



Fitness landscapes among many options under social influence

Camila C.S. Caiado^a, William A. Brock^{b,c}, R. Alexander Bentley^{d,*}, Michael J. O'Brien^e

^a Department of Mathematical Sciences, Durham University, Durham, UK

^b Department of Economics, University of Missouri, Columbia, MO 65211, USA

^c Department of Economics, University of Wisconsin, Madison, WI 53706, USA

^d Department of Comparative Cultural Studies, University of Houston, Houston, TX 77204, USA

^e Department of Anthropology, University of Missouri, Columbia, MO 65211, USA

HIGHLIGHTS

- A learning model and fitness-landscape function; agents choose from N options.
- Its three key factors: social learning, transparency of choice, change through time.
- Our hill-climbing algorithm finds expected optimal decisions as landscape peaks.
- Multiple equilibria at each point on landscape, which is rugged even for $N=3$ choices.
- Initial conditions, path dependence underlie optimal behavior among social organisms.

ARTICLE INFO

Article history:

Received 2 June 2015

Received in revised form

14 December 2015

Accepted 17 December 2015

Keywords:

Discrete choice

Fitness landscape

Individual learning

Payoffs

Social learning

ABSTRACT

Cultural learning represents a novel problem in that an optimal decision depends not only on intrinsic utility of the decision/behavior but also on transparency of costs and benefits, the degree of social versus individual learning, and the relative popularity of each possible choice in a population. In terms of a fitness-landscape function, this recursive relationship means that multiple equilibria can exist. Here we use discrete-choice theory to construct a fitness-landscape function for a bi-axial decision-making map that plots the magnitude of social influence in the learning process against the costs and payoffs of decisions. Specifically, we use econometric and statistical methods to estimate not only the fitness function but also movements along the map axes. To search for these equilibria, we employ a hill-climbing algorithm that leads to the expected values of optimal decisions, which we define as peaks on the fitness landscape. We illustrate how estimation of a measure of transparency, a measure of social influence, and the associated fitness landscape can be accomplished using panel data sets.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Imagine a human decision scenario, modern or ancient, such as a person choosing a cereal at a grocery store or a prehistoric forager deciding which stand of trees to visit to gather hazelnuts. We tend to think of the former as economics and the latter as human ecology, but in each case, the decision has many similar options and depends on (1) the transparency of how good each option is, (2) the intrinsic utility of each option, and (3) the social utility of each option.

As researchers, we can observe the proportion of individuals who choose each option and, based on that information, attempt to infer these three quantities. Leaving social utility aside for a

moment, consider just transparency and intrinsic utility. We expect that if the intrinsic utility of each choice is highly transparent, then the probability distribution of decisions is in good accord with the fitness landscape, and there will be a single peak at the highest-utility option. As transparency decreases, the probability distribution flattens out as the fitness landscape becomes less visible and utility differences can no longer be discerned among the different options. At zero visibility, the probability distribution approaches a uniform distribution, and we effectively have random choice.

Now add back in social utility. For example, a shopper chooses the brand that he just saw someone else choose, or perhaps a forager follows her kinfolk to a particular stand of trees. Then aggregate those decisions over time and/or people. With social utility added to the mix, herding effects are possible, and the most popular option among the aggregated observations need not have

* Corresponding author.

E-mail address: rabentley@uh.edu (R.A. Bentley).

the highest utility. Indeed, if social utility is high and intrinsic transparency low, the respective distributions of choice probability and intrinsic utility among the options could differ significantly.

This exemplifies the complexity of modeling a fitness landscape of discrete choice under social influence. Here we build on previous work by examining cases where there may be multiple observed distributions of choices for transparency, utility, and social influence. We sketch an approach to estimation of the fitness landscape in the presence of multiple equilibria. This approach builds on a two-dimensional map we recently presented to track decision making as it relates to learning and transparency (Bentley et al., 2014; Brock et al., 2014).

To place our study in a broader context, we summarize below some of the recent work that has been done in the area of cultural learning, given that it forms the foundation of the horizontal axis of our map. We stress that the econometric studies we borrow from do not deal with the actual estimation of fitness functions, nor do they deal with the computation of equilibria or provide a theory of *which* equilibria are likely to be observed when actual estimation is conducted in the presence of multiple equilibria. Addressing these issues represents our contribution to the formulation of fitness functions and their estimation. We necessarily use terms such as “transparency,” “social conformity,” and “social interactions,” which are subject to the imprecision of words in contrast to the precision of the mathematical concepts for these terms that we develop later.

2. Social influence: a key element in decision making

When agents are faced with making a decision that involves multiple options, they can do one of two things. They can either learn individually, where they attempt to think things through by themselves, or they can learn socially by using other agents as sources of information. Within any population, the precise mixture of individual, or independent (asocial), learners versus social learners—a dichotomy sometimes referred to as information “producers” versus information “scroungers” (Mesoudi, 2008; Rendell et al., 2011)—may be crucial to a group's ability to climb a rugged fitness landscape (Rogers, 1995; Mesoudi and Whiten, 2008; Rendell et al., 2010; O'Brien and Bentley, 2011; O'Brien et al., *in press*). The reason for this is that whereas social learning spreads behaviors, it depends on individual learning to generate them in the first place. The question is, how does an agent integrate social and individual learning (Perreault et al., 2012), and how do their collective decisions affect fitness? Several studies have examined this question (e.g., Giraldeau et al., 2002; Kendal et al., 2009), many building on the work of Rogers (1988), who proposed that environmental change lowers group fitness when social learners copy outdated environmental information (Enquist et al., 2007; Rendell et al., 2011; Rieucou and Giraldeau, 2011). If the environment does not change, group fitness tends to increase as social learners copy optimal behaviors. Similarly, natural selection favors agents who place heavy weight on social cues when the environment changes slowly or when its state cannot be well predicted using individual learning (Perreault et al., 2012).

A population will ideally contain an optimally adaptive mix of the two learning strategies, but there is no assurance that this optimal mix will occur, as other steady-state mixes might exist. Numerous studies suggest that about 5% of informed individuals are enough to guide a social group (e.g., schooling fish) to a destination (Dyer et al., 2009; Herbert-Read et al., 2013; Wolf et al., 2013; Kurvers et al., 2014). Among that minority, this “pied piper” effect is augmented by intensity of direction (Couzin et al., 2011), which we might generalize as the “intensity of choice” (Bentley et al., 2014), or the accumulation of knowledge (Gomes, 2006).

In traditional human societies, social learning is usually transparent, as experts in different essential categories of adaptive knowledge (medicinal plants, hunting, fishing, cultivation) are well known to the group members (Henrich and Broesch, 2011). Over generations, well-directed social learning increases collective knowledge—teachers to students, parents to children, experts to general communities. As a consequence, the benefits of social learning are substantial enough for it to have been a key factor in human evolution (Hruschka, 2010; Hoppitt and Laland, 2013; Christakis and Fowler, 2014). Small groups can outperform even the most skilled/knowledgeable individual on complex tasks (Woolley et al., 2010) and in remembering information (Clément et al., 2013).

Transparency is not assured, however. Even among social animals, if misinformation invades the social-learning process, it can spread (Couzin et al., 2005). As information spreads between, say, Facebook or Twitter friends (Aral et al., 2009; Bond et al., 2012; Garcia-Herranz et al., 2014), expertise is not necessarily transparent to all members of the networks. In cases where expertise is not transparent, a good strategy might be to copy recent success (Laland, 2004; Rendell et al., 2010). Schools and flocks may be seen as “copying the recent”: when flocking agents are copying their neighbors' current direction of travel, the information is available practically instantaneously (Couzin et al., 2005).

An empirical challenge is in characterizing social-learning strategies from data aggregated at a broader scale than individual-agent motivations. For example, social-psychology experiments (e.g., Salganik et al., 2006; Lorenz et al., 2011) show that providing information about what others are doing often reduces the diversity of independent judgments within trials but increases variance between trials, thereby reducing the predictability and accuracy of the aggregated mean of those judgments.

Social influence is best recorded by close observation of each agent and its interactions through time (Hobaiter et al., 2014). If the observational data are more aggregated, however, it is difficult to demonstrate social learning without resorting to strong a priori assumptions (Shalizi and Thomas, 2011; Thomas, 2013). These sorts of aggregated datasets are common, but can they be used to distinguish between genuine social influence and individual discovery? The “three-degrees-of-influence” hypothesis concerning behaviors that spread within human social networks beyond one's immediate friends (Christakis and Fowler, 2013) can also be explained by simple autocorrelation through individual discovery combined with homophily—the tendency for individuals with similar traits to co-associate (Brock and Durlauf, 2001; Aral et al., 2009; Thomas, 2013). We return to the issue of homophily later.

3. A bi-directional map of decision making

Following our discussion above, we focus on two important factors, or “dimensions,” in terms of how decisions are made in the face of multiple options: the magnitude of social influence in the learning process and the transparency of costs and payoffs to either social learning or individual learning. These two dimensions, together with how they change over time, are the essence of discrete-choice theory with social influence (Brock and Durlauf, 2001; Brock et al., 2014). We chose discrete-choice theory as an exploratory vehicle because of its relationship to other theories of decision making, both individually and in groups, such as replicator dynamics and Bayesian updating and information theory (Krakauer, 2011) and statistical mechanics (Durlauf, 1999).

This led us to propose a theoretical framework grounded in a bi-axial map that extracts, from observational data, the transparency of decisions and the extent to which a behavior is acquired socially versus individually (Fig. 1) (Bentley et al., 2011a, 2011b,




Fig. 1. A four-quadrant map for understanding different domains of human decision making, based on whether a decision is made individually or socially (horizontal axis) and the transparency of options and payoffs that inform a decision (vertical axis) (after Bentley et al., 2014).

2014; Bentley and O'Brien, 2011, 2015; Brock et al., 2014). The horizontal axis represents the learning continuum, as we aim to identify popularity-data signatures that distinguish individually motivated actions from those driven by social influence. The vertical axis captures how transparent the payoffs of actions and/or their role models are.

The map has characteristic “quadrants.” Rational choice, for example, maps in the northwest, and the popularity of variables tends to be normally (Gaussian) distributed as a result of cost/benefit constraints underlying them. Well-informed social learning—for example, using high-ranking individuals or renowned experts as models (Henrich and Gil-White, 2001; Atkisson et al., 2012)—lies in the northeast. There, the popularity of variables tends to follow a log-normal distribution of proportional advantage, as agents judge a behavior by its popularity rather than by its intrinsic utility.

In the southeast, socially based decisions are made, but payoffs among different options begin to lack transparency. In the extreme, agents copy each other in an unbiased manner, not intentionally copying skill or even popularity. It is as if each person points to someone else and says, “I’ll have what she’s having” (Bentley et al., 2011a). This is not to say that agents behave randomly in the southeast; rather, it says that in the pattern at the population scale, their biases and individual rationales balance out.

Random choice lies in the southwest quadrant, where agents interact minimally and select from among many similar options that lack transparent payoffs. Referring to our examples above, the southwest might characterize a model in which specific nuts in a forager’s basket are shared with each group member. Although behavioral ecologists probably are not interested in which nuts are shared, modern market researchers certainly are, in terms of predicting consumer choices among literally thousands of extremely similar products and information sources (Ehrenberg, 1959; Evans and Foster, 2011; Sela and Berger, 2012).

To resolve both axes of the map simultaneously requires traditional and novel time-series analysis of the form and dynamics of popularity distributions cross-referenced with studies from individuals and populations (Brock et al., 2014). We have

parameterized the functions that underlie the map so that we can estimate paths through it. The vertical axis, which we parameterize as b_t , represents the transparency of an individual’s decision and its consequences—costs and payoffs—from absolute transparency along the northern edge ($b_t = \infty$) to complete opaqueness along the southern edge ($b_t = 0$). The horizontal axis, measured by parameter J_t , represents the extent to which a decision is made, from purely individually at the western edge ($J_t = 0$) to purely social decision making, or copying, at the eastern edge ($J_t = \infty$). This allows a parameterization in terms of how probability of choice k (P_k , as opposed to the null-choice probability, P_0) depends on transparency of choice and social influence. Here the null choice, zero, is introduced to allow a choice made outside the set of options $1, 2, \dots, N$ to serve as a useful baseline option. It is also useful in tidying up notation when writing log-odds regression equations in discrete-choice estimation theory.

We follow Anderson et al. (1992, chap. 2, especially app. 2.10.4) in developing the relatively standard background from discrete-choice theory that eventually leads to the estimation equation discussed later (see also Greene (2003) and especially Amemiya (1985, chap. 9)). At each date t assume there are $i = 1, 2, \dots, I$ persons selecting from a set $\{0, 1, 2, \dots, N\}$ of choices. An individual, i , is assumed to face N primary choices of interest plus another choice, denoted by zero. We assume the payoff of any given choice k to person i at date t consists of a deterministic term, $U_t(i, k)$, and a random term, $\tilde{\varepsilon}_t(i, k)$. We assume the latter is distributed identically and independently across people, choices, and dates. We also add the restriction that the distribution of the random term is a double-exponential distribution with shape parameter μ_t (Anderson et al., 1992, Eqs. 2.30 and 2A.6). Note that μ_t is allowed to change over time. Anderson et al. (1992, eq. 2.30) show that if there are no social effects, i.e., $J(\varphi_2, y_{it}) = 0$, the probability that person i chooses choice k at date t is given by

$$P_t(i, k) = \frac{1}{Z_t} e^{(1/\mu_t)U_t(i, k)}, k = 0, 1, 2, \dots, N$$

$$Z_t \equiv \sum_{j=0}^N e^{(1/\mu_t)U_t(i, j)}, t = 0, 1, 2, \dots, T. \quad (3.1)$$

We define transparency as $b_t \equiv 1/\mu_t$ (this is commonly denoted *intensity of choice* in other literature). The variance in the decision making of person i is proportional to μ_t^2 (Anderson et al., 1992, Eq. 2A.13). When we use the word “transparency,” we are referring to b_t and will specify a functional form for b_t as a function of a parameter vector and observable characteristics (we do not claim that this is the most useful definition of “transparency” in all contexts, but it enables us to make quantitative progress in some contexts).

The larger the transparency, b_t , is, the smaller the variance in decision making across the alternatives is. In other words, $b_t = 0$ corresponds to the lowest level of transparency—noise in choice is so large that choice is completely random over the choice set, and each option is chosen with the same probability. The opposite polar case is $b_t = \infty$, which corresponds to the highest level of transparency—the relative values of payoffs of each choice are so high that there is no doubt as to which choice yields the highest payoff.

One might ask whether the intensity-of-choice measure, b_t , from discrete-choice theory is a useful measure to an econometrician or statistician interested in using observational data on a set of decision makers to estimate and measure “transparency.” Discrete-choice theory, or the *random utility model*, originated from attempts to model a single decision maker whose state of mind is randomly changing (Anderson et al., 1992, chap. 2). Randomness can arise from many causes, ranging from incomplete understanding of the values of the various options, to learning

about the values of the various options, to inherent changes in the values of the options to the decision maker. If none of this randomness were present in the mind of the decision maker, he or she would simply rank the options and choose the best one. This is the polar case of infinite intensity of choice, $b_t = \infty$.

Another interpretation of the discrete-choice model holds that uncertainty is a result of the lack of information available to the modeler (Anderson et al., 1992; Manski, 1977). This could be the result of nonobservable characteristics, nonobservable variations in individual utilities, measurement errors, and functional misspecifications. Fortunately, “the two approaches lead to the same choice probabilities” (Anderson et al., 1992, p. 33), which is why we decided to use the intensity-of-choice measure, b_t , as our measure of “transparency.”

Returning to Eq. (3.1), note that it can be written in the equivalent log-odds form as

$$\ln \left(\frac{P_t(i, k)}{P_t(i, 0)} \right) = \left(\frac{1}{\mu_t} \right) [U_t(i, k) - U_t(i, 0)] = b_t [U_t(i, k) - U_t(i, 0)], \quad (3.2)$$

which is handier for estimation purposes. We now extend this development to include social effects, which we can include by replacing $U_t(i, k)$ with

$$V_t(i, k) = U_t(i, k) + J_t P_t(k), \quad (3.3)$$

where $P_t(k)$ denotes the fraction of the community that chooses choice k at date t .

Suppose one has a set of characteristics of the $N+1$ different available choices at dates $t=1, 2, \dots, T$, denoted by $\{x_{ijt}, i=1, 2, \dots, I, j=0, 1, 2, \dots, N, t=1, 2, \dots, T\}$. Suppose also that one has a data set on individuals, $i=1, 2, \dots, I$, that consists of variables that may impact the variance of decision making at each date, denoted by $\{z_{it}, i=1, 2, \dots, I, t=1, 2, \dots, T\}$. Finally, suppose that one has a set of variables, $\{y_{it}, i=1, 2, \dots, I, t=1, 2, \dots, T\}$, that should enter the function $J(\varphi_2, y_{it})$, which parameterizes the social influence on choices made by individual i . We also assume for estimation purposes that we have data on the average choice fractions, $P_t(j), j=0, 1, 2, \dots, N$, made by the community. Given these data, we write the estimating equation from Brock et al. (2014, Eq. 5) that is appropriate for implementation by non-linear least squares (NLLS) as

$$\ln \left(\frac{P_t(i, k)}{P_t(i, 0)} \right) = b(\theta, z_{it}) \{ \varphi_1(x_{ikt} - x_{i0t}) + J(\varphi_2, y_{it})(P_t(k) - P_t(0)) \}, \quad (3.4)$$

where $(\theta, \varphi_1, \varphi_2)$ is a vector of parameters to be estimated and $U_t(i, k) - U_t(i, 0) = \varphi_1(x_{ikt} - x_{i0t})$.

Recall that the characteristic x_{ikt} is a scalar here. It is easy to generalize the treatment of characteristics to the case where they are vectors.

Here we repeat that our estimating equation is motivated by the underlying theory from, for example, Eqs. (3.2) and (3.3). Brock et al. (2014) developed this framework and showed how estimation of the parameter vector $(\theta, \varphi_1, \varphi_2)$ can be done in practice by trying it out on simulated data, but they did not address conceptualizing and estimating a useful specification of fitness function or how to deal with multiple equilibria. We take up these issues below, after briefly discussing potential biases we face in conducting estimation procedures using field data.

Once we have estimates of the parameter vector, $(\theta, \varphi_1, \varphi_2)$ —denote these estimates as $(\hat{\theta}, \hat{\varphi}_1, \hat{\varphi}_2)$ —we may insert them into the social-surplus function, $S(V)$, defined by Anderson et al. (1992, Eq. 2.37), which we slightly modify in Eq. (3.5)

$$F(p; b, J) = \mu \ln \left[\left(\frac{1}{N+1} \right) \sum_{k=0}^N e^{V_k/\mu} \right] = \left(\frac{1}{b} \right) \ln \left[\left(\frac{1}{N+1} \right) \sum_{k=0}^N e^{bV_k} \right] \quad (3.5)$$

to obtain an estimated fitness function, $\hat{F}_t$. The function can change over time because although the estimate of the parameter vector $(\hat{\theta}, \hat{\varphi}_1, \hat{\varphi}_2)$ does not change, the vector of covariates can. In this way, we obtain not only an estimate of a path $\left\{ b(\hat{\theta}, z_{it}), J(\hat{\varphi}_1, y_{it}) \right\}_{t=0}^T$ but also an estimate, $\hat{F}_t$, over time for fitness along the estimated path $\left\{ b(\hat{\theta}, z_{it}), J(\hat{\varphi}_1, y_{it}) \right\}_{t=0}^T$ on the map. We discuss this in more detail below.

Parameterization of J_t by the function $J_t = J(\varphi_2, y_{it})$ is for estimation purposes. Regarding the social-learning dimension, J , we assume that social utility is positively influenced by relative popularity by letting $J \geq 0$ denote the strength of social influence on decision making. Although for specificity we will emphasize the case of positive association with relative popularity, note that in Eq. (3.4) $J(\varphi_2, y_{it})$ could be negative, which might, for example, represent anticonformity.

Our goal has been to extract these dimensions from observational data. Using simulated data, we showed in an earlier paper (Brock et al., 2014) how functions representing b_t and J_t could be specified in terms of data and a vector of parameters and how these parameters could be estimated by NLLS. We generated noisy random values of “real” variables x , y , and z and used NLLS and our parameterization above to see how well we could recover the dimensions b (north–south) and J (east–west). In Fig. 2, the red dots are simulated data, and the blue dots are our estimates. One can see that the largest discrepancies between model and data lie in the north–south dimension. This is because the expected value of the optimal decision, made over N possible choices, needs to be computed for a given set of parameters of the payoffs to the decision makers. The problem lies in the multiple equilibria that appear as social influence becomes strong in an environment where there are many possible choices.

Here we show how to compute the expected value of the optimal decision as a function of the basic parameters of the environment for deciding among N possible choices. Our objective is to estimate the functions $b(\cdot)$ and $J(\cdot)$ simultaneously (Brock et al., 2014), but we need to be aware of the difficulties in distinguishing between homophily and social causation. As far as we know, Goldbaum and Mizrach (2008) are the only ones to estimate the function $b(\cdot)$ as a function of observable covariates. However, they did not discuss estimation of social influence and potential biases caused by selection bias and correlated unobservables or estimation of a fitness function. In fact, estimates of the function $J(\cdot)$ may be picking up any kind of “social” correlation. More research is needed on methods to remove such biases in estimations of $b(\cdot), J(\cdot)$, as they can affect estimation of the corresponding fitness function, $\hat{F}_t$.

4. Optimal decisions among many options under social influence

Samuelson (1941, 1947) argued in his correspondence principle that an equilibrium would not be observed in a field setting unless it were stable with respect to a plausible out-of-equilibrium adjustment process. This notion, however, has been debated ever since (e.g., Echenique, 2008), the main point of contention being what should serve as a plausible out-of-equilibrium adjustment process. We formulate a process below that seems to make sense. For ease of presentation, we change the notation slightly by deleting the person index, i , and the date index, t . We will also think of the out-of-equilibrium adjustment dynamics as occurring on a faster time scale than the time scale of change in the observable variables that drive changes in the intensity of choice and social influence. Looking ahead, we argue that unstable equilibria




Fig. 2. Map comparing nonlinear least-squares estimates (blue dots) against simulations (red dots) of the model defined by Eq. (3.4) (after Brock et al., 2014, Fig. 4). The simulated data for variables x_{i1gt} , x_{i0gt} , y_{igt} , and z_{igt} were each independently chosen from values normally distributed through time t with mean 10, with the variance for the payoffs x_{i1gt} and x_{i0gt} set at 0.01 and the variance for y_{igt} and for z_{igt} set at 1. For individual i in group g at time t , the payoff difference between options 1 and 0 is represented by $x_{i1gt} - x_{i0gt}$, the presence of social influence is measured by y_{igt} , and z_{igt} represents how variable the choices were through time. As described by Brock et al. (2014), estimating the parameter vector θ , along with the scalar observable z , determines the transparency of choice, $b(\cdot)$, and estimating the parameter vector, ψ_2 , specifies the social-influence function, $J(\cdot)$. Each simulation used 30 time steps, 100 groups, and 200 agents per group and noise component $\varepsilon_{i1gt} - \varepsilon_{i0gt}$, with mean 0 and $\sigma = 0.1$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

with respect to our out-of-equilibrium adjustment process are unlikely to appear in any observational data set. This is why we spend time below exploring properties of the adjustment process. Finally, we note that it is notationally more convenient to drop the “baseline” choice labeled zero and work with choices labeled $k = 1, 2, \dots, N$.

Consider an environment consisting of N possible decisions, each with payoff $U_k, k = 1, 2, \dots, N$. As previously, we let $b \geq 0$ denote the transparency of the payoff values of the decisions. With zero transparency ($b=0$), decision makers have zero confidence in their evaluation of payoffs, so they select choice $k = 1, 2, \dots, N$ with probability $1/N$, i.e., completely randomly. With complete transparency, $b \rightarrow \infty$, agents can choose the highest payoff option with probability one.

In the discrete-choice literature, the utility of different choices consists of a deterministic utility, U_{kt} , plus an element of randomness, $\mu_t \tilde{\varepsilon}_{kt}$, as

$$\tilde{U}_{kt} = U_{kt} + \mu_t \tilde{\varepsilon}_{kt}, k = 1, 2, \dots, N, t = 1, 2, \dots, T. \quad (4.1)$$

Here we have placed subscripts to indicate that when discussing estimation we allow the utilities as well as the parameter μ_t to change over time. In instances where we are simply explaining the basics of discrete-choice theory, we drop the date subscript to reduce notation.

The deterministic part of each payoff, U_k , can be ranked as $U_1 > U_2 > \dots > U_N$, such that, without loss of generality, when μ is small, most choices will be clustered around choice number 1. As μ

becomes large, the choice probabilities spread out across the N choices, approaching a uniform distribution as μ becomes very large. Because our intensity-of-choice function, $b(\cdot)$, is inversely related to μ , the farther north we go on the map, the more tightly the distribution of choices clusters around the maximum deterministic payoff. This is what motivates our estimation of $b(\cdot)$ as a function of observable covariates (Brock et al., 2014).

As previously, the deterministic component of utility when social effects are present is given by

$$V_k = U_k + Jp_k, k = 1, 2, \dots, N. \quad (4.2)$$

Actual utility of choice, k , is random and consists of the deterministic component, V_k , and a random component, $\tilde{\varepsilon}_k$, that is distributed double exponentially with scale parameter μ . The random component is distributed independently and identically across choices. This implies that the pairwise difference of the random components is distributed logistically with zero mean and variance, $\sigma^2 = \pi^2 \mu^2 / 6$ (Anderson et al., 1992, p. 60). Hence, we write the random utility as

$$\tilde{V}_k = V_k + \tilde{\varepsilon}_k, k = 1, 2, \dots, N. \quad (4.3)$$

An important and standard quantity from discrete choice theory (Anderson et al., 1992, pp. 60–61), which we shall use later, is the expectation of the maximum of the random utilities across choices,

$$E\{\max_{k \in \{1, 2, \dots, N\}} \tilde{V}_k\} = S(V) = \mu \ln \left[\sum_{k=1}^N e^{(V_k/\mu)} \right] = \left(\frac{1}{b} \right) \ln \left[\sum_{k=1}^N e^{bV_k} \right], \quad (4.4)$$

where $b \equiv 1/\mu$. The quantity $S(V)$ can be thought of as a welfare measure for an individual facing this particular choice environment.

Our preliminary goal is to show how $(b(\cdot), J(\cdot))$ in a population of agents facing N possible choices determines the fitness of the optimal decision in the population. For such a population, the maximum payoff for each fixed value of $(b(\cdot), J(\cdot))$ is the expected value of the optimal decision, i.e., the “fitness function.” This is a novel problem because the optimal decision depends not only on intrinsic utility, U_k , but also on transparency, b , and social learning, J , as well as on the relative popularity, p_k , of each of the N possible choices. This recursive relationship means that multiple equilibria can exist and that the fitness function may also change over time as $(b(\cdot), J(\cdot))$ changes.

To keep things simple to start, we fix the values of b and J . We can think of (b, J) as changing on a slower time scale than the faster dynamics that give rise to the optimal decision. By assuming that (b, J) are constant over this faster time scale, we will introduce an algorithm that computes equilibria as well as points on our fitness landscape. Even with fixed b and J , we still face the problem of multiple equilibria, so we start by computing equilibrium solutions for the extreme values (zero and infinity) of b and J . Then, when there are multiple equilibria, we argue that certain “natural” equilibria will be unique for each pair of b and J .

We start with the multinomial logit framework for the general case of N choices (Brock and Durlauf, 2001), in which the probability of choice k (from 1 to N) is given by

$$p_k = \frac{1}{\sum_{i=1}^N e^{bU_i + bJp_i}} e^{bU_k + bJp_k}, \quad (4.5)$$

where $k = 1, 2, \dots, N$ indexes the N different choices. It is a fixed-point equation because we must find a vector $(p_1, p_2, \dots, p_N)$ of probabilities that simultaneously satisfies both sides of Eq. (4.5) for $k = 1, 2, \dots, N$. We have explained how the choice probabilities given in Eq. (4.5) are derived from the random utilities (Anderson et al., 1992, Eq. 2.30). Recall that μ is the shape parameter of the double-exponential distribution function (Anderson et al., 1992,

Eq. 2A.6) and that $b \equiv 1/\mu$ by definition. We add an extra term to the usual expected value of the optimal decision, $S(V)$ (Anderson et al., 1992, Eq. 2.37), in order to build the fitness function that we want to compute. Our fitness function is given by

$$F(p; b, J) \equiv E\left\{\max_{i \in \{1, 2, \dots, N\}} \tilde{V}_k\right\} + \frac{1}{b} \ln\left(\frac{1}{N}\right) = S(V) + \frac{1}{b} \ln\left(\frac{1}{N}\right) \\ = \frac{1}{b} \ln\left[\left(\frac{1}{N}\right) \sum_{k=1}^N e^{bU_k + bJp_k}\right]. \quad (4.6)$$

As an interim measure of “fitness” in Eq. (4.6), we have taken the standard discrete-choice measure of welfare, $S(V)$, and added the term $\frac{1}{b} \ln(\frac{1}{N})$. Note that this does not change the partial derivatives with respect to U_k or p_k for $k = 1, 2, \dots, N$. Here, $p \equiv (p_1, \dots, p_N)$. If we knew how to compute a “natural” choice of equilibrium vector, call it $p^*(b, J)$, to solve the N Eq. (4.5), we could define our fitness-landscape function as $W(b, J) \equiv F(p^*(b, J); b, J)$. We will discuss what we mean by a “natural” equilibrium in Section 5.

This modification of adding the constant term $\frac{1}{b} \ln(\frac{1}{N})$ to the formula for the term $E\left\{\max_{i \in \{1, 2, \dots, N\}} \tilde{V}_k\right\}$ in Anderson et al. (1992, eq. 2.37) turns out to have some useful properties. One is that we can address the problem of multiple equilibria first by computing equilibrium solutions for the extreme values $b=0, \infty$ and $J=0, \infty$. For $J=0$, the solution of (4.5) is unique. Further, for $J=0$,

$$b \rightarrow 0, \text{ implies, } F(p; b, 0) \rightarrow \left(\frac{1}{N}\right) \sum_{k=1}^N U_k \\ b \rightarrow \infty, \text{ implies, } F(p; b, 0) \rightarrow \max\{U_k\}. \quad (4.7)$$

(See Appendix 1 for a short proof of (4.7).) Suppose one plots the “fitness function” (4.6) on the vertical axis and plots (b, J) on the horizontal axis of a graphical display of a fitness landscape. We can see right away from our results (4.7) that if one draws the peaks for the far west side of the map, where $J=0$, one will see the height of the peak rising as b increases. For $J=0$, the height will rise from the average payoff at the far south, i.e., where $b=0$, to the maximum of the payoffs as one heads north.

For finite $J > 0$, writing $p_k(b)$ to emphasize the dependence of p_k on the value of b , we have

$$b \rightarrow 0, \text{ implies, } F(p; b, J) \rightarrow \left(\frac{1}{N}\right) \sum_{k=1}^N (U_k + Jp_k(0)) \\ p_k(0) = \frac{1}{N}, k = 1, 2, \dots, N \quad (4.8)$$

$$b \rightarrow \infty, \text{ implies, } F(p; b, 0) \rightarrow \max\{U_k + Jp_k(\infty)\}$$

$$p_k(\infty) = 1.$$

Suppose we rank the deterministic payoffs as mentioned above, with $U_1 > U_2 > \dots > U_N$, and U_1 as the maximum payoff. We have already seen for the case $J=0$ that one can prove that $p_1 \rightarrow 1$ as $b \rightarrow \infty$. However, for $0 < b < \infty$, the next result shows that multiple equilibria will appear when J becomes large enough. Somewhat surprisingly, one can show that when $J \rightarrow \infty$, for any choice k there is an equilibrium where the probability that k is chosen is one. This is because we are restricting ourselves to settings where the social utility of a choice correlates positively with its popularity (we recognize there will be many settings where this is not the case). As choice k becomes more popular, the social payoff (4.2) rises to the point where it surpasses the intrinsic part of the payoff and an agent can simply conform to the majority choice of the community (McElreath et al., 2008).

To summarize, we see that for $J \rightarrow \infty$, the equilibria, i.e., limiting fixed points, in (4.5) occur at $p_k=1$, and all other probabilities zero

for $k = 1, 2, \dots, N$. If $U_k=U_j=0$, then for $J \rightarrow \infty$, the equilibria of (4.5) occur with $p_k=1/2, p_j=1/2$ for every pair $k \neq j$. If $U_k=U_j=U_l=0$, then for $J \rightarrow \infty, p_k=1/3, p_j=1/3, p_l=1/3$ for every triplet $k \neq j \neq l$, and so on for all possible mixed-strategy equilibria. Parenthetically, it might be argued that a discussion of limiting behavior and multiple equilibria, including mixed-strategy equilibria, is extraneous to the main goal of estimation of intensity of choice, social influence, and fitness function, but we believe that it is important to have a solid understanding of the theory that lies behind the objects one is trying to estimate.

5. The fitness landscape over multiple equilibria

For large values of the social-learning parameter, J , there can be a plethora of equilibria. We need to evaluate the fitness function $F(p; b, J)$ at each of these equilibria in order to complete the fitness landscape discussed in Section 4. We can start with small values of $N=2, 3, 4, \dots$ and then evaluate $F(p; b, J)$ at each one of these equilibria in order to show the height of the fitness landscape at each value of (b, J) . We start with the case of $N=2$, where we can compute the equilibria as $b > 0$ and $J \geq 0$ vary by plotting a graph and using the constraint $p_1 + p_2 = 1$. For $N=3$ we can use $p_1 + p_2 + p_3 = 1$ to reduce the problem to solving two equations in two unknowns, (p_1, p_2) , for each value of the vector (b, J) , and in the general case we solve the equations below, where we have added a set of differential equations whose steady states are equilibria:

$$p_k = \frac{e^{bU_k + bJp_k}}{\sum_{j=1}^N e^{bU_j + bJp_j}} \equiv f_k(p; b, J), p_k \geq 0, k = 1, 2, \dots, N \\ \sum_{k=1}^N p_k = 1 \\ \frac{dp_k}{dt} = f_k(p; b, J) - p_k, k = 1, 2, \dots, N. \quad (5.1)$$

Given that multiple equilibria may appear for $b > 0$ as $J \rightarrow \infty$, where the fitness function itself may not be well defined or have infinite limit, which equilibrium should we use for the value of $F(p; b, J)$? We face the task of computing all the equilibria for each $b > 0, J \geq 0$ and possibly picking the one that gives the largest value of $F(p^*(b, J); b, J)$. However, we argue below that a better option is to specify a plausible out-of-equilibrium dynamic process in the form of a set of differential equations, i.e., a dynamical system whose steady states are equilibria (solutions of equations such as (5.1) for the general case of N choices), and to simulate the dynamical system to identify equilibria that are local attractors for the dynamical system (5.1). We labeled steady states that satisfy this local stability property as “natural” equilibrium points. The motivation for this choice is that an “unnatural” equilibrium is unlikely to be observed, as was argued by Samuelson (1941, 1947).

We use a Lyapunov function for this exploration. The specific choice of Lyapunov function is not so important, as long as its gradient, where zero, yields a solution to Eqs. (4.5) and (5.1). We propose the following as our Lyapunov function:

$$H(p; b, J) \equiv F(p; b, J) - \frac{J}{2} \sum_{k=1}^N p_k^2. \quad (5.2)$$

We claim that the partial derivatives of H with respect to $p_k, k = 1, 2, \dots, N$ are all zero on equilibria and that $L(p; b, J) \equiv -H(p; b, J)$ acts like a Lyapunov function for the dynamical system,

$$\frac{dp_k}{dt} = \frac{\partial H}{\partial p_k} = J \left(\frac{e^{bU_k + bJp_k}}{\sum_{j=1}^N e^{bU_j + bJp_j}} - p_k \right). \quad (5.3)$$

Note that this system is the same as the one in (5.1) except for the constant J , which does nothing but change the “speed” of the

dynamics, provided that $J > 0$, which we assume to be the case from this point on. This system is essentially the same as the one in (5.1), except for the units of “time,” and is what we will use as our out-of-equilibrium adjustment process in order to implement Samuelson’s (1941, 1947) correspondence principle. We chose (5.3) because, as we will see, it has very nice mathematical properties and is, hence, quite easy to analyze.

The dynamical system is a “gradient dynamical system,” that is, the right side of (5.3) is the gradient of a scalar function $H(p; b, J)$. It yields a handy “pseudo-Lyapunov” function, $L(p; b, J) = -H(p; b, J)$. We call L a “pseudo-Lyapunov” function because it does not satisfy the usual definiteness conditions for a Lyapunov function; however, it is still useful for computing equilibria.

Intuitively, the dynamical system (5.3) “climbs” the function $H(p; b, J)$ because it moves in the direction of the gradient vector $\partial H(p; b, J)/\partial p$. Hence, we expect the system to stop climbing and come to rest on local maxima of the function, $H(p; b, J)$. We call these equilibria of (5.1), which are local maxima of $H(p; b, J)$, “natural equilibria.” Again, Samuelson’s correspondence principle demands that an equilibrium be “natural,” i.e., that it have a nonzero basin of attraction with respect to an out-of-equilibrium dynamical adjustment process, which we take to be the system (5.3). The system will converge to that equilibrium if it starts near enough (in terms of that equilibrium). The proof is simple, so we give it here. We must show for $L(p; b, J) = -H(p; b, J)$ that $dV/dt < 0$ on solutions of (5.3) and is zero on solutions of (4.5), i.e., is zero on the set of equilibria. Or, put another way, it is zero on steady-state solutions of the differential Eq. (5.1). We have from (5.3) and the definition of L ,

$$\frac{dL}{dt} = -\frac{dH}{dt} = -\sum_{k=1}^N \left(\frac{\partial H}{\partial p_k} \right) \left(\frac{dp_k}{dt} \right) = -\sum_{k=1}^N \left(\frac{\partial H}{\partial p_k} \right)^2 < 0 \quad (5.4)$$

Hence,

$$\frac{dL}{dt} = 0 \text{ iff } \frac{\partial H}{\partial p_k} = 0, k = 1, 2, \dots, N. \quad (5.5)$$

Note that the equilibrium that makes the function $F(p; b, J)$ largest on the set of equilibria, call it $p^*(b, J)$, may not be found by the process (5.3) because $p^*(b, J)$ may not be a steady state of local attraction for the process (5.3), i.e., it is unstable with respect to the process (5.3). We exclude such equilibria for the “Samuelsonian” reasons given above. Recall Eq. (4.4), where we let $p^{**}(b, J)$ denote the equilibrium that gives the largest value of $F(p; b, J)$ over the set of natural equilibria.

The function $F(p^{**}(b, J); b, J)$ appears to be the most natural candidate for a fitness-landscape function. If so, this raises the question of whether the set of equilibria found by the hill-climbing algorithm (5.3) and the bottom-seeking version of (5.3), where we multiply the right side of (5.3) by -1 , will contain the equilibrium that corresponds to $p^{**}(b, J)$. Provided that we grid the simplex $\{(p_1, \dots, p_N) | \sum_{k=1}^N p_k = 1\}$ of non-negative probability vectors finely enough and solve the dynamical system (5.3) from each of these gridded initial conditions, we can be fairly certain that we will find all the “natural” equilibria and hence find the equilibrium we seek, $p^{**}(b, J)$.

Some may argue for the use of a process other than (5.3), and we reiterate that it is only one plausible candidate. Whatever one’s view is on the usefulness of this particular process for implementing Samuelson’s correspondence principle, any estimation method of the fitness function,

$$F(p; b, J) \equiv \mu \ln \left[\left(\frac{1}{N} \right) \sum_{k=1}^N e^{V_k/\mu} \right] = \left(\frac{1}{b} \right) \ln \left[\left(\frac{1}{N} \right) \sum_{k=1}^N e^{bV_k} \right],$$

which we implemented for estimation purposes above as the quantity $\hat{F}_t$ using observational data, is likely to be estimated only

on equilibria that are “stable” with respect to some out-of-equilibrium adjustment process, assuming the system is in equilibrium in the first place. The vexing issue of usefully modeling how particular equilibria are actually approached by a social system remains an open problem.

6. Computing multiple equilibria

In summary, we have defined a fitness function $F(p; b, J)$, which is a modification of the “social surplus” function (Anderson et al., 1992, Eq. 2.37), which itself is the expected value of the optimal decision, at a particular b, J coordinate with an optimal division of popularities, given by the vector of probabilities, p_k , for each choice $k = 1, 2, \dots, N$ among N options. We use a Lyapunov function, $H(p; b, J)$, to find each natural equilibrium with respect to (5.1), or, equivalently, (5.3), which is the same system up to speed of convergence. Because there might be multiple peaks that are locally stable, we initialize our dynamical system over a range of $p_k(0)$. In this way, we are able to find all of the natural equilibria and evaluate $F(p; b, J)$ at each of these. If $p^{**}(b, J)$ is the choice probability vector that gives the largest value of $F(p; b, J)$ over the set of natural equilibria, the fitness landscape will plot as $F(p^{**}(b, J); b, J)$. This entire landscape depends on U_k , the vector of utilities for all the different choices $k = 1, 2, \dots, N$, and could potentially be quite different for a different vector of intrinsic utilities of the N different choices.

To sum up the relations among these variables, we have

$$\begin{aligned} p_k &= \frac{e^{bU_k + bJp_k}}{\sum_{i=1}^N e^{bU_i + bJp_i}} \\ &= \frac{e^{bV_k}}{\sum_{i=1}^N e^{bV_i}} \\ &= f_k(p; b, J), k = 1, 2, \dots, N, \end{aligned} \quad (6.1)$$

$$F(p; b, J) = \frac{1}{b} \ln \left[\frac{1}{N} \sum_{k=1}^N e^{bU_k + bJp_k} \right], \text{ and} \quad (6.2)$$

$$H(p; b, J) = F(p; b, J) - \frac{J}{2} \left(\sum_{k=1}^N p_k^2 \right). \quad (6.3)$$

Note that

$$\sum_{k=1}^N f_k(p; b, J) = 1 \quad (6.4)$$

for all $(p; b, J)$, such that $\sum_{k=1}^N p_k(t) = 1$ for all dates t . For given values of b and J , we use computer simulation to follow the gradient given by the Lyapunov function. In this simulation, we must first specify the intrinsic utilities of all N choices by the vector U . These intrinsic utility values do not change through the simulation. We then choose an initial set of p_k values from $p_1, p_2, p_3, \dots, p_N$, the sum of which is one. In the case of our diagrams for $N=2$ and $N=3$, respectively, this means

$$\begin{aligned} p_1(t) + p_2(t) &= 1 \\ p_1(t) + p_2(t) + p_3(t) &= 1 \end{aligned} \quad (6.5)$$

for all dates t for the dynamics (we deleted t in the notation to avoid confusion).

Each value of N corresponds to a different fitness-landscape function of (b, J) for each N . We can plug into Eq. (5.3) each pair of b and J values to obtain a partial derivative for each p_k value. The resulting vector of partial derivatives, $(\partial H/\partial p_1, \partial H/\partial p_2, \dots, \partial H/\partial p_N)$, multiplied by our chosen scalar step size h (we used $h=0.02$), points us one step up the hill of $H(p; b, J)$. We follow this iteratively until the gradient vector approaches zero, at which point we have




Fig. 3. Illustration of hill-climbing process for $N=2$ choices, of $U_1=0.511$ and $U_2=0.489$, respectively, at a specific map coordinate, $b=5$ and $J=1$. Among the four plots, the top row shows how the Lyapunov algorithm converges on p_1 and on F over the t time steps, and the bottom left shows $\Delta p_1/dt$ against p_1 . The bottom right shows the equilibrium probability at the end of the hill-climbing algorithm. In this example, the fitness was maximized through 100% choice 2.





Fig. 4. A vector figure for $N=2$, $b=5$, and $J=1$, with utility values $U_1=0.5$ and $U_2=0.4$. The line represents $p_1 + p_2 = 1$, along which each initial $p_1(0)$ shows arrows that point at that line in the direction of dp_1/dt starting at $p_1(0)$. Each arrow has a length proportional to the size of dp_1/dt . Note that $dp_2/dt = -dp_1/dt$.

a fairly sharp transition from positive to negative hill-climbing gradient $\partial H/\partial p_1$ (Fig. 6, lower right). When J is larger, $1 \leq J \leq 2$, however, we find a step transition in final fitness at lower p_1 , such that a small increase in initial p_1 yields a substantial increase in final fitness (Fig. 6, upper left). Along this discontinuity, there is also a slight slope such that a small change in social transparency, J , would also abruptly change the fitness, but only very close to the discontinuity (Fig. 6, upper left).

The sensitivity to initial p vectors shows that in order to account for potential multiple fitness peaks at each b, J coordinate, we need to find *all* equilibria, which could number N or fewer at each coordinate. Hence, in order to generate a final fitness map in b, J space for $N=3$ (Fig. 7), we need to choose a sufficient range of different starting vectors, $p(0)$, to cover the space of possible p vectors and then, for each different $p(0)$, follow the gradient of $\partial H/\partial p_1$, $\partial H/\partial p_2 \dots \partial H/\partial p_N$ to its corresponding fitness peak. After checking for multiple equilibria and retaining the maximum fitness value for each, this algorithm is repeated for discrete choices of b and J to fill in the map in pixelated fashion (Fig. 7). The computation time increases exponentially with the number of choices, N , which is why Fig. 7 shows the b, J map for a relatively modest value of $N=3$. Keep in mind this map corresponds to a

found an equilibrium fitness peak at that map coordinate of (b, J) . Fig. 3 shows features of the algorithm for $N=2$, $b=5$, and $J=1$.

The choice of $p(0)$ can matter even for a binary choice, $N=2$. Fig. 4 (after Lahkar and Sandholm, 2008) shows, for $N=2$, how the equilibrium point changes as a function of the initial starting value, p_0 , for different pairs of utility values at the same coordinate of $b=5$ and $J=1$. Arrows show the direction of flow toward attractors and away from repellers. For Fig. 4 the utility values of the two choices were set at 0.5 and 0.4, respectively, and the attractors in this case are 100% for either choice, such that if p_1 is greater than about 40% at time 0, it goes to the peak at 100% p_1 ; otherwise it goes to the other peak at 100% p_2 . The fitness value in this case, describing the highest overall fitness, is at 100% p_1 , which has the higher intrinsic utility.

When we increase to $N=3$, the terrain becomes considerably more rugged, with basins of attraction and areas of repulsion within a space defined by (p_1, p_2, p_3) . Fig. 5 shows these landscapes for the selected combinations of b and J , with each set at 1, 5, or 10. Utilities for the three choices were fixed at $U_1=0.5$, $U_2=0.4$, and $U_3=0.3$. In each plot, the axes show the values of p_1 and p_2 , the combination of which determines also $p_3 = 1 - (p_1 + p_2)$. Within the space, arrows point uphill toward the peak defined by the Lyapunov gradient. Empty spaces, or “corridors,” show areas of repulsion, with all arrows pointing away.

Under these parameters, the space for $b=1$ and $J=1$ shows one attractor at approximately $(p_1, p_2, p_3) = (0.38, 0.32, 0.29)$. From this situation, if we increase b to 5, keeping $J=1$, the one attractor in the middle becomes three attractors at the corners, near 100% p_1 , p_2 or p_3 , respectively, with clear corridors of repulsion dividing their respective basins of attraction (Fig. 5). Interestingly, the effect is similar when we increase J to 5, keeping $b=1$, or increase both b and J to 5 (Fig. 5): attractors at 100% for one of the three choices, with clear boundaries in between. As we increase b and J toward 10 for each, these three areas resolve themselves very clearly (Fig. 5).

To explore how this transition happens, we can hold b constant and look at what happens when we vary the east–west coordinate, J , the strength of social influence. Holding $b=5$ constant, as we increase J , we find some interesting developments in terms of both the final fitness values (Fig. 6, left) and the hill-climbing gradient $\partial H/\partial p_1$ (Fig. 6, right). When social influence is small, $J \leq 0.1$, we see




Fig. 5. Plots for $N=3$, with utility values $U_1=0.5$, $U_2=0.4$ and $U_3=0.3$. The nine plots in the grid show three different values for both b and J . For plots on the left, the vertical axis shows final fitness, F ; plots on the right show hill-climbing gradient $\partial H/\partial p_1$.

specific assignment of utility values for the three different choices; different utility values would yield a different b, J map.

As we can see, the space of potential fitness landscapes is enormous, and further exploration will need to be strategic. In one exploration we might determine the maximum fitness peak at each b, J coordinate for higher numbers of choices ($N > 3$) without retaining the results for “sublandscapes” of hill climbing, like those in Fig. 5, in finding these peaks at every b, J coordinate. Each landscape depends on the vector of utility values as well, so we would then plot a landscape in b, J space for $N=3$, $N=4$ and so on, for each specified vector of intrinsic utility values, $U_1, U_2, \dots, U_N$. Alternatively, we may also want to focus on the sublandscapes that emerge at these higher N values to determine whether fitness peaks appear at intermediate combinations of (p_1, p_2, p_3) rather than at the corners close to 100% for one of the choices (Fig. 5).

7. Conclusion

We applied discrete-choice theory to construct a fitness landscape over a two-dimensional surface representing transparency of choice and social influence, heuristically represented as orthogonal dimensions. In Section 3 we presented an estimating Eq. (3.4) that can be parameterized and showed how the parameters $\theta, \varphi_1, \varphi_2$ can be estimated once one has the appropriate data sets, which we discussed only briefly, given that actual estimation is beyond the scope of this article. We defined a “transparency” function, $b(\theta, z_{it})$, as the inverse of the standard deviation of the random component of utility. By experimenting with different specifications of the function $b(\theta, z_{it})$ and splitting the data set into subperiods of time, one can test whether decision making becomes more or less precise, i.e., $b(\theta, z_{it})$ increases or decreases with time. Also, one can test whether certain observable characteristics of individual decision makers are associated with more or less precision in decision making. Whereas we used linear specifications of the personal component of the utility difference,

other specifications can be explored. Finally, because the function $J(y_{it}, \varphi_2)$ can be specified, the parameter vector, φ_2 , can be estimated and used to test hypotheses involving the potential presence of social-interaction effects. For example, one could test the null hypothesis of zero social interaction effects, $J(y_{it}, \varphi_2) = 0$.

The decision map we created has a close analogy with field studies of social animals, and we expect it to interest ethologists as well. As a case in point, in their study of how wild olive baboons in Kenya decide how and where to move across the landscape, Strandburg-Peshkin et al. (2015) use a two-dimensional map that plots directional agreement on the vertical axis and the number of initiators of movement on the horizontal axis. We view directional agreement as analogous to transparency, (b), and number of initiators as close to degree of social influence, (J). It is a convenient coincidence that Strandburg-Peshkin et al. (2015) oriented their axes in the same way as ours (Fig. 1). In addition, the observed probability function of a baboon following the movements of one subgroup over another follows a logistic dependence on the numerical difference between the two subgroups (Strandburg-Peshkin et al., 2015), which is equivalent to the right hand side of Eq. (3.4). Analogous to the northeast corner of our map, Strandburg-Peshkin et al. (2015) model the highest predictability of baboon-group movement in the space of many initiators and high directional agreement. Interestingly, increasing the number of initiators (social influence) without increasing directional agreement (transparency) may actually decrease the predictability of followers in the baboon study.

Once we identified social influence and transparency of choice as two key factors in how decision frequencies change through time among multiple options, we explored the problem of optimizing decision fitness at the population scale. This is a novel challenge because fitness depends on both the intrinsic utilities of each choice and the transparency of social learning in terms of relative popularity of the choice taken. Using a hill-climbing algorithm, we then explored how this recursive relationship creates multiple equilibria, which we define as peaks on the fitness landscape. Among the more surprising results is just how rugged




Fig. 6. Image plots for $N=3$, with $b=5$, where p_1 is varied but we set $p_3=3p_2$; that is, $p_2=(1-p_1)/4$ and $p_3=3(1-p_1)/4$. In the plots on the left, the colors show the final fitness values, with numbers next to “+” symbols showing the value along selected contours. On the right, colors and contour values indicate the hill-climbing gradient $\partial H/\partial p_1$. In the lower right, where $J \leq 0.1$, note the fairly sharp transition from positive to negative hill-climbing gradient $\partial H/\partial p_1$. In the upper right, where $1 \leq J \leq 2$, note the step transition in final fitness at lower p_1 , such that a small increase in initial p_1 would yield a substantial increase in final fitness. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the landscape becomes as one moves east on the fitness landscape, as sensitivity to social influence increases. The ruggedness is such that we struggle, even through computational methods, to define the fitness landscape for even $N=4$ different options. The complexity of human choice among similar options is therefore formidable, even in experimental contexts. For example, Salganik et al. (2006) had subjects choose among $N=48$ online music tracks, the popularity of which was either visible or invisible.

In future work we could incorporate diminishing returns from growing popularity in the social portion of our fitness function. In other cases, however, the social fitness we used here may be

appropriate, especially where popularity does contribute more or less directly to overall fitness. In the world of humans, obvious candidates include financial investments or a communication or entertainment technology whose utility grows with more users. In animal culture we might see this be the case, as in chimpanzee tool culture, with our model implying that social learning among even a few or several different options could bring about a rugged fitness landscape and hence multiple regions of different cultural traditions (e.g., Whiten et al., 1999). For these reasons, we see considerable value in continuing to explore fitness landscapes of discrete choice with social influence.




Fig. 7. Contour plot of the b, J map showing the maximum fitness for $N=3$. To find these maximum fitness values, the hill-climbing process was followed at b, J coordinate. At each coordinate, multiple hill-climbing searches were made over a range of initial probability vectors for the three choices: $p_1(0)$, $p_2(0)$ and $p_3(0)$. The maximum fitness was then retained at each b, J coordinate.

Acknowledgments

We thank Andrew Morozov for inviting us to submit this paper and several anonymous reviewers for their substantial assistance in helping us clarify and sharpen our discussion. The example that we use to open the paper came from one of the reviewers.

Appendix A

Consider the expression

$$F(p; b, 0) \equiv E\left\{\max_{k \in \{1, 2, \dots, N\}} \tilde{V}_k\right\} + \frac{1}{b} \ln\left(\frac{1}{N}\right) = S(\underline{V}) + \frac{1}{b} \ln\left(\frac{1}{N}\right) \\ = \frac{1}{b} \ln\left[\frac{1}{N} \sum_{k=1}^N e^{bU_k}\right]. \quad (\text{A.1})$$

Note that when b goes to zero in (A.1), one gets the limiting form 0/0. Apply L'Hospital's Rule of basic calculus and differentiate both the numerator and the denominator of (A.1) with respect to b to obtain

$$\left\{\frac{1}{N} \sum_{k=1}^N U_k e^{bU_k}\right\} / \left\{\frac{1}{N} \sum_{k=1}^N e^{bU_k}\right\} \rightarrow \frac{1}{N} \sum_{k=1}^N U_k, \text{ as } b \rightarrow 0. \quad (\text{A.2})$$

This proves the first line of (4.5). For the second line, suppose without loss of generality, that $\max_k \{U_k\} = U_1$. Then,

$$\frac{1}{b} \ln\left(\sum_{k=1}^N e^{bU_k}\right) = \frac{1}{b} \ln\left(e^{bU_1} \sum_{k=1}^N e^{b(U_k - U_1)}\right) = U_1 + \frac{1}{b} \ln\left(\sum_{k=1}^N e^{b(U_k - U_1)}\right) \\ \rightarrow U_1, \text{ as } b \rightarrow \infty \quad (\text{A.3})$$

because $U_k - U_1 < 0$, for all $k \neq 1$. This proves the second line of (4.7).

References

Amemiya, T., 1985. Advanced Econometrics. Harvard University Press, Cambridge, MA.

- Anderson, S., de Palma, A., Thisse, J.F., 1992. Discrete Choice Theory of Product Differentiation. MIT Press, Cambridge, MA.
- Aral, S., Muchnik, L., Sundararajan, A., 2009. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proc. Natl. Acad. Sci. USA* 106, 21544–21549.
- Atkisson, C., O'Brien, M.J., Mesoudi, A., 2012. Adult learners in a novel environment use prestige-biased social learning. *Evol. Psychol.* 10, 519–537.
- Bentley, R.A., Earls, M., O'Brien, M.J., 2011. I'll Have What She's Having: Mapping Social Behavior. MIT Press, Cambridge MA.
- Bentley, R.A., O'Brien, M.J., 2011. The selectivity of social learning and the tempo of cultural evolution. *J. Evol. Psychol.* 9, 125–141.
- Bentley, R.A., O'Brien, M.J., 2015. Collective behavior, uncertainty and environmental change. *Philos. Trans. R. Soc. A* 373, 20140461.
- Bentley, R., O'Brien, M.J., Brock, W.A., 2014. Mapping collective behavior in the big-data era. *Behav. Brain Sci.* 37, 63–119.
- Bentley, R.A., O'Brien, M.J., Ormerod, P., 2011. Quality versus mere popularity: a conceptual map for understanding human behavior. *Mind Soc.* 10, 181–191.
- Bond, R.M., Fariss, C.J., Jones, J.J., Kramer, A.D.I., Marlow, C., Settle, J.E., Fowler, J.H., 2012. A 61-million-person experiment in social influence and political mobilization. *Nature* 489, 295–298.
- Brock, W.A., Bentley, R.A., O'Brien, M.J., Caiado, C.S.S., 2014. Estimating a path through a map of decision making. *PLoS One* 9 (11), e111022.
- Brock, W., Durlauf, S., 2001. Discrete choice with social interactions. *Rev. Econ. Stud.* 68, 235–260.
- Christakis, N.A., Fowler, J.H., 2013. Social contagion theory: examining dynamic social networks and human behavior. *Stat. Med.* 32, 556–577.
- Christakis, N.A., Fowler, J.H., 2014. Friendship and natural selection. *Proc. Nat. Acad. Sci. USA* 111, 10796–10801.
- Clément, R.J.G., Krause, S., von Englehardt, N., Faria, J.J., Krause, J., Kurvers, R.J.H.M., 2013. Collective cognition in humans: groups outperform their best members in a sentence reconstruction task. *PLoS One* 8 (10), e77943.
- Couzin, I.D., Krause, J., Franks, N.R., Levin, S.A., 2005. Effective leadership and decision-making in animal groups on the move. *Nature* 433, 513–516.
- Couzin, I.D., Ioannou, C.C., Demirel, G., Gross, T., Torney, C.J., Hartnett, A., Conradt, L., Levin, S.A., Leonard, N.E., 2011. Uninformed individuals promote democratic consensus in animal groups. *Science* 334, 1578–1580.
- Durlauf, S., 1999. How can statistical mechanics contribute to social science? *Proc. Natl. Acad. Sci. USA* 96, 10582–10584.
- Dyer, J.R.G., Johansson, A., Helbing, D., Couzin, I.D., Krause, J., 2009. Leadership, consensus decision making and collective behaviour in human crowds. *Philos. Trans. R. Soc. B* 364, 781–789.
- Echenique, F., 2008. The correspondence principle. In: Durlauf, S., Blume, L. (Eds.), *The New Palgrave Dictionary of Economics*, second ed. Palgrave Macmillan, New York.
- Ehrenberg, A.S.C., 1959. The pattern of consumer purchases. *J. R. Stat. Soc. C* 8, 26–41.
- Enquist, M., Eriksson, K., Ghirlanda, S., 2007. Critical social learning: a solution to Rogers's paradox of nonadaptive culture. *Am. Anthropol.* 109, 727–734.
- Evans, J.A., Foster, J.G., 2011. Metaknowledge. *Science* 331, 721–725.
- Garcia-Herranz, M., Moro, E., Cebrian, M., Christakis, N.A., Fowler, J.H., 2014. Using friends as sensors to detect global-scale contagious outbreaks. *PLoS One* 9 (4), e92413.
- Giraldeau, L.-A., Valone, T.J., Templeton, J.J., 2002. Potential disadvantages of using socially acquired information. *Philos. Trans. R. Soc. B* 357, 1559–1566.
- Goldbaum, D., Mizrach, B., 2008. Estimating the intensity of choice in a dynamic mutual fund allocation decision. *J. Econ. Dyn. Control* 32, 3866–3876.
- Gomes, O., 2006. Organizational learning: a discrete choice approach. *J. Econ. Soc. Res.* 8, 39–59.
- Greene, W.H., 2003. *Econometric analysis*, fifth ed. Prentice Hall, Upper Saddle River, NJ.
- Henrich, J., Broesch, J., 2011. On the nature of cultural transmission networks: evidence from Fijian villages for adaptive learning biases. *Philos. Trans. R. Soc. B* 366, 1139–1148.
- Henrich, J., Gil-White, F.J., 2001. The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evol. Hum. Behav.* 22, 165–196.
- Herbert-Read, J.E., Krause, S., Morrell, L.J., Schaerf, T.M., Krause, J., Ward, A.J.W., 2013. The role of individuality in collective group movement. *Proc. R. Soc. B* 280, 20122564.
- Hobaiter, C., Poist, T., Zuberbühler, K., Hoppitt, W., Gruber, T., 2014. Social network analysis shows direct evidence for social transmission of tool use in wild chimpanzees. *PLoS Biol.* 12 (9), e1001960.
- Hoppitt, W., Laland, K.N., 2013. *Social Learning*. Princeton University Press, Princeton, NJ.
- Hruschka, D.J., 2010. *Friendship*. University of California Press, Berkeley.
- Kendal, J., Giraldeau, L.-A., Laland, K.N., 2009. The evolution of social learning rules: payoff-based and frequency-dependent biased transmission. *J. Theor. Biol.* 260, 210–219.
- Krakauer, D., 2011. Darwinian demons, evolutionary complexity, and information maximization. *Chaos* 21, 037110.
- Kurvers, R.H.J.M., Wolf, M., Krause, J., 2014. Humans use social information to adjust their quorum thresholds adaptively in a simulated predator detection experiment. *Behav. Ecol. Sociobiol.* 68, 449–456.
- Lahkar, R., Sandholm, W.H., 2008. The projection dynamic and the geometry of population games. *Games Econ. Behav.* 64, 565–590.
- Laland, K.N., 2004. Social learning strategies. *Learn. Behav.* 32, 4–14.

- Lorenz, J., Rauhut, H., Schweitzer, F., Helbing, D., 2011. How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci. USA* 108, 9020–9025.
- Manski, C., 1977. The structure of random utility models. *Theor. Decis.* 8, 229–254.
- 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. *Philos. Trans. R. Soc. B* 363, 3515–3528.
- Mesoudi, A., 2008. An experimental simulation of the “copy-successful-individuals” cultural learning strategy: adaptive landscapes, producer–scrounger dynamics, and informational access costs. *Evol. Hum. Behav.* 29, 350–363.
- Mesoudi, A., Whiten, A., 2008. The multiple uses of cultural transmission experiments in understanding cultural evolution. *Philos. Trans. R. Soc. B* 363, 3489–3501.
- O'Brien, M.J., Bentley, R.A., 2011. Stimulated variation and cascades: two processes in the evolution of complex technological systems. *J. Archaeol. Method Theor.* 18, 309–335.
- O'Brien, M.J., Boulanger, M.T., Buchanan, B., Bentley, R.A., Lyman, R.L., Lipo, C.P., Madsen, M.E., Eren, M.I., 2016. Design space and cultural transmission: case studies from Paleoindian eastern North America. *J. Archaeol. Method Theor.* <http://dx.doi.org/10.1007/s10816-015-9258-7> (in press).
- Perreault, C., Moya, C., Boyd, R., 2012. A Bayesian approach to the evolution of social learning. *Evol. Hum. Behav.* 33, 449–459.
- Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M.W., Fogarty, L., Ghirlanda, S., Lillicrap, T., Laland, K.N., 2010. Why copy others? Insights from the Social Learning Strategies Tournament. *Science* 328, 208–213.
- 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 Cogn. Sci.* 15, 68–76.
- Rieucau, G., Giraldeau, L.-A., 2011. Exploring the costs and benefits of social information use: an appraisal of current experimental evidence. *Philos. Trans. R. Soc. B* 366, 949–957.
- Rogers, A., 1988. Does biology constrain culture? *Am. Anthropol.* 90, 819–831.
- Rogers, E.M., 1995. *Diffusion of Innovations*, 4th ed. Free Press, New York.
- Salganik, M.J., Dodds, P.S., Watts, D.J., 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 854–856.
- Samuelson, P.A., 1941. The stability of equilibrium: comparative statics and dynamics. *Econometrica* 9, 97–120.
- Samuelson, P.A., 1947. *Foundations of economic analysis*. Harvard University Press, Cambridge, MA.
- Sela, A., Berger, J., 2012. Decision quicksand: how trivial choices suck us in. *J. Consum. Res.* 39, 360–370.
- Shalizi, C.R.S., Thomas, A.C., 2011. Homophily and contagion are generically confounded in observational social network studies. *Sociol. Methods Res.* 40, 211–239.
- Strandburg-Peshkin, A., Farine, D.R., Couzin, I.D., Crofoot, M.C., 2015. Shared decision-making drives collective movement in wild baboons. *Science* 348, 1358–1361.
- Thomas, A.C., 2013. The social contagion hypothesis: comment on ‘Social contagion theory: examining dynamic social networks and human behavior’. *Stat. Med.* 32, 581–590.
- Whiten, A., Goodall, J., McGrew, W.C., Nishida, T., Reynolds, V., et al., 1999. Cultures in chimpanzees. *Nature* 399, 682–685.
- Wolf, M.R., Kurvers, H.J.M., Ward, A.J.W., Krause, S., Krause, J., 2013. Accurate decisions in an uncertain world: collective cognition increases true positives while decreasing false positives. *Proc. R. Soc. B* 280, 20122777.
- Woolley, A.W., Chabris, C.F., Pentland, A., Hashmi, N., Malone, T.W., 2010. Evidence for a collective intelligence factor in the performance of human groups. *Science* 330, 686–688.