoffs (in the decision and game theory traditions). The aBOP model simply combines these in x_i, a summary judgment of the proportion of the evidence favoring one option over the other. When the threshold parameter $a = .50$, the actor will be inclined toward a "preponderance of evidence" standard, which can be a vague tendency or a strict rule, depending on the clarity parameter k. A stricter threshold (e.g., $a = .95$) is more like a "beyond a reasonable doubt" standard.

MacCoun (2012, Appendix A) shows how the bBOP model can be derived as a strict utility model (Luce, 1959) or a random utility model (McFadden, 1974) depending on whether one interprets the model psychophysically or statistically. The same reasoning applies to the aBOP model, except that the independent variable reflects evidence strength rather than norm strength. This changes the interpretation of the clarity and threshold parameters, as they now describe one's responsiveness to evidence rather than one's responsiveness to the popularity of the choices.

Signal detection theory (SDT; Swets, 1988) is the dominant model of choice in perceptual tasks. As a threshold model, aBOP is closely related to SDT; although most application of SDT assume normally distributions of signal and noise, applying SDT to logistic distributions produces almost identical results when a slight adjustment factor is used (DeCarlo, 1998). Specifically, in SDT, $d = (\psi_s - \psi_n)/\tau$ is the distance between the modes of the signal and noise distributions, respectively. For normal distributions, $\tau = 1$. For the logistic distribution, $\tau = \sigma\sqrt{3}/\pi \approx .5513\sigma$ (DeCarlo, 1998).

A common label for the SDT threshold is c, but to avoid confusion with bBOP's clarity parameter, I will use a. DeCarlo (1998) shows that for the logistic distribution, the probability of correctly identifying a signal when it is present is

$$p(Y = 1|Signal) = \frac{1}{1 + \exp\left[\frac{a - \psi_S}{\tau}\right]}$$

We can easily translate this into aBOP. MacCoun (2012, footnote 14) shows that the BOP clarity parameter equals approximately $\frac{\sigma}{1.8}$, and since $\tau \approx .5513\sigma$ and $1/.5513 = 1.8138$, we can see that $\tau \approx 1/k$ and

$$p(Y = 1|Signal) = \frac{1}{1 + \exp[k(a - \psi_S)]} = \frac{1}{1 + \exp[-k(\psi_S - a)]}$$

Thus aBOP gives the probability of a "true positive" when x equals the mode of the signal distribution, ψ_S. In the same way, the probability of a false positive rate is:

$$p(Y = 1|Noise) = \frac{1}{1 + \exp[-k(\psi_N - a)]}$$

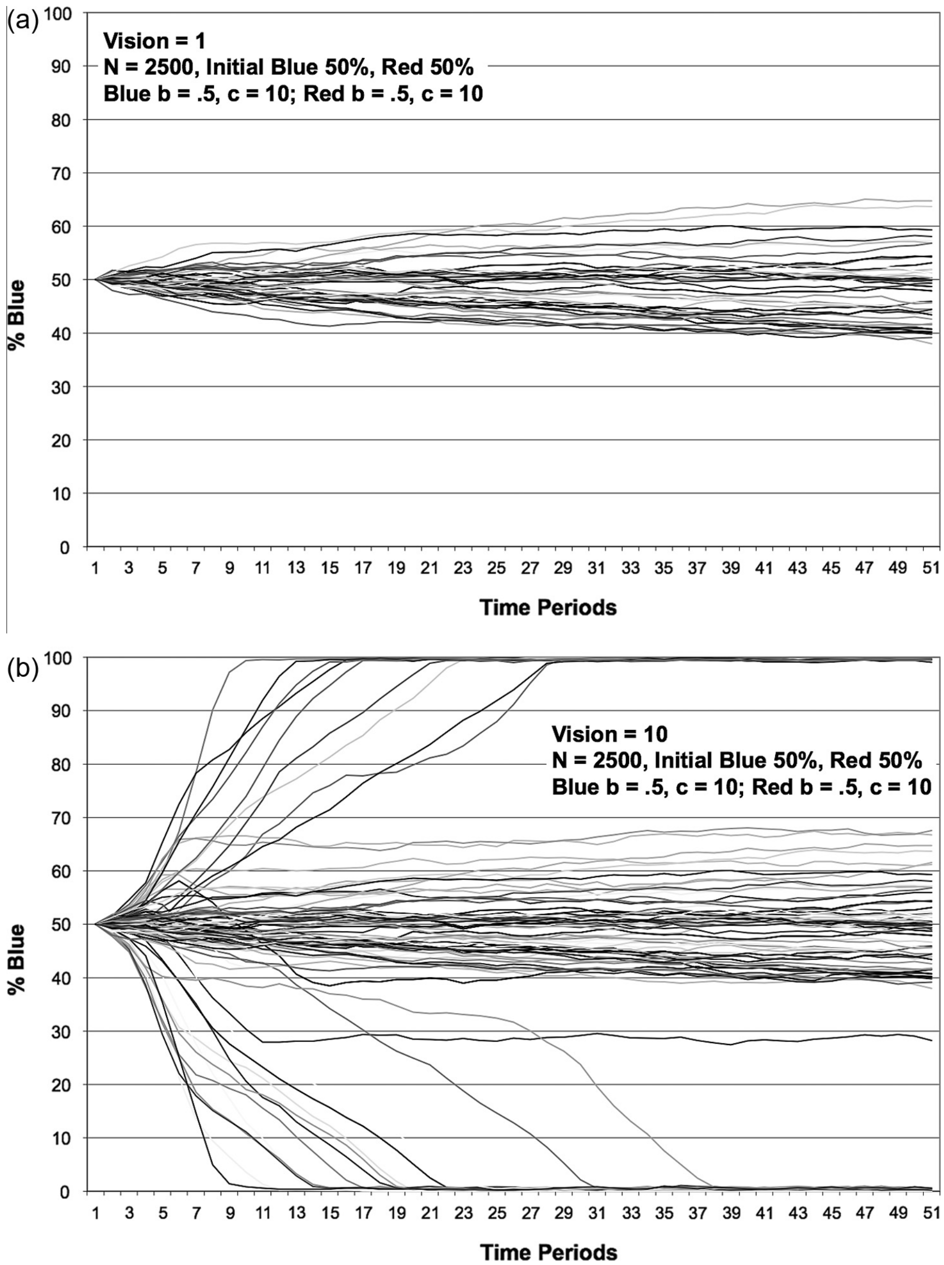

Fig. 5. Agent-based modeling trajectories (see MacCoun, 2012, p. 359): (a) Vision = 1 (i.e., each agent can see only the immediate 8 neighbors, N, NE, E, SE, S, SW, W, NW); (b) Vision = 10 (i.e., each agent can see 10 neighbors in each of the 8 cardinal directions).

Prospect theory (Kahneman & Tversky, 1979) offers what is arguably the dominant psychological model of evaluative choice. It proposes the following value function:

$$v(x) = \begin{matrix} |x-r|^{\alpha} & \text{if } x \geqslant r \\ \lambda|x-r|^{\beta} & x < r \end{matrix}$$

where r is a salient reference point (often either 0 or else the status quo), $\alpha, \beta \approx .88$ and the loss aversion parameter, λ, has been estimated at around 2.25. Taken alone (i.e., without a decision weighting function), this value function is a model of riskless choice (Tversky & Kahneman, 1991), with a deterministic choice rule: Select the prospect with the most favorable weighted value. The aBOP model is stochastic; it yields the predicted probability that the more favorable choice will be selected. But if we interpret a as the reference point (i.e., $a = r$), and let λ indicate the same loss aversion parameter, we get the following choice function:

$$p(choice) = \begin{matrix} \frac{1}{1+\exp[-k(x-r)]} & \text{if } x \geqslant r \\ \frac{1}{1+\exp[-\lambda k(x-r)]} & \text{if } x < r \end{matrix}$$

Although the equations look different, this model and prospect theory's riskless choice model can produce striking similar asymmetric *S*-shaped functions that drop more steeply in the domain where $x < r$.

An appealing feature of the aBOP (and bBOP) formalism is that it readily provides several different indices of the choice response. In addition to the propensity for a given choice (p), there is the threshold (a), the threshold's clarity (k), the threshold's asymmetry ($|a - .5|$), and the absolute distance between the stimulus value and the threshold ($|x - a|$). Conventionally, the latter value is taken as an index of confidence (Thomas & Hogue, 1976). It would be interesting to examine how these indices map onto meta-cognitive judgments like certainty, importance, intensity, ambiguity, and ambivalence (Visser, Bizer, & Krosnick, 2006).

For risky or uncertain prospects, prospect theory proposes a nonlinear decision-weighting function that gives extra weight to very small non-zero probabilities, and slightly underweights moderate to high probabilities below 1. Note that these weights, though probabilistic, do not change the deterministic nature of the theory's choice rule. The version of aBOP presented in this paper has no explicit representation of probability weights, but one can easily replace x_i with px_i or $w(p)x_i$, where $w(p)$ is some non-linear transformation of probabilities.

Thus, aBOP is closely related to the Luce choice model, the McFadden discrete choice model, signal detection theory, and prospect theory. Indeed, it is a useful framework for making more explicit how these major models of choice relate to each other. Of course, aBOP is not a detailed process model of how evidence is cognitively (and perhaps affectively) integrated to form choices, but it is not intended to provide a rival to other accounts. Rather, like bBOP, it is intended to provide (a) a useful skeletal structure for summarizing key parameters of the choice process, and (b) a useful estimation tool for inferring operative thresholds and their clarity given actual choice data.

Balancing Evidence and Norms (BEAN)

The aBOP model complements the bBOP model by providing an account of evidence evaluation that closely mirrors the account of norm evaluation. But this of course begs the question of how the two models are related and how they might jointly shape choice and action.

One obvious candidate for an integrative model would be a weighted average of the aBOP and bBOP models:

$$p(Option\ 1) = wp(Option\ 1|aBOP) + (1-w)p(Option\ 1|bBOP)$$

But this model has some significant drawbacks. It has an extravagant five free parameters (w, k, a, c, and b), making it desirable to find a more parsimonious account. And conceptually, by treating evidence evaluation and norm evaluation as independent processes, it ignores the possibility that the two are intertwined: The evidentiary foundation for one's position is likely to be a major impetus for resistance to social pressures from opponents, and hence the weight of the evidence (x_i) is likely to be an important input to the b threshold. Similarly, the knowledge that a position will be very unpopular (s_i) is likely to significantly raise the standard of evidence one demands to change one's position (a).

A second obvious candidate would be a standard logit model, with the proportion of evidence favoring Option 1 and the proportion of norms favoring Option 1 as independent variables:

$$p(Option) = \frac{\exp(w_0 + w_N s_1 + w_E x_0)}{1 + \exp(w_0 + w_N s_1 + w_E x_0)}$$

This model has only 3 free parameters and has the virtue of simplicity. It also allows one to estimate relative weights for norms and for evidence.

But an alternative parameterization, which I will call BEAN (for "Balancing Evidence and Norms"; see Table 2, row 4), retains the threshold character of the aBOP and bBOP models, and contains each of them as a special case.[4] To do so, first express evidence strength in terms of support for the position (Option 0) opposite that indexed by norm strength (Option 1); viz. $x_0 = 1 - x_1$. Second, replace the regression coefficients for evidence and norms with a single parameter expressing the clarity of the comparison process; viz., $\beta(s_1 - x_0 + \alpha)$. Finally, as with bBOP, the *ceiling* parameter allows for any resistance to Option 1 that is not captured by the threshold comparison process.

Although BEAN is more general than either aBOP or bBOP, it should not be considered more applicable or more useful. Judging from the current research literature, reflecting decades of study, few data sets have parametric information on both descriptive norms (s_1) and evidence (x_0); even fewer (if any) involve careful experimental manipulation of both factors. As a result, opportunities to operationalize BEAN (or the related logit model) are rare. The usefulness of its component models, aBOP and bBOP, is that they allow the analyst to infer something about whichever factor (perceived norms or evidence, respectively) is missing from a dataset.

Rather, BEAN's value is heuristic; it is a reminder of what is left out when we study evidence alone, or norms alone. When we have data on relative evidence strength, but none on norms, then BEAN reduces to aBOP, and its a threshold parameter "absorbs" the s variable – viz., $a = s + \alpha$. Likewise, when we have data on relative norm strength, but none on evidence, then BEAN reduces to bBOP, and its b threshold parameter absorbs the x variable – viz., $b = x + \alpha$. To the

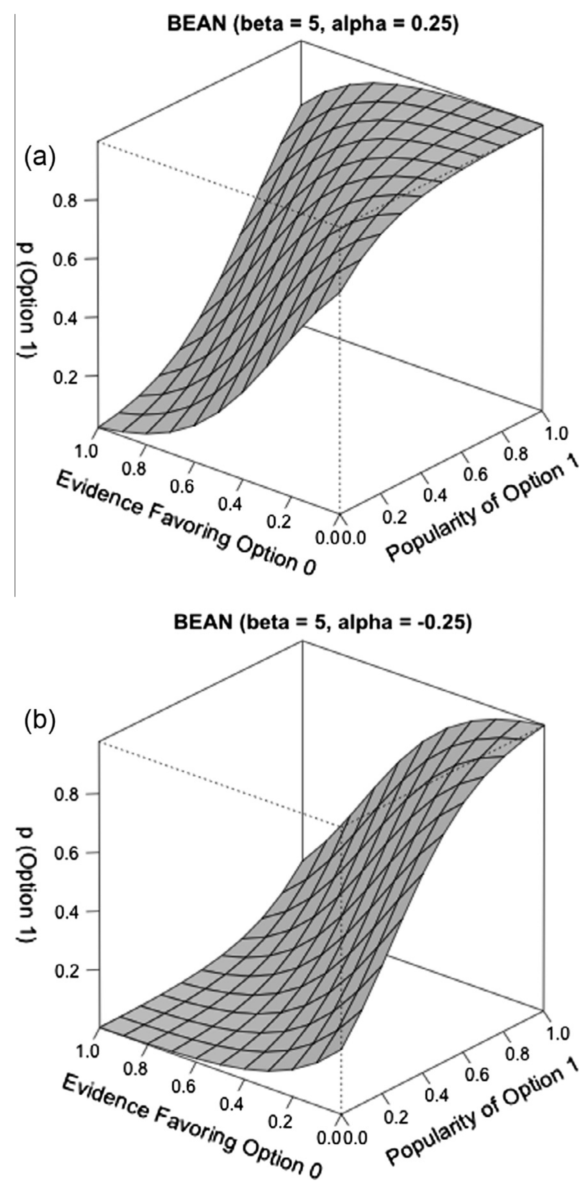

Fig. 6. Predictions of the BEAN model for (a) $\beta = 5$, $\alpha = .5$ and (b) $\beta = 5$, $\alpha = -.5$.

[4] This model is close but not identical to the logit formulation. The clarity parameter of BEAN corresponds to the average of the two logit weights, and the alpha parameter will be proportional to the difference in the two weights.

extent that we can operationalize both x and s, the α parameter will capture any additional "burden of proof" placed on Option 0 (when α is positive) or on Option 1 (when α is negative). We can think of α as a bias – a thumb on the scale – favoring Option 1 when it is positive (Fig. 6a) and Option 0 when it is negative (Fig. 6b).

The BEAN model has both similarities and important differences with the family of models in the "theory of reasoned action" (TRA) tradition (Ajzen & Fishbein, 1980). Both models suggest that norms combine with personal beliefs to influence behavioral propensities, and both suggest that the combination is additive/subtractive. However, models in the TRA tradition are linear whereas the BEAN model is non-linear; this in turn reflects the fact that the BEAN model predicts the probability of a choice (on a 0–1 metric) whereas TRA provides a scale of intention that is at best only correlated with choices. And the conception of norms is different; TRA's "subjective norms" are injunctive norms (what others think I should do) whereas BEAN (via bBOP) focuses on descriptive norms (what others actually do; Cialdini, Kallgren, & Reno, 1991). In a meta-analysis of 21 estimates, Rivis and Sheeran (2003) found that descriptive and injunctive norms were similarly predictive of intentions (weighted r's of .45 and .44, respectively) and the two types of norms were correlated with each other ($r = .38$).

Growing evidence from multiple research teams (Fischer et al., 2009; Shteynberg, Gelfand, & Kim, 2009; Zou et al., 2009; for a review and discussion see Morris, 2014) suggests that shared cultural biases are better explained by perceived descriptive norms (e.g., "Americans rise for the National Anthem") than shared values ("It is important to honor our country by rising for the National Anthem"). This idea might appear inconsistent with evidence from the TRA literature: For example, a meta-analysis of over one hundred estimates (Armitage & Conner, 2001) found that behavioral intentions are more strongly correlated with attitudes than with subjective norms (weighted $r = .49$ vs. .34), but unlike attitudes, norms were "typically measured by a single item" (p. 484), which would tend to attenuate correlations due to low reliability, and these studies assess injunctive rather than descriptive norms. Moreover, most of these studies are correlational, so the attitudes might themselves have been shaped by norms. Several lines of evidence suggest that the relative impact of attitudes vs. norms varies as a function of personality (Carver & Scheier, 1998; Finlay, Trafimow, & Moroi, 2006; Fu et al., 2007; Trafimow & Finlay, 1996), nation (Gelfand et al., 2011; Savani, Morris, & Naidu, 2012), and situational salience and framing (Carver & Scheier, 1998; Cialdini et al., 1991; Pillutla & Chen, 1999).

Comparison to other models of cultural evolution

Boyd and Richerson (1985, 2005)

What do the present models offer relative to existing models in the cultural evolution literature? In this section, I attempt to show how a logistic threshold account can capture important features of other models, while also accounting for observable patterns they do not readily describe.

Perhaps the most prominent account of cultural evolution is the large and multifaceted body of theory and research by Boyd and Richerson (1985, 2005) and their colleagues (e.g., Henrich, 2000; Henrich & McElreath, 2003; McElreath et al., 2008). In several decades of work, they have conducted pathbreaking analyses of a variety of central questions in cultural evolution, using a mix of new formal theories, interpretations of classic and new field data, and even some laboratory experiments. Here I examine a small portion of their work that most directly pertains to the tension

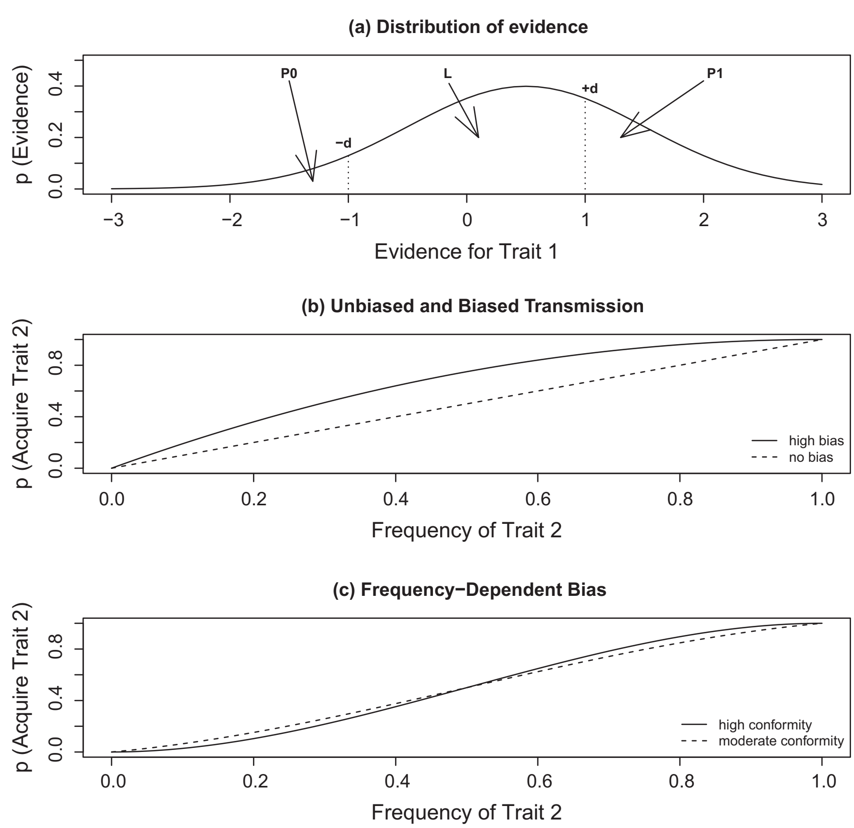

Fig. 7. Boyd and Richerson's (1985, 2005) theory: (a) Distribution of evidence and probabilities of Option 0 (P_0), imitation (L), and Option 1 (P_1); (b) Probability of acquiring Option 1 under unbiased and biased transmission; and (c) Probability of acquiring Option 1 under frequency-dependent transmission.

between individual learning and social imitation, but exclude any discussion of their important efforts to model the relative adaptive benefits of alternative learning strategies or the interplay between cultural and biological evolution.

The Boyd and Richerson model BR1 (my label for convenience) in Table 3 gives the simple unbiased form of transmission, using notation defined in Table 1. Their account begins with a prototypical situation in which each individual draws at random an environmental signal about the relative value of two alternative behaviors or traits (see Fig. 7a). The individual has two thresholds for assessing evidence, such that the individual will adopt Option 1 when $x > +d$, and Option 0 when $x < -d$. The two thresholds are symmetrical around the 0 point (rather than, say, around the mean of the distribution). These thresholds create three regions. The area in region P_1 gives the probability that an individual will choose Option 1, the area P_0 gives the probability that an individual will choose Option 0, and the area of the middle region L gives the probability that the individual will be undecided. Boyd and Richerson (2001) argue that the area of region L gives the probability that the actor, failing to reach a decision on the evidence, will imitate a randomly chosen individual. Because $L = 1 - P_1 - P_2$, this implies that the probabilities P_1 and P_0 are "subadditive"—they can sum to less than 1.

The dual-threshold aspect of Boyd and Richerson's theory is somewhat reminiscent of Shafer's (1976) theory of belief functions, in which belief in a proposition and belief in its negation need not sum to 1.[5] The use of two thresholds also bears some resemblance to Sherif and Hovland's (1961) proposal that the attitude continuum is segmented into distinct latitudes of acceptance and rejection and non-commitment. But its not clear that two thresholds are needed for their account; if we replot the normal distribution in Fig. 7a as a cumulative probability distribution, we would find a sigmoid shape, with a single inflection point, quite like the BOP models. (BOP's logistic has somewhat fatter tails than a normal distribution, which is arguably a good feature for a system that adapts well to uncertain evidence; see Bentley & O'Brien, 2011; Bentley, Ormerord, & Batty, 2011.)

Henrich (2000) shows that this BR1 model, taken alone, will always produce an *r*-shaped or concave growth pattern, and is hence inconsistent with the typical *S*-shaped growth seen in empirical studies of the diffusion of innovations. To produce the typically seen *S*-shaped growth curves, the model must incorporate *biased* transmission, as seen in model BR2 in Table 3. *Direct-biased transmission* occurs when some feature of a trait makes it intrinsically more attractive to us due to some feature of our individual psychology. *Prestige-biased transmission* occurs when we are more likely to adopt a given trait because certain specific individuals with high-status traits are seen to model it.

For Boyd and Richerson, *frequency-dependent (or conformist) transmission* occurs when we are differentially attracted to whichever trait is favored by a majority of others. This is captured in model BR3 in Table 3, which replaces the direct bias term (B) with a weighted average of B and a frequency-dependent term, $2s - 1$. Note that the latter term is below 0 when opponents are in the minority, 0 when both factions are equal in size, and above 1 when opponents are in the majority.

Even this abbreviated presentation results in a complex model, but for present purposes, there are three qualitative predictions of interest, shown in Fig. 7b and c (adapted from Boyd & Richerson, 1985, Fig. 7.1). First, unbiased transmission produces a simple 45-degree response function – the dashed line in Fig. 7b. This is identical to the *proportionality* decision scheme in social decision scheme theory (Davis, 1973; Kerr et al., 1996), as well as the predictions of Mullen's (1983) *other-total ratio* model. MacCoun (2012) shows that this proportionality function is useful mostly as a baseline; it does not adequately describe actual social influence data. Second, direct-biased transmission produces a slightly concave function – the solid line in Fig. 7b. This is similar to the pattern seen in the Milgram et al. (1969) data in Fig. 3a. And third, frequency-dependent transmission produces a mild sigmoid function as seen in Fig. 7c.

With the right parameter values, bBOP and BEAN produce the same qualitative features shown in Fig. 7b and c. But the Boyd–Richerson models differ in several ways. First, notwithstanding the probabilistic nature of the distribution of evidence (Fig. 7a), the BR2 and BR3 models are deterministic whereas the BOP and BEAN models generate probabilities, or what might be called behavioral propensities. Second, the BOP and BEAN models readily produce asymmetric thresholds, a feature that is readily apparent in actual group behavior (see Figs. 3 and 4). I return to this latter point below.

McElreath et al. (2008)

McElreath et al. (2008) have formalized the Boyd–Richerson models differently, shifting from the state-updating approach of growth models to probabilistic formulations that are similar to BOP and BEAN. MEA1 (again, my label) in Table 2 shows their model of unbiased choice under individual learning (with notation defined in Table 1). They refer to this as a "softmax" rule using a common label in computational modeling; psychologists will recognize this as an expression of the Luce (1959) choice rule. Taken alone, this model produces behavioral propensities that are equivalent to aBOP with $a = .5$, which might be called a "mere preponderance of evidence" rule, where their λ parameter produces the same range of soft to hard thresholds as BOP. Fig. 2a is identical to the behavior of MEA1, if one substitutes lambda for c and evidence strength for norms. But their model does not allow for an asymmetric threshold, as might be expected under "reasonable doubt" or any conceptual scheme that is more skeptical toward evidence favoring one of the two options.

MEA2 is their formalization of frequency-dependent learning. Ignoring the exponents, it is a simple proportionality model (see Mullen, 1983), which provides only a rough approximation of the social influence patterns observed empirically (MacCoun, 2012). But the exponents allow the model to produce a range of behavior almost identical to MEA1 (and bBOP with $b = .5$ in Fig. 2a). Surprisingly, they do not mention this similarity to MEA1; and since they did not derive the model deductively, it is curious that they didn't choose the same softmax formulation as MEA1 to simplify their account.[6] Like MEA1, this model lacks any parameter permitting the threshold shifts that produce the more asymmetric pattern shown in Fig. 2b and commonly observed in empirical studies of the effects of faction size. Finally, MEA3 is a weighted average model that combines MEA1 and MEA2. They show that this hybrid model outperformed either MEA1 or MEA2, although a more complex hybrid that also involved a comparison of mean outcomes did slightly better.

Like BEAN, MEA3 has 3 free parameters, and the two models are capable of producing similar soft or hard thresholds. But BEAN (like aBOP and bBOP) has a more direct mechanism for allowing *shifting* or asymmetric thresholds. A rigorous test of these alternative models would require experimental data in which evidence strength and norm strength are independently manipulated. To my knowledge, no such experiments have been conducted, and doing so con-

[5] In Shafer (1976), $bel(A) \leqslant p(A) \leqslant pl(A)$ and $pl(A) = 1 - bel(\text{notA})$, where *bel()* is a belief function and *pl()* is a plausibility function.

[6] Unlike the softmax rule, MEA2 includes a simple proportionality rule (when $f = 1$), but it is not clear why that is needed to model norms but not evidence.

vincingly would be challenging because these two factors tend to be correlated in ordinary social situations.

As a more provisional test, I fit MEA3 to the data from the Asch (1956) conformity study and the Kerr and MacCoun (1985) jury deliberation study. Since neither study manipulates evidence strength, I allowed A_1 to vary as a fourth free parameter (letting $A_0 = 1 - A_1$). For the Asch data, the best fitting parameters were $A_1 = 0$, $\lambda = 12.86$, $f = 1.94$, and $\gamma = .30$. That $A_1 = 0$ is consistent with Asch's (1956) finding that, when participating alone, his respondents selected correct responses about 99 percent of the time. The estimate of $\gamma = .30$ suggests that participants were mostly under the influence of individual learning rather than conformist learning; this is also close to bBOP's estimated ceiling parameter of $m = .325$ for this study, which has a similar interpretation. But the MEA3 model did not fit the Asch data very well ($r^2 = .78$, adjusted $r^2 = .77$), especially relative to the fit of the bBOP model (r^2 and adjusted $r^2 = .99$).

MEA3 fared somewhat better in fitting the Kerr and MacCoun (1985) mock jury data, with $A_1 = 0$, $\lambda = 12.68$, $f = 2.81$, and $\gamma = .87$ ($r^2 = .88$, adjusted $r^2 = .88$). As in the Asch case, the fit was best when $A_1 = 0$, but in this study that makes less sense because the trial evidence was fairly closely balanced (see Kerr & MacCoun, 1985). As with the Asch data, the bBOP model provided a better fit (r^2 and adjusted $r^2 = .99$).

Szabó and Töke (1998), Traulsen, Pacheco, and Nowak (2007)

Several models in the cultural evolution literature have proposed social learning rules that appear similar to both BOP and to the McElreath et al. (2008) models. For example, Szabó and Töke (1998) proposed that people engage in a pairwise comparison process by judging the recent payoff of their own choice to that received by a randomly chosen neighbor:

$$p = 1/(1 + \exp(-\beta[\pi_A - \pi_A])$$

They refer to this as a Fermi equation, and offer a statistical physics interpretation of β as "temperature." Traulsen et al. (2007) use the same basic model with different notation. This model is mathematically similar to the BOP models, but focuses on the difference between two outcomes rather than the difference between some value and a personal threshold. It is similar to bBOP in that a randomly chosen individual has the collective's average as its expected value; it is similar to aBOP in modeling evidence (payoffs) rather than positions (popularity). Their pairwise comparison of outcomes is of course very reminiscent of models in the distributive justice and social exchange traditions in psychology.

An assessment and some future directions

The BOP and BEAN models share some essential features with models in the Boyd and Richerson program. First, all these models allow for "direct bias" favoring one option over the other irrespective of evidence and/or norms. In BOP/ BEAN, direct bias takes the form of an asymmetry in the threshold (i.e., $|a - .5|$ in aBOP, $|b - .5|$ in bBOP, $\alpha < 0$ or $\alpha > 0$ in BEAN). Henrich (2000) suggests that models should also allow for "prestige bias," in which case the bias is linked to an influence source rather than the position he or she advocates. Both direct and prestige bias are consistent with the idea of a shared conceptual scheme for recognizing the better option; in the case of prestige, there's a shared conceptual scheme (e.g., chain of command, seniority, caste, reputation) that makes some advocates carry more weight. But an important difference between a direct bias and the notion of a shared conceptual scheme is that the former leaves unexplained why everyone with the bias doesn't *already* hold the preferred position. A shared conceptual scheme allows one or two people to quickly persuade the rest of us of something we initially failed to recognize.

Second, the models all produce a sigmoid "frequency-dependent" conformity response. The sigmoid shape implies that factions can have disproportionately more or less influence than one would predict from their relative size. But the BOP models readily allow for considerable asymmetry in the threshold, and I have argued that this asymmetry is an essential feature of actual choice and influence behavior (also see Kerr & MacCoun, 2012; MacCoun & Kerr, 1988).

The clarity parameters in these threshold models allow for change that is gradual (when thresholds are soft) or abrupt (when they are sharp). As MacCoun (2012) notes, sharp thresholds imply highly correlated behavior across actors; the fact that thresholds can be soft or hard helps to explain why some cultural changes appear to be "tipping points" but others play out nearly undetected over longer time courses. If you and I share the same threshold (for evidence in aBOP or for norms in bBOP), we will "tip" at roughly the same moment in response to new physical or social information in the environment.

The cultural evolution models of Boyd and Richerson (1985, 2005) and their colleagues (e.g., Henrich, 2000; McElreath et al., 2008) are deduced systematically from reasonable assumptions. Their work is rigorous, transparent, and quite fecund, having enabled them to analyze a great many important questions about individual learning, cultural evolution, adaptation to ecosystems, and the relationship between cultural and biological evolution. While the analyses presented here don't even begin to capture the ambition and reach of the full Boyd and Richerson program, I believe the bBOP model, and perhaps the aBOP and BEAN models, improve upon some components of the Boyd–Richerson framework, and they do so in a way that retains the notions of direct- and frequency-dependent bias.

Theoretically, the BOP models are closely related to other choice models, including signal detection theory (Swets, 1988), the Luce (1959) choice model, and social decision scheme theory (Davis, 1973). The bBOP model integrates important insights from Schelling (1978), Granovetter (1978), Latané (1981) and other theorists, while providing a good fit to data from the conformity, deliberation, and diffusion of innovation paradigms. Finally, aBOP and bBOP are transformed logit models, and relatively easy to use as estimation tools in empirical applications. They also provide simple response rules that facilitate Monte Carlo and agent-based simulations.

The BEAN model is more difficult to apply empirically than either aBOP or bBOP, because it requires more information. It is fairly simple to implement an experiment that parametrically varies the evidence strength modeled in aBOP. It is more difficult to parametrically vary the "strength in numbers" modeled in bBOP. But to vary both types of influence in a single experiment, in a convincing fashion, could be quite challenging. And in non-experimental settings, it is possible to estimate BEAN's parameters, but they must be interpreted with caution, as evidence strength and strength in numbers may confounded with each other or with other variables.

The strength in arguments (or evidence) and the strength in numbers do not exhaust the list of influences on choice, by any means. As noted, Henrich (e.g., 2000) discusses the important role of prestige. French and Raven (1960) distinguish other forms of "social power," including coercive power, reward power, reference power, legitimate power, and expert power. The BEAN model could be expanded to incorporate such features by further parsing the α parameter, at least under the assumption that they are additively separable components,.

In this article I have drawn links between the cultural evolution literature and the social influence tradition in social psychology. But similar links should be established to the diffusion of innovation literature in marketing and in science and technology studies (Mahajan et al., 1995). The Bass model, and the BR3 and MEA3 models examined here share in common the use of weights representing the relative impact of individual learning and social learning. In the Bass model, these weights are known as the coefficients of innovation and imitation, respectively, but they are not required to sum to 1. Mahajan et al. (1995) note that these coefficients are generally estimated at around .03 and .30–.50. If so, that would be equivalent to a weight of .056 (viz., .03/(.03 + .30)) to .091 on individual learning and a weight of .909 to .943 on frequency-dependent social learning in the cultural evolution models – a fairly strong pair of empirical constraints. To the extent that the estimated weights of the McElreath models depart from these constraints, it might suggest that there is an importance difference between the assumptions of the two modeling traditions. But given the strictly correlational nature of most diffusion studies, it is possible that the imitation term is picking up other variables (see MacCoun, 2012, p. 355).

It is exciting that a growing community of scholars in the cultural evolution tradition (including Bentley & O'Brien, 2011; Bentley et al., 2011; Boyd & Richerson, 1985, 2005; Hastie & Kameda, 2005; Henrich, 2000; McElreath et al., 2008; Szabó & Töke, 1998; Traulsen et al., 2007) have demonstrated that "grand" questions about the human experience can be addressed in an analytically rigorous way that incorporates contributions from many different disciplines. It is hoped that the models sketched here contribute to that enterprise and illustrate how to build closer links to theory and research in cognitive and social psychology. Granted, all such models provide a grossly oversimplified portrayal of the rich and complex tapestry of an evolving culture, but hopefully they illuminate some important features.

Acknowledgments

I thank Katrina Fincher, Barbara Mellers, Michael Morris, Daria Roithmayr, Sydney Scott, Phil Tetlock, and two anonymous reviewers for comments and advice.

References

Ajzen, I., & Fishbein, M. (1980). *Understanding attitudes and predicting social behavior*. Englewood Cliffs, NJ: Prentice-Hall.
Armitage, P. C., & Conner, M. (2001). Efficacy of the theory of planned behaviour: A meta-analytic review. *British Journal of Social Psychology, 40*, 471–499.
Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs, 70* (Whole no. 416).
Bass, F. (1969). A new product growth model for consumer durables. *Management Science, 15*, 215–227.
Bentley, R. A., & O'Brien, M. J. (2011). The selectivity of social learning and the tempo of cultural evolution. *Journal of Evolutionary Psychology, 9*, 125–141.
Bentley, R. A., Ormerord, P., & Batty, M. (2011). Evolving social intelligence in large populations. *Behavioral Ecology and Sociobiology, 65*, 537–546.
Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
Boyd, R., & Richerson, P. J. (2005). *The origin and evolution of cultures*. Oxford.
Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. New York: Cambridge University Press.
Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton.
Chang, S. W. C., Gariepy, J., & Platt, M. J. (2013). Neuronal reference frames for social decisions in primate frontal cortex. *Nature Neuroscience, 16*, 243–250.
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.
Cooksey, R. W. (1996). *Judgment analysis: Theory, methods, and applications*. Bingley, UK: Emerald Group.
Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences, 13*, 36–43. http://dx.doi.org/10.1016/j.tics.2008.10.002.
Cumming, G. (2014). The new statistics: Why and how. *Psychological Science, 25*, 7–29.
Davis, J. H. (1973). Group decision and social interaction: A theory of social decision schemes. *Psychological Review, 80*, 97–125.
DeCarlo, L. T. (1998). Signal detection theory and generalized linear models. *Psychological Methods, 3*, 186–205.
Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *Journal of Abnormal Social Psychology, 51*, 629–636.
Eagly, A., & Chaiken, S. (1993). *The psychology of attitudes*. Fort Worth, TX: Harcourt Brace Jovanovich.
Epstein, J. (1975). *Jacob Epstein: An autobiography*. Arno Press.
Festinger, L. (1954). A theory of social comparison processes. *Human Relations, 7*, 117–140.
Fidler, F., Thomason, N., Cumming, G., Finch, S., & Leeman, J. (2004). Editors can lead researchers to confidence intervals, but they can't make them think: Statistical reform lessons from medicine. *Psychological Science, 15*, 119–126.
Finlay, K. A., Trafimow, D., & Moroi, E. (2006). The importance of subjective norms on intentions to perform health behaviors. *Journal of Applied Social Psychology, 29*, 2381–2393.
Fischer, R., Ferreira, M. C., Assmar, E., Redford, P., Harb, C., Glazer, S., et al. (2009). Individualism-collectivism as descriptive norms: Development of a subjective norm approach to culture measurement. *Journal of Cross-Cultural Psychology, 40*, 187–213.
Fisher, R. A. (1928). *Statistical methods for research workers*. London: Oliver & Boyd.
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: Harper and Row.
Fu, J. H. Y., Morris, M. W., Lee, S. L., Chao, M., Chiu, C. Y., & Hong, Y. Y. (2007). Epistemic motives and cultural conformity: Need for closure, culture, and context as determinants of conflict judgments. *Journal of Personality and Social Psychology, 92*, 191–207.
Gelfand, M. J., Raver, J. L., Nishii, L., Leslie, L. M., Lun, J., Lim, B. C., et al. (2011). Differences between tight and loose cultures: A 33-nation study. *Science, 332*, 1100–1104.
Granovetter, M. (1978). Threshold models for collective behavior. *American Journal of Sociology, 83*, 1420–1443.
Hastie, R., & Kameda, T. (2005). The robust beauty of majority rules in group decisions. *Psychological Review, 112*, 494–508.
Henrich, J. (2000). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased transmission is the predominant force in cultural change. *American Anthropologist, 103*, 992–1013.
Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology, 12*, 123–135.
Hoppitt, W., & Laland, K. N. (2013). *Social learning: An introduction to mechanisms, methods, and models*. Princeton University Press.
Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–291.
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.
Kerr, N. L., & MacCoun, R. J. (2012). Is the leniency asymmetry really dead? Misinterpreting asymmetry effects in criminal jury deliberation. *Group Processes and Intergroup Relations, 15*, 585–602.
Kerr, N., MacCoun, R. J., & Kramer, G. (1996). Bias in judgment: Comparing individuals and groups. *Psychological Review, 103*, 687–719.
Latané, B. (1981). The psychology of social impact. *American Psychologist, 36*, 343–356.
Laughlin, P. R. (2011). *Group problem solving*. Princeton University Press.
Leary, M. R., & Baumeister, R. F. (2000). The nature and function of self-esteem: Sociometer theory. In M. P. Zanna (Ed.), *Advances in experimental social psychology (Vol. 32)* (pp. 1–62). San Diego, CA, US: Academic Press.
Lewin, K., Dembo, T., Festinger, L., & Sears, P. (1944). Level of aspiration. In J. Hunt (Ed.), *Personality and the behavior disorders* (pp. 333–378). Oxford: Ronald Press.
Lorge, I., & Solomon, H. (1955). Two models of group behavior in the solution of Eureka-type problems. *Psychometrika, 20*, 139–148.
Luce, R. D. (1959). *Individual choice behavior*. NY: Wiley.
MacCoun, R. J. (2012). The burden of social proof: Shared thresholds and social influence. *Psychological Review, 119*, 345–372.
MacCoun, R. J. (2014). Alternative maps for the world of collective behavior. *Behavioral and Brain Sciences, 37*, 88–90.
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.
Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology, 28*, 143–166.
Mahajan, V., Muller, E., & Bass, F. (1995). Diffusion of new products: Empirical generalizations and managerial uses. *Management Science, 14*, G79–G88.
McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J., & Waring, T. (2008). Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B, 363*, 3515–3528.
McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics* (pp. 105–142). NY: Academic Press.
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.
Morris, M. W. (2014). Values as the essence of culture: Foundation or fantasy? *Journal of Cross-Cultural Psychology, 45*, 14–24.
Mullen, B. (1983). Operationalizing the effect of the group on the individual: A self-attention perspective. *Journal of Experimental Social Psychology, 19*, 295–322. http://dx.doi.org/10.1016/0022-1031(83)90025-2.

Nowak, A., Szamrej, J., & Latané, B. (1990). From private attitudes to public opinion: A dynamic theory of social opinion. *Psychological Review, 97*, 362–376.

Parks, C. D., & Kerr, N. L. (1999). Twenty-five years of social decision scheme theory. *Organizational Behavior and Human Decision Processes, 80*, 1–2.

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

Pillutla, M., & Chen, X. P. (1999). Social norms and cooperation in social dilemmas: The effects of context and feedback. *Organizational Behavior and Human Decision Processes, 78*, 81–103.

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

Rendell, L., Fogarty, L., Hoppitt, W. J. E., Morgan, T. J. H., Webster, M. M., & Laland, K. N. (2011). Cognitive culture: Theoretical and empirical insights into social learning strategies. *Trends in Cognitive Sciences, 15*, 68–76.

Rivis, A., & Sheeran, P. (2003). Descriptive norms as an additional predictor in the theory of planned behaviour: A meta-analysis. *Current Psychology, 22*, 218–233.

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

Savani, K., Morris, M. W., & Naidu, N. V. R. (2012). Deference in Indians' decision making: Introjected goals or injunctive norms? *Journal of Personality and Social Psychology, 102*, 685–699.

Schelling, T. C. (1978). *Micromotives and macrobehavior*. W.W. Norton and Co.

Shafer, G. (1976). *A mathematical theory of evidence*. Princeton University Press.

Sherif, M., & Hovland, C. I. (1961). *Social judgement*. New Haven, CT: Yale University Press.

Shteynberg, G., Gelfand, M. J., & Kim, K. (2009). Peering into the "Magnum Mysterium" of culture: The explanatory power of descriptive norms. *Journal of Cross-Cultural Psychology, 40*, 46–69.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63*, 129–138.

Swets, J. A. (1988). Measuring the accuracy of diagnostic systems. *Science, 240*, 1285–1293.

Szabó, G., & Töke, C. (1998). Evolutionary prisoner's dilemma game on a square lattice. *Physics Review E, 58*, 69–73.

Thomas, E. A. C., & Hogue, A. (1976). Apparent weight of evidence, decision criteria, and confidence ratings of juror decision making. *Psychological Review, 83*, 442–465.

Trafimow, D., & Finlay, K. A. (1996). The importance of subjective norms for a minority of people: Between subjects and within-subjects analyses. *Personality and Social Psychology Bulletin, 22*, 820–828.

Traulsen, A., Pacheco, J. M., & Nowak, M. A. (2007). Pairwise comparison and selection temperature in evolutionary game dynamics. *Journal of Theoretical Biology, 246*, 522–529.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *Quarterly Journal of Economics, 106*, 1039–1061.

Visser, P. S., Bizer, G. Y., & Krosnick, J. A. (2006). Exploring the latent structure of strength-related attitude attributes. *Advances in Experimental Social Psychology, 38*, 1–67.

Zou, X., Tam, K.-P., Morris, M. W., Lee, S.-L., Lau, I. Y.-M., & Chiu, C.-Y. (2009). Culture as common sense: Perceived consensus versus personal beliefs as mechanisms of cultural influence. *Journal of Personality and Social Psychology, 97*, 579–597.