of the computer simulation by a questionnaire study of patients in the emergency department of Hahnemann Hospital, associated with the Drexel University College of Medicine.

The key point for the Bentley et al. map is that learning in the simulations reported, although based on a probabilistic choice model such as Bentley et al. postulate, relied only on individual experience to learn and estimate the benefits from the different strategies available to the agents. Different individual signals from the individual's own past experience provide the different signals necessary to support a Nash equilibrium in an anti-coordination game. Thus, these simulations belong in the extreme northeast of the Bentley et al. map. Indeed we do see, as predicted, an *r*-shaped adoption curve. However, in preliminary simulations that did assume social learning, no tendency to converge to a Nash equilibrium was observed. The shared average experience of each type of agent provides no different signal that could support the choice of different strategies in the congestion game, as a Nash equilibrium requires in such a game. These simulations also, consequently, disagree with the evidence from the questionnaire study, which were consistent with the hypothesis of Nash equilibrium.

Here is the conclusion. On the one hand, the game-theoretic discussion and the agent-based computer simulations indicate that social learning is inappropriate to congestion games, in that it cannot provide signals that lead different agents to choose different strategies, whereas a reliance on individual experience may do so. On the other hand, the evidence suggests that the Nash equilibrium hypothesis is descriptive of the actual experience of patients in emergency departments. Somehow, people seem to have focused their attention and learning on the sort of information that could give rise to individually satisfactory outcomes, so far as the "game" permits. Perhaps the Bentley et al. map, which seems to take the social or individual bias of decision-making as given, needs some refinement on the basis of non-cooperative game theory.

Cultural evolution in more than two dimensions: Distinguishing social learning biases and identifying payoff structures

doi:10.1017/S0140525X13001805

Alex Mesoudi

Department of Anthropology, Durham University, Durham DH1 3LE, United Kingdom.

a.a.mesoudi@durham.ac.uk

https://sites.google.com/site/amesoudi2/

Abstract: Bentley et al.'s two-dimensional conceptual map is complementary to cultural evolution research that has sought to explain population-level cultural dynamics in terms of individual-level behavioral processes. Here, I qualify their scheme by arguing that different social learning biases should be treated distinctly, and that the transparency of decisions is sometimes conflated with the actual underlying payoff structure of those decisions.

The target article by Bentley et al. provides an innovative conceptual map of human decision-making in a social context. It may well provide a valuable guide to researchers who are beginning to analyse "big data", such as aggregate online purchasing decisions or peer-to-peer social interactions. I see their scheme as complementary to calls from myself and others (Gintis 2007; Mesoudi 2011; Mesoudi et al. 2006; Richerson & Boyd 2005) to restructure the social and behavioral sciences around an evolutionary framework. Evolutionary "population thinking" concerns the exact problem that is addressed throughout the target article: how individual-level processes aggregate to form population-level patterns. In biology, the individual-level processes are natural selection, genetic mutation, Mendelian inheritance, and so on, and the population-level patterns include adaptation, speciation, adaptive radiation, serial founder effects, etc. Several decades of research has identified equivalent (but often different) individual-level processes in cultural evolution (Boyd & Richerson 1985; Cavalli-Sforza & Feldman 1981; Mesoudi 2011), such as conformist- or prestige-based social learning biases (Boyd & Richerson 1985), or the non-random generation of new cultural variation according to content-based inductive biases (Griffiths et al. 2008). Major advances in the social sciences can be made by borrowing tools from evolutionary biology to both explore the population-level consequences of these individual-level biases (e.g., population-genetic-style mathematical models; see Bentley et al. 2004), and quantitatively identify and measure those population-level patterns in real cultural datasets (e.g., phylogenetic methods; see O'Brien et al. 2001).

The conceptual map presented in Figure 1 of the target article similarly links individual-level decisions (the degree to which individuals rely on social or asocial learning, represented by the east–west axis) made within different environments (the transparency in payoffs, represented by the north–south dimension) to different population-level patterns, such as the popularity distributions shown in Figure 2. While recognising the heuristic value in such a simple conceptual map, I caution that these particular dimensions may over-simplify and obscure some key issues, which the cultural evolution literature has identified in recent years as being particularly important for understanding cultural change.

First, regarding the east–west axis, it seems problematic to treat all social learning as equivalent, or at least as having broadly similar population-level consequences. Different social learning biases, such as the aforementioned conformist, prestige and inductive biases, may have very different population-level signatures. Models suggest that prestige bias can generate a runaway process towards maladaptively extreme values, while conformity generates particularly pronounced within-group behavioral homogeneity (Boyd & Richerson 1985). Culturally driven copycat-suicide clusters require a particular combination of social learning processes in order to occur, primarily the rapid one-to-many transmission characteristic of the mass media plus a celebrity-driven prestige bias (Mesoudi 2009). Even restricting ourselves to the popularity distributions that Bentley et al. focus on, Mesoudi and Lycett (2009) showed that conformity and anti-conformity have very different consequences on population-level frequency distributions of discrete traits such as first names, with conformity creating a "winner-take-all" distribution where popular traits are made even more popular, and anti-conformity favoring traits of intermediate frequency. In sum, knowing that social (as opposed to individual) learning is at work might be useful, but knowing what kind of social learning is operating seems to be crucial, too.

Second, the north–south "transparent-opaque" dimension appears to conflate feedback error in agents' decisions with the actual payoff structure that underlies different decisions. The north–south dimension is said to represent "the extent to which there is a transparent correspondence between an individual's decision and the consequences (costs and payoffs) of that decision" (target article, sect. 2, para. 1). In other words, it concerns feedback error: high feedback error equals opaque decision-making, while low feedback error equals transparent decision-making. This is captured formally in the b_t parameter of equation 1. Yet this does not address the actual payoff functions underlying different choices (denoted by the function U in the equation, but unaddressed in the map).

With respect to actual payoffs, there are several possibilities: There may be a single objectively best option and many bad options, or there may be several equally good options, or there may be no functional correspondence between choice and payoff whatsoever. In adaptive landscape terms (Wright 1932), these correspond to a unimodal, a multimodal (or rugged), and a flat landscape, respectively. The shape of this underlying adaptive landscape is logically independent to how well that payoff structure can be perceived by agents (i.e., the vertical transparency-opaqueness dimension).

I would argue that one cannot understand the consequences of transparent versus opaque feedback error without also considering the actual shape of the underlying adaptive landscape. Opaque feedback in a flat (neutral) landscape will be unproblematic, because all options are equivalent and feedback error is unimportant. However, opaque feedback in a rugged landscape will be very problematic, given the need to find one of a small number of fitness peaks. Conversely, perfectly transparent feedback may be problematic in a rugged landscape because it may lead learners to locally optimal but globally sub-optimal peaks/decisions, whereas the error intrinsic in slightly opaque feedback might lead learners, by chance, off their sub-optimal peak and onto a higher peak elsewhere in the landscape.

Experiments and models show that the shape of the adaptive landscape can significantly affect both people's choices and the aggregate outcome of those choices, quite independently of feedback error (Mesoudi 2008; Mesoudi & O'Brien 2008a; 2008b). Yet Bentley et al. appear to conflate these two distinct dimensions. For example, the neutral models discussed in section 2.3.1 (and analysed in Bentley et al. 2004) surely concern the case where the actual payoffs of all possible choices are equivalent, rather than where payoffs are opaque.

Naturally, all heuristic schemes such as the one presented by Bentley et al. are simplifications, and their value lies in that simplicity, as researchers grapple with the enormous datasets generated in the modern age. At the same time, oversimplification can sometimes lead to the wrong answer. I suspect that distinguishing between different social learning biases, and considering payoff structure as well as feedback error, might be crucial in avoiding those wrong answers.

Using big data to predict collective behavior in the real world[1]

doi:10.1017/S0140525X13001817

Helen Susannah Moat,[a,b] Tobias Preis,[b] Christopher Y. Olivola,[c] Chengwei Liu,[b] and Nick Chater[b]

[a]*Department of Civil, Environmental and Geomatic Engineering, University College London (UCL), London, WC1E 6BT, United Kingdom;* [b]*Behavioural Science Group, Warwick Business School, The University of Warwick, Coventry, CV4 7AL, United Kingdom;* [c]*Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA 15213.*

Suzy.Moat@wbs.ac.uk Tobias.Preis@wbs.ac.uk olivola@cmu.edu
Chengwei.Liu@wbs.ac.uk Nick.Chater@wbs.ac.uk
http://www.wbs.ac.uk/about/person/suzy-moat/
http://www.wbs.ac.uk/about/person/tobias-preis/
https://sites.google.com/site/chrisolivola/
http://www.wbs.ac.uk/about/person/chengwei-liu/
http://www.wbs.ac.uk/about/person/nick-chater/

Abstract: Recent studies provide convincing evidence that data on online information gathering, alongside massive real-world datasets, can give new insights into real-world collective decision making and can even anticipate future actions. We argue that Bentley et al.'s timely account should consider the full breadth, and, above all, the predictive power of big data.

Modern everyday life is threaded with countless interactions with massive technological systems that support our communication, our transport, our retail activities, and much more. Through these interactions, we are generating increasing volumes of "big data," documenting our collective behavior at an unprecedented scale.

Bentley et al. provide a timely account of the role of big data in the study of collective behavior. They offer a comprehensive analysis of what our interactions on the Internet, in particular using social network sites such as Facebook and Twitter, can tell us about how information flows throughout the large and complex network of human society. While we agree that this insight into the structure of social connections is important, we emphasize that big data do not only come from online social networks. We note a number of recent studies providing evidence that big data can tell us much more about real-world collective decision making than has been acknowledged in Bentley et al.'s account, and can even allow us to better anticipate collective actions taken in the real world.

For example, human decision making often involves gathering information to determine the consequences of possible actions (Simon 1955). Increasingly, we turn to the Internet, and search engines such as Google in particular, to provide information to support our everyday decisions. Can massive records of our search engine usage therefore offer insight into the previously hidden information-gathering processes which precede real-world decisions taken around the globe? Recent results suggest that they can. A series of studies have shown that search engine query data "predict the present," providing a measurement of real-world behavior often before official data are released (Choi & Varian 2012). Correlations between search engine query data and real-world actions have been demonstrated across a range of areas such as motor vehicle sales, incoming tourist numbers, unemployment rates, reports of flu and other diseases, and trading volumes in the U.S. stock markets (Askitas & Zimmerman 2009; Brownstein et al. 2009; Choi & Varian 2012; Ettredge et al. 2005; Ginsberg et al. 2009; Preis et al. 2010).

Further studies have illustrated that data on online information gathering can also anticipate future collective behavior. Goel et al. (2010) demonstrated that search query volume predicts the opening weekend box-office revenue for films, first-month sales of video games, and chart rankings of songs. Our own investigations have suggested that changes in the number of searches for financially related terms on Google (Preis et al. 2013) and views of financially related pages on Wikipedia (Moat et al. 2013) may have contained early warning signs of stock market moves.

In a recent study, we exploited the global breadth of Google data to compare information-gathering behavior around the world. Our analysis uncovered evidence that Internet users from countries with a higher per capita gross domestic product (GDP) tend to search for more information about the future rather than the past (Preis et al. 2012). For 45 countries in 2010, we calculated the ratio of the volume of Google searches for the upcoming year ("2011") to the volume of searches for the previous year ("2009"), a quantity we called the "future orientation index." We found that this index was strongly correlated with per capita GDP. In ongoing work, we seek to better understand whether these results reflect international differences in decision-making processes. Perhaps, for example, a focus on the future supports economic success.

Aside from search data, other research has provided evidence that the massive datasets generated by our everyday actions in the real world can also support better forecasting of future behavior (King 2011; Lazer et al. 2009; Mitchell 2009; Vespignani 2009). Large-scale datasets allow us to look for patterns in collective behavior which might recur in the future, similar to the way in which we as individuals rely on the statistical structure we have observed in the world when trying to forecast consequences of decisions (Giguère & Love 2013; Olivola & Sagara 2009; Stewart 2009; Stewart et al. 2006). For example, analysis of data collected through daily police activities has shown that the occurrence of a burglary results in a short-term increase in the probability that another burglary will occur on the same street, with implications for behavioral models of how these crimes are committed (Bowers et al. 2004; Johnson & Bowers 2004; Mohler et al. 2011). Such insights have been captured in predictive policing systems which aim to deploy police to areas before an offence