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	$p(Option\ 1)_{FD} = n^f_{1,t}/(n^f_{1,t} + n^f_{0,t})$
	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.