

Horizon

Where diversity comes from and why it matters?

SCOTT E. PAGE*

Center for the Study of Complex Systems, University of Michigan, USA

Abstract

In this essay, I describe some of the benefits of cognitive diversity in a complex world as well as the origins of that diversity. The essay has two main parts sandwiched between a brief description of what I mean by diversity and complexity, as well as a brief discussion of whether social systems produce sufficient diversity. In the first part, I describe models that provide insight into why we see the levels of diversity that we do. These models rest on social psychological foundations but borrow ideas from economics as well as population genetics. In the second part, I describe the functional benefits of diversity. I show how diverse predictive models can make a collection of people better able to make accurate predictions, how diverse perspectives and heuristics can enable groups of problem solvers to find innovative new solutions to problems, and how diverse behaviors and representations of the world can make a society more robust. Copyright © 2014 John Wiley & Sons, Ltd.

I don't like work,—no man does—but I like what is in the work—the chance to find yourself. Your own reality—for yourself, not for others—what no other man can ever know. They can only see the real show and never tell what it really means.

—Joseph Conrad, *The Heart of Darkness*

Two questions animate the research that I will describe in this essay: why do people differ in their behaviors and mental models to the extent that they do and how do these differences matter? Does diversity merely create problems, leading to conflict and a lack of trust, or do our differences make us collectively more innovative and more creative, better able to solve difficult problems and make sense of our complex world? Do they make our societies more resilient? These are large and messy questions, so the contributions that I make will necessarily be modest both in scope and in depth.

Clearly, the functional implications of diversity cut both ways. Whether one looks at the level of communities or nations, one finds that trust (Putnam, 2007), contributions to public goods, and measures of social capital often decline with identity heterogeneity.¹ Looking across countries, one finds that cultural diversity increases the likelihood of conflict, but economic diversity promotes trade. If one looks within a country, one can see evidence of tension, criminal behavior, and lower levels of funding for public goods. These all result from identity diversity. At the same time, however, as Alberto Alesina and Eliana La Ferrara (2005) note, the most culturally diverse cities, here they consider New York and Los Angeles, but one

could also mention London, Paris, or Hong Kong, are also “constant producers of innovation in the arts and business.” Further, when one separates birthplace diversity from ethno linguistic diversity, the former positively correlates with aggregate economic productivity at the city level (Alesina, Harnoss, & Rapoport, 2013). Overall then, the evidence would seem to suggest that differences in backgrounds help the economy, provided that people can just get along.²

Evidence and theory also demonstrate that cognitive diversity improves problem solving and prediction abilities for small groups (Page, 2007). This raises the possibility that political systems in diverse countries could be higher functioning than in homogeneous societies, but this requires knowledge of how to leverage that diversity (Laitin & Jeon, 2013; Landemore, 2013). This way of thinking presents diversity as a part of the solution as opposed to the underlying cause of political problems (Huckfeldt et al., 2004). To wit, the opening sentence of Jim Johnson and Jack Knight's (2011) *The Priority of Democracy* reads, “Politics, in large part, is a response to diversity.” I accept that differences in preferences and identities underlie the splitting up of so many pies, but I claim too that cognitive diversity can enlarge the pies before they are allocated.

Let me turn briefly to the proposition that people differ. Stated as such, the claim is indisputable but lacks scholarly focus. One must choose a set of dimensions on which to identify differences and a set of mechanisms that cause them. We differ in our preferences and interests, in our beliefs, in how we perceive and categorize the world around us, and in the heuristics and tools we possess to solve problems. We also differ in our personalities, our cultural traits, and our physical

*Correspondence to: Scott E. Page, Center for the Study of Complex Systems, 321 West Hall, Ann Arbor, MI 48104A.
E-mail: spage@umich.edu

¹See Aligica (2014) Chapter 1 for an excellent, brief summary.

²Other models posit that productivity is single peaked in diversity—see Ashraf and Galor (2013).

characteristics. And, as Conrad so aptly observes, apart from our physical traits, these differences remain largely opaque to others. Whether contemplating a book, a political campaign, a Federer–Nadal rally, or the fantastic imagery of Hieronymus Bosch, each person has a unique, only partly communicable experience. I might add that the lack of translatability extends to one's research. People only see the final draft—the real show—they do not see the contradictory ideas and impulses that lie hidden beneath.

In this essay, my focus will be on cognitive and behavioral differences—differences in how people solve problems, make predictions, and behave in strategic settings. Not personality. Not identity. Those types of diversity will correlate with and often cause the types of diversity I consider, but I will leave them for others with greater expertise. My focus will be on those dimensions that interest economists, political scientists, and organizational theorists: differences in how people think about problems. The research I describe differs from the standard approaches taken by social psychologists in two ways: I tend to rely on formal models, and I emphasize the role of strategic incentives.

My desire to study diversity using models represents a slight departure from the norm. Models rely on simplifications, so as to elucidate the fundamental causal structures. Within economics at least, a common assumption has been to rely on *representative agents*—to assume that every person is identical, to sweep all our differences under a giant rug. Representative agent models have proven of great use, but they have limits. And because of those limits, most models, even in macro economics, now allow for some heterogeneity.

The value of embracing diversity and including it in one's models can be found in the work of Elinor and Vincent Ostrom. They spent decades exploring empirically as well as theoretically how heterogeneity—in beliefs, talents, endowments, and preferences—influenced efforts to solve collective action problems and common pool resource problems (Aligica, 2014; Ostrom, 2005). Collective action problems range in scope from small villages managing water usage to the current global effort to reduce carbon emissions. The Ostroms found that workable solutions require deep understandings of human differences and how those differences facilitate and frustrate institutional interventions. Without an understanding of human differences, no institution, no Leviathan if you will, can prevent our lives from being cold, brutish, and short.

Taking an even broader perspective, if one accepts that first, the modern world is both complex and growing more complex at each moment, and second that individuals are limited in what they can understand, then one is led rather quickly to the idea that our continued success and perhaps even our survival depend on our ability to leverage our differences (Axelrod & Cohen, 2001). Thus, heterogeneity—rather than seen as a nuisance that creates variation and might even undo statistical significance—should be seen as of functional value to social systems.

The opportunity to describe these functional contributions motivated my writing this essay. So often within the social sciences, when the topic of diversity arises, discussions turn to the many moral, political, and legal dimensions that arise within a society as a result of identity diversity and, in particular, the historical injustices that often occurred as a result of those differences.

My focus here will be orthogonal to those concerns, but is in no way intended to diminish them. My interest will be in showing how cognitive differences can improve the complex systems in which we live by making them more innovative and productive, by showing that diverse ways of seeing the world can help predict outcomes of complex processes, and that even though diversity can contribute to large events such as market crashes and social uprisings, it (paradoxically) also contributes to robustness by creating multiple routes of adaptation (Bednar, 2009; Frank, 1997).

At this point, I will take a moment to define what I mean by complexity, which has literally scores of formal mathematical definitions (Mitchell, 2009; Page, 2010). A good stepping off point is to distinguish between complexity and the complex systems that produce it. Complex systems consist of diverse, networked actors with adaptive interdependent behaviors (Miller & Page, 2007; Page, 2008a). Economies, ecosystems, the brain, and middle schools can be classified as complex systems. In this way, diversity and complexity go hand in hand (Page, 2010). Without diversity, one cannot have much complexity. But without complexity, diversity becomes mere variation, noise about the mean. Within a complex system, diversity creates patterns and helps drive transitions (Nowak & Vallacher, 1998).

Complex systems produce four classes of outcomes: equilibria, simple patterns, randomness, and complexity (Wolfram, 2002). Wolfram defines complexity as lying between order, such as the alternating sequence 1010101010 and randomness, such as might be produced by flipping a coin. A complex sequence lies between these two in that it has some patterns, but the patterns are longer and not predictable. Others define an outcome or process to be complex if it is difficult to describe, engineer, emerge, explain, or predict (see Page, 2010). For example, many define complexity as the *minimal description length*, the number of words or symbols that are needed to describe a process or an outcome. Definitions like this require agreement on an alphabet, often just numbers. The sequence 0000000 can be described in two numbers as 0^7 , which means seven zeros. The sequence 10101010 requires three. It is written as $(10)^5$ or five 10s. The sequence 010010001 would require more than three numbers; hence, it is more complex. This third sequence would also be harder to engineer or explain, and it would require a more complex system to produce from the bottom up; that is, it would seem less likely to emerge.

These two definitions, between ordered and random, and difficult to explain, engineer, emerge, or predict, will be what I mean when I say complex. Note that these definitions loosely align. Something that is between ordered and random tends to be hard to explain, engineer, or predict. What is perhaps equally important to keep in mind is that complexity differs from *chaos*: extreme sensitivity to initial conditions, *complicated*: lots of interdependence but no adapting parts and therefore no surprise, *uncertain*: environments with a random component, and *difficult*: problems that require many searches or steps to solve (Page, 2008b). Some examples help to clarify the distinctions. A double pendulum is chaotic, the roll of a die is uncertain, a mechanical watch is complicated, and solving a crossword puzzle is difficult. None are complex.

Complexity science differs from more “normal” science by focusing on multiple processes and the interdependencies

between them in a dynamic setting (Nowak & Vallacher, 1998). Outcomes in complex systems are often, again by definition, difficult to predict. In mathematics, one plus one equals two, but within a complex system, one plus one can equal three. And sometimes, to extend Philip Anderson's (1972) "more is different" meme, the result of interactions can produce new emergent phenomena, so one plus one equals penguin or even penne arrabiata.

Given the mysterious aggregation properties of complex systems, a second goal in this essay will be to describe how considering multiple process simultaneously—these could be at the individual or group level—can lead to different results than if one considers them in isolation. To that end, I will describe a model that combines two rather commonly assumed social psychological motives—a desire for individual level consistency and a drive for social conformity. Each of these forces on its own produces homogeneity. Together they will produce and maintain diversity. Before presenting that model, I will begin by contextualizing that result in a broader discussion of how people choose strategies, actions, and behaviors. I will conclude the paper by describing in some detail two functional properties of cognitive diversity: its potential contribution to collective wisdom and its value in innovative problem solving. Neither of these properties will surprise social psychologists, but my methodological route to establishing them may. My approach will also produce unexpected implications. It will suggest that as our world becomes more complex—and by that I mean less predictable, higher dimensional, and more interdependent, the existence of collective cognitive diversity will become even more important. Thus, understanding what produces and maintains cognitive diversity takes on practical importance as the challenges that we face—climate change, poverty, and managing the world's financial system—become more complex. To survive, we must do more than contain multitudes, we must leverage them.

I recognize the implicit optimism in my perspective. By focusing on the diverse cognitive tools that people bring to bear on problems and situations, I am ignoring differences in personality types, preferences, and identities. In any given group interaction, all of these dimensions of diversity may be in play. Even if functional cognitive diversity correlates with better outcomes, other types of diversity may impede progress. Even more troubling, separating out the distinct effects of each type of diversity may be difficult (van Knippenberg & Schippers, 2007). These caveats notwithstanding, controlled experiments do find interactions between functional diversity and performance (Pelled, Eisenhardt, & Xin, 1999). Further, recent analyses that consider enormous data sets involving academic papers, patents, and collective problem solving demonstrate the benefits of cognitive diversity rather convincingly (Adamic et al., 2010; Schilling & Green, 2011; Uzzi, Mukherjee, Stringer, & Jones, 2013), as does decades of economic research on the performance of workers in cities (Florida, 2002; Glaeser, 2010).

The remainder of this essay consists of three parts. In the first part, I describe how diversity might be produced and maintained within complex systems. In the second part, I highlight three functional contributions of diversity: more accurate predictions, better solutions to hard problems, and greater resilience. And finally, in a brief discussion at the end, I describe some possible future research questions.

THE PRODUCTION AND MAINTENANCE OF DIVERSITY

When studying human behavior, social scientists, be they economists, sociologists, or psychologists, tend to focus on a particular domain such as whether people take risks, vote, cooperate, or act altruistically—perhaps allowing someone with just a few items to cut ahead in the supermarket line. Within a given domain, human behavior can be explained with a variety of theories. Consider cooperative behavior. One can find evidence of intrinsic altruism (Fehr & Schmidt, 1999), of conformity (Asch, 1951; Ross, Bierbrauer, & Hoffman, 1976), and of rational self-interest stemming from expected future interactions (Myerson, 1997). In many situations, all three may operate.

In this section, I will consider a person interacting in multiple domains. Thus, I will characterize an individual as a vector of behaviors, with each variable corresponding to a behavior in a specific context. This will enable me to discuss diversity both within and across individuals. A population of individuals can then be characterized as a collection of vectors, or bundles if you will, of behaviors.

If we assign letters to behaviors (A = altruism, C = cooperate, S = selfish, etc.), then we can compare bundles from a purely mathematical perspective. An individual's bundle might consist of all As or all Cs, the latter representing someone who always acts cooperatively. A bundle could consist of a mix of A's, C's, and S's. This would represent someone who is inconsistent or accepting of dissonance. We can think of each of these behaviors being linked to contexts or domains. This enables us to look at the population level and see whether people have coordinated their behavior.³ In the models that I describe here, my co-authors and I focus on strategic situations in which individuals benefit materially as well as psychologically from coordinating. This construction includes only a small desire for *distinctiveness* on the part of individuals (Brewer, 1991), a point that I return to later after presenting our theoretical results.

Aside: Measuring Diversity

Lord Kelvin famously said that *when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind*.⁴ Many (including me) find this an overstatement of the relative contribution of quantitative science, but for the types of questions that I ask—*theoretical* questions—quantitative measures are necessary to derive results. Measures also facilitate the communication of logical relationships as well. A claim that increases in connectivity decrease diversity can be more effectively communicated if both connectivity and diversity have formal definitions.

One can choose from a variety of diversity measures. These range from the number of types to more complicated entropy-based measures from information theory.⁵ Here, I will use the *diversity index*.⁶ To construct the measure, assume that there

³In economics, this sort of coordination occurs whenever each player chooses the same equilibrium strategy in a game.

⁴Lord Kelvin, *Popular Lectures and Addresses* Vol. 1.

⁵See my book *Diversity and Complexity* for a survey of diversity measures.

⁶This measure might be familiar to readers as Simpson's Index, Hirschman-Hirshman Index, the Effective Number of Parties, or the Effective Number of Types.

exists a population of individuals each of whom has a type. The assumption that types exist can be problematic. Boundaries can be blurry—biologists argue at length over what constitutes a species—or even arbitrary—philosophers dispute whether types exist or are socially constructed. Those caveats notwithstanding, we can often characterize people by types (democratic, cooperater, voter, fan of the opera, or punisher) with sufficiently high degrees of certainty to make the distinction between types meaningful.

The *diversity index* equals the inverse of the probability that two randomly chosen members of the population have the same type.⁷ The measure is constructed so that if there exist five equally likely types, the diversity index equals five.

If the types are not equally distributed, then the diversity index will be less than the number of types. For example, a recent survey showed that the colors of new cars sold in the USA in 2012 were distributed as follows: white (22%), silver (20%), black (19%), gray (12%), red (9%), beige/tan (8%), blue (7%), green (2%), and all others (1%).⁸ If we count “all others” as a type, then there exist 10 types. But because the cars are not spread evenly across those 10 types, the diversity index equals 6.2. So, even though cars come in 10 colors, it is “as if” there were six colors of cars.

This index plays a prominent role in how scientists view the health of systems. Here are just three examples. Economists use the inverse of the diversity index to determine the competitiveness of a market. This measure provides a crude measure of monopoly power. Up until relatively recently, the United States Department of Justice horizontal merger guidelines set thresholds for changes in the concentration index that trigger a presumption of too much market power. Low diversity therefore implies exploitation. Similarly, political scientists use the diversity index as a measure of political competitiveness, where the parties represent types and the percentage of each party equals its vote percentage. A lack of sufficient diversity can be seen as a signal of potential inefficiency. Thus, a democracy with a diversity index near one may be a democracy in name but not in practice. Third, ecosystems with less biodiversity are seen by many as more fragile (Tilman, 1999).

GameS Theory and Behavioral Ensembles

The first behavioral diversity model that I describe considers individuals who play multiple games at the same time with different people. This model will provide a partial explanation for why we see diversity across groups. I shall refer to this model as the gameS theory model (Bednar & Page, 2007). To highlight the innovation in our approach, consider by way of contrast, the canonical game experiment run by economists, political scientists, or psychologists. A population of people plays a game, such as the Prisoners’ Dilemma. Each person learns some strategy to play in that game, and the experimenter measures outcomes, rates of learning, and whether the population attains a Nash

equilibrium—where each person’s strategy is optimal given the strategies of the others.

In Bednar and Page (2007), we extend the standard setup to allow for *game ensembles*. In our formulation, each person either chooses or learns strategies to play in each of several games. Importantly, rather than analyze each game in isolation, we consider the potential for contextual effects in which behavior in one game bleeds over into others.⁹

Our first investigation used a mathematical game theoretic model as well as an agent-based model. This second methodology enabled us to experiment with multiple environments and contexts and to nail down our core logical argument (Brown et al., 2005; Miller & Page, 2007). In the original paper, we referred to our approach as gameS theory to highlight our inclusions of multiple games. The strategies that individuals choose in these multiple games were called *behavioral repertoires*.

To give this discussion more formal structure, consider how an individual might construct her repertoire. She could treat each game within the ensemble in isolation and either select or evolve an optimal strategy for that game. That assumption would allow for an individual to use radically different strategies in different games. Alternatively, she could choose general purpose strategies that apply across multiple games, or she might choose distinct but similar strategies that balance increases in cognitive costs with reductions in payoffs.

In our initial paper (Bednar & Page, 2007), we relied on a representation of strategies called a Moore Machine in which strategies consist of *states* and *transitions* between those states. We let the number of states serve as a proxy for cognitive costs.¹⁰ The model produced substantial ensemble effects on behavior. We found that behavior in a particular game often depended on the other games in the ensemble.

Taken at face value, these findings provide a more bottom-up explanation for empirical evidence of behavioral differences in similar games across different cultures (Henrich et al., 2001). Those studies highlight correlation between levels of cooperation and equal sharing with the extent of market concentration and cooperative types of economic activities. The gameS theory model suggests an alternative causal chain to explain that regularity: market institutions induce equal splitting behaviors (game theorists will be quick to note that equal splitting is the Nash equilibrium of a bargaining game), and cooperative games should similarly induce cooperative behaviors. Confronted with novel situations, individuals invoke familiar behavioral routines, that is, cooperation or equal splitting.

A deeper question, one that lies in the intersection of cognitive psychology and social psychology, is the extent to which so-called human cognitive biases might be contingent on culture and even be partly a product of the informational and strategic situations within which people interact. Answering that

⁹The assumption that people attempt to use similar strategies across games builds from theories in personality research and social psychology (e.g. cognitive dissonance), suggesting that a person will prefer to act in a consistent fashion. Evidence for dissonance reduction exists in many studies, although it is not as strong as many believe (Ajzen & Fishbein, 1980; Festinger, 1957; Festinger & Carlsmith, 1959; van Overwalle & Jordens, 2002; Ross & Ward 1995). Evidence from the political realm shows that an individual may ignore the opinions of experts that do not align with his or her world view (Darmofal, 2005).

¹⁰Samuelson (2001) used a similar approach in considering allocation of cognitive effort across two games.

⁷If p_1 is the probability that a person is of type, then the probability that two are of type one equals $p_1 * p_1$. For explanatory purposes, I am assuming here that the two members are chosen with replacement or that the population is infinite. For those that desire more formality, given the distribution across types (p_1, p_2, \dots, p_n), the diversity index equals $DI(\vec{p}) = \frac{1}{\sum_{i=1}^n p_i^2}$.

⁸Source: PPG Industries.

<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr> <td style="text-align: center; padding: 5px;">7,7</td> <td style="text-align: center; padding: 5px;">2,10</td> </tr> <tr> <td style="text-align: center; padding: 5px;">10,2</td> <td style="text-align: center; padding: 5px;">4,4</td> </tr> </table>	7,7	2,10	10,2	4,4	<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr> <td style="text-align: center; padding: 5px;">7,7</td> <td style="text-align: center; padding: 5px;">4,14</td> </tr> <tr> <td style="text-align: center; padding: 5px;">14,4</td> <td style="text-align: center; padding: 5px;">5,5</td> </tr> </table>	7,7	4,14	14,4	5,5	<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr> <td style="text-align: center; padding: 5px;">7,7</td> <td style="text-align: center; padding: 5px;">4,11</td> </tr> <tr> <td style="text-align: center; padding: 5px;">11,4</td> <td style="text-align: center; padding: 5px;">5,5</td> </tr> </table>	7,7	4,11	11,4	5,5
7,7	2,10													
10,2	4,4													
7,7	4,14													
14,4	5,5													
7,7	4,11													
11,4	5,5													
Prisoners' Dilemma	Strong Alternation	Weak Alternation												

Figure 1. Examples of games in ensembles

question would require many cross-cultural experimental studies. Here, I am not attempting to make an empirical point but instead emphasizing the logic of our model, which is builds from the ideas of (Gigerenzer, 2000). We assume that people develop behaviors in response to their environment and that for the most part, those behaviors function adequately if not nearly optimally. Yet, at times, because the behaviors are rule based, they will fail to be optimal. And, therefore, they can appear to be biases.

This idea that psychology needs to branch out and consider more diverse populations before characterizing a behavior as a *human bias* is not our own. It finds much fuller voice in an article by Medin, Bennis, and Chandler (2010) who argued for the broadening of subject pools.¹¹ Pushing that line of argument even further, Henrich, Heine, and Norenzayan (2010) characterized the Americans who have provided subject pools for many of the studies on cognitive biases as the “weirdest people in the world.” To reiterate, we believe that “weirdness” in a given setting may be partly a function of the other domains within which one finds oneself. In other words, Americans may be weird in the lab because of the lives we live outside of it.

The theoretical papers that I have discussed rely on mathematics and simulations. The specific learning models assumed rather primitive behavior such as best response rules and replicator dynamics. We therefore felt obliged to study those same contexts experimentally. Together with Yan Chen and Tracy Liu, Jenna Bednar, and I took the same basic framework to the laboratory (Bednar, Chen, Liu, & Page, 2012). There, we found that individuals often played common strategies across games even if those strategies may not be optimal in both games. Consider the three games shown in Figure 1. The first is the familiar Prisoners' Dilemma; the second, a game we call Strong Alternation, as individuals do best by alternating who gets the better prize; and the third, a game we call Weak Alternation. Notice that in Weak Alternation, alternating in the off diagonal boxes gives a higher payoff than playing in the upper left but only by a small amount.

In the experiments, individuals played two games, each with a different player. We found strong empirical evidence for behavioral spillovers in these games. How people played in one game was not independent of what other games they were playing. Perhaps more striking than the more general result were the findings from particular games. For example,

when individuals played two Weak Alternation games, or when they played one Weak Alternation game and one Strong Alternation game, they were twice as likely to alternate between the off diagonal corners (the optimal strategy was played in each case approximately 36% of the time) than when they played Weak Alternation with the Prisoners' Dilemma game, where they alternated only 18% of the time. In other words, if you are alternating in one context, you are twice as likely to alternate in another.

To summarize the experimental results, the behaviors that emerge, and I use that word purposefully as we believe them to be more the emergent result of a social process than rationally chosen by the individual subjects, depend on the games within the ensemble. Independent and related multi game experiments by others support that intuition: namely that people appear to juggle many strategies in their heads and that what strategy comes to dominate in one game can be influenced by the strategies used in other games (Cason, Savikhin, & Sheremeta, 2012). Thus, ensemble effects may contribute to an understanding of the behavioral diversity that we see across populations.

As a follow up, we have begun to explore whether behavior depends on the order in which games were added to an ensemble. This question is of interest to people who study institutions because evidence from natural experiments shows that early institutions can have a large effect on future behavioral patterns as does research on founder effects in organizational behavior and even regional behavior. Dell (2010), for instance, found that the former presence of *mita*, a form of forced labor arrangements in mines, within a region persists today. In modern day Peru, regions that had *mita* exhibit less equality and provide fewer public goods. This is an example of *behavioral path dependence*—where previous games influence current outcomes (Page, 2006).

In one preliminary experiment, we have found that individuals who learned to alternate in the Weak Alternation game tended to sustain that alternation when playing a Prisoners' Dilemma game *even though they were playing the Prisoners' Dilemma game with a different player*.¹² An informal survey of experimental economists revealed that they had rarely if ever seen individuals alternate in the Prisoners' Dilemma when that game is played in isolation except for brief periods.

¹²We see this phenomenon as distinct from the sort of uncertainty reduction motivation identified by Sherif and Harvey (1952) and having more to do with learning transfer.

¹¹See also Kitayama, Snibbe, Markus, and Suzuki (2004).

Conformity and Consistency

In both our theoretical and experimental research, we were surprised by the level of diversity within populations. This piqued our interest. Even if we ran our computer models for thousands of iterations, the populations often did not fully coordinate as predicted by the mathematical model. In the experiments, when playing multiple games, unless the game had what game theorists call a *dominant strategy efficient equilibrium* (in layman's terms—a totally obvious best choice), the population kept churning. Individuals kept switching their actions. Here was an example of where more was different. That difference merited further study.¹³

We had constructed the model to explore diversity across populations, and it was producing diversity within populations as well. Jenna Bednar deserves credit for the key observation. She noticed that the gameS theory framework created dual incentives—a drive towards *consistency* at the individual level in the form of consistent strategies across the games in the ensemble—and a drive toward *coordination* at the population level in the form of equilibrium behavior. These incentives do not necessarily compete—individuals could coordinate on a common, consistent set of behaviors.

To explore why the model did not converge, we constructed a simplified version of the game ensemble model that assumes the two forces: one socially oriented that we refer to as *conformity* and one self-oriented that we call *consistency* (Bednar, Bramson, Jones-Rooy, & Page, 2010). In the starkest version of the model, behavior in a given context takes a value of either zero or one. For example, zero might represent selfish behavior, and one might represent cooperative behavior.

We initialize the model by randomly assigning each individual random strings of zeros and ones to represent his or her behavior. One person therefore might be represented by 1011110 and another by 0111010. We then apply the self-oriented and social-oriented psychological processes. We capture social conformity as follows: we randomly choose two people and an attribute. If the two people have the same value on the attribute, nothing happens. If they disagree on that attribute, one of the two people switches its attribute to match that of the other. To make the model more realistic, we assume a small probability of randomly choosing an attribute's value (less than a 1% chance). In this way, some interactions will produce two people with different attribute values. This randomness can be thought of as error, experimentation, or even the mildest desire for distinctiveness.

Our assumption of such a large probability of conforming requires an explanation. In the model, individuals have many attributes and switch at most one attribute. We could lower the probability of coordinating or becoming consistent, but that would only slow the process. It would not alter the results. Second, our interest was for games in which people have strong incentives to coordinate. When I meet a friend, I have strong incentives to greet him or her in the same way he or she greets me—lest when I put my hand to shake, I poke out his or her eye when he or she begins her bow. I also have strong incentives to drive on the same side of the road as everyone else or to go

the same party as my friends. I am less likely to switch my opinions on gun control or abortion. Other choices, such as what clothes I choose to wear, what books I read, and what music I listen to lie between these two extremes.

One need not break out the mathematics to know what happens in a model with only social conformity operating. After a relatively short period, for any given attribute, most individuals converge on the same value. Across attributes, however, the values will differ. For example, a particular run of the model might produce a population in which almost everyone can be represented by the string 1011101. The percentage of people who do not take the more common value on an attribute increases linearly with the amount of randomness. Therefore, such a model does not produce much diversity across a society, but it does produce diversity within each individual.

To capture the desire for consistency, we randomly choose an individual and randomly select two of that individual's attributes. If those two attributes have the same value, he or she does nothing. If they disagree, then he or she changes one to match the other. As before, we also assume a small amount of randomness. If only consistency operates, then the process will be similar to the conformity process with the number of attributes playing the role of the number of individuals. The result will be that each individual tends to assign the same value to each attribute but that people differ in which attribute; for example, one person might be 0000000, whereas another will be 111111, with the caveat that the noise will cause slight deviations.

To summarize, the conformity model works at the social level. It produces homogeneity across individuals. The consistency model operates at the individual level. It produces homogeneity within individuals. When one turns on both processes, one would expect homogeneity at both levels: consistent people who are all consistent in the same way. After all, neither process promotes diversity. One creates individual level consistency. The other produces social conformity.

Yet, when we turn on both conformity and consistency, the model does not produce homogeneity. Instead, it produces diversity. We have two well-established forces in play—an individual desire for consistency and a social desire for conformity. Each produces homogeneity but at different levels. Yet, when you construct a society with both forces in operation, you find that small errors magnify to such an extent that you do not get homogeneity at the population level. You get diversity. Why?¹⁴

The answer lies in the interaction between the processes. Diversity persists because random fluctuations propagate in two directions. First, fluctuations spread across the population as one person copies another. Second, they spread within a person as the inner drive for consistency can cause a behavior in one context to be applied in another. Once applied in that new context, that behavior could be applied by someone else and spread.

Here is an example of how that could happen. In Kuala Lumpur, many people ride motorbikes to work. Some people wear helmets. Some younger males wear baseball hats, and

¹³We measured strategic churning using an entropy measure over outcomes. We found high levels of entropy over outcomes in the multiple game experiments, which imply that outcomes did not converge.

¹⁴As an example, in one set of numerical simulations, with 100 individuals, 10 attributes, and 6 possible attribute values, we found the diversity index to equal to approximately three. The most it can be is six, and if people had coordinated, it would equal to one. See Page, Sander, and Schneider-Mizell (2007) for a mathematical analysis of convergence.

these are typically worn backward so as to not fly off into the air. That behavior is no different than in many other cities. Let us label that behavior *wear clothing backward on motorbike*. That might not seem like a behavior that would spread to another domain, yet it has.¹⁵ Kuala Lumpur has a humid tropical climate and moderate to substantial pollution. To protect their clothes and to avoid overheating, many people wear thin jackets backwards when riding their motorbikes. This functional adaptation protects their clothing from the dirty air and occasional splash but at the same time keeps them cool. Does it look silly to wear a coat backwards? Yes. But once everyone is doing it, it becomes acceptable. The social acceptance is beneficial as well given the functionality of the behavior.

The type of disconnect our model produces, where the macro level result (diversity) does not neatly line up with the micro level forces, is endemic to complex systems, an idea highlighted in Thomas Schelling's *Micromotives and Macrobehaviors*. Although individuals seek to be internally consistent and to conform, small errors accumulate in such a way to prevent them from accomplishing either. This paradox of aggregation suggests the value from considering how social psychological processes aggregate.

In the model that I have described, all behaviors are equally linked. More nuanced models might consider either clusters of behaviors based on context (work, family, social, etc.) or might place a network structure on the behaviors themselves. In other words, not all behaviors might share a link. One might even construct an influence parameter that captures how much one behavior is influenced by another. In such a model, one would expect clustered, linked, or close behaviors to converge to similar or identical values and those that are far apart not to do so. Such a model would be worth exploring and would almost certainly generate even greater intra-individual diversity.

Let me return now to Brewer's (1991) research on optimal distinctiveness. In trying to achieve distinctiveness, people might have to put forth great effort to differentiate themselves within communities that exert tremendous pressures to conform. Or people could find that it is relatively easy to be distinct, that even a tiny amount of experimentation maintains a level of personal heterogeneity necessary for a feeling of self-worth. Our findings suggest that achieving distinctiveness may be relatively easy. Coordinating may be the difficult part.

Ecologies of Mental Models

The coordination consistency model provides one approach for demonstrating how micro level interactions produce macro level diversity. An alternative approach borrows ideas and techniques from population genetics. The idea here is to think of an *ecology of mental models*. What do I mean by that? I mean that each person has a mental model that they carry around in their head. A person's mental model may change over time because of interactions with the mental models of others. As a modeling conceit, you can think of each type of mental model as you might a species in an ecology. Just as a savannah might contain lions and elephants, so might a cocktail party contain donkeys (democrats) and elephants (republicans).

¹⁵I have no evidence that it spreads from hats to coats. I am speculating.

The reason for adopting the ecological metaphor is that one can use tools from population genetics to solve for population level diversity as a function of the features of the world that matter to social scientists, namely, the structure of social interactions, the rates at which people copy from one another, the number of possible types, and the population size. Thus, it is at least possible that models from ecology may help us understand levels of cognitive or behavioral diversity in the social world just as they help ecologists to make sense of species diversity. This is what is meant by transdisciplinary: a single model applies to disparate disciplines.

Along with Evan Economo, Lu Hong and I, we have shown how an extension of the standard population genetics model can also produce a *diversity carrying capacity*, a total amount of maintainable diversity within a social systems (Economo, Hong, & Page, 2012). Imagine a population of individuals each of whom has a type. As indicated earlier, these types could be anything—behaviors, mental models, preferences, or problem-solving approaches. They could even correspond to how you wear your jacket on your motorbike. Assume that initially, these types are random. In subsequent periods, at some rate, individuals will update their types in one of two ways. They can copy the type of someone else. Or, they can randomly choose a new type. So, like in the previous model, there exists conformity and experimentation.

In population genetics models, the level of diversity in the population is considered as an aggregate statistic in a dynamical system. I will do my best here to make some rather cumbersome mathematics intuitive. Recall that to measure diversity, we need only to keep track of the probability that two randomly chosen individuals "match", that is, have the same type (the diversity index equals the inverse of the probability of matching). One uses a mathematical equation to keep track of how the probability of matching changes over time. At an *equilibrium*, the probability of matching—and therefore the diversity index—attains a specific value. People continue changing their types, but the probability that two randomly chosen people match types reaches an equilibrium. Thus, the equilibrium is *stochastic*, not *static*.¹⁶

The equation that captures how diversity change over time builds from two observations.

Observation 1: *If either of the two individuals experiments, the probability that they have the same type equals the probability of two randomly chosen types matching.*

Observation 2: *If both individuals copy someone from the population, they can either copy the same person in which case they will have the same type, or they can copy different people and the probability they have the same type will be the same as that for the entire population.*

Notice that using this framework, the number of types and the population size act as constraints. The relative propensities to copy and experiment describe how the social process unfolds. I will be the first to admit that the model that I have described so far is too stark to capture a real social process. It oversimplifies how ideas and behaviors spread and leaves out social structure entirely. This model is not the end of a research agenda. It is an opening. In Economo et al. (2012),

¹⁶See Economo et al. (2012) for the full version.

we extend the basic model that I just described to include social structure and show how social structure influences diversity.

Using a more general model, we can show that breaking a population into non-interacting subgroups does not produce a very large increase in diversity because each subgroup has a substantially lower diversity carrying capacity. Thus, whereas a group of 400 might maintain 40 types, two groups of 200 might each only be able to maintain 25 types, and their type overlap might be sufficient that overall diversity does not increase.¹⁷

The larger point, the one that I hope does not get lost in the minutia of these models, is that in a social system, multiple social psychological processes occur contemporaneously. These can include efforts to be individually consistent, to be unique or interesting, or to conform with a group. Collectively, these processes produce a distribution, a diversity, of behaviors, beliefs, ideas, and even skill sets. That diversity, as I show next, has functional value. Therefore, it is worth contemplating whether these process overproduce or underproduce diversity, a thought I return to in the discussion at the end of the essay.

THE FUNCTIONAL CONTRIBUTIONS OF DIVERSITY

The previous models have provided an overview of some techniques for understanding how much diversity we might expect to see in a population and how it depends on constraints and processes. In this section, I will briefly describe some models that explain why cognitive diversity plays such an important role, particularly in a society that confronts complex problems. I will do so in the context of two tasks: *prediction* and *problem solving*.¹⁸ I will then make some brief comments on the relationship between diversity and resilience.

Prediction

On a predictive task, collections of individuals often prove more accurate than individuals, a phenomenon popularized in the book *the Wisdom of Crowds* (Surowiecki, 2005). The relationship between diversity of predictions and collective accuracy can be stated using mathematics based on expected errors using what I call the *Diversity Prediction Theorem*.¹⁹ The theorem, a mathematical identity, can be written as follows:

$$\text{Crowd's Error} = \text{Mean Individual Error} - \text{Predictive Diversity}$$

In this equation, the crowd's error equals distance (the squared difference) between the crowd's prediction and the truth.²⁰ The mean error of individuals equals the average distance of each person in the crowd from the truth. The predictive diversity equals the average distance from the individual

predictions to the average prediction. Formally, this is the variance in the predictions, but I use the word diversity to highlight that this is the realized differences in their predictions.

The equation tells us that wise crowds, that is, crowds with small errors, rely as much on collective diversity as they depend on individual accuracy. Why? Consider a crowd of two people predicting annual snowfall (in inches) in Buffalo. Assume each has a squared error equal to nine, that is, each misses the true amount by 3 inches. If both predict 3 inches, too many or 3 inches too few, that is, if they exhibit no diversity, then the crowd's prediction equals their common individual prediction, and the crowd also has an error of nine. (The theorem can be written as $9 = 9 - 0$.) If instead, one predicts 3 in. too much and the other predicts 3 in. too few, then the crowd will be exactly correct. The errors cancel because of the diversity of the predictions. (The theorem can be written as $0 = 9 - 9$).

The point of introducing this theorem is to lay a foundation for deeper questions about the nature of the predictions themselves. I have been particularly interested in how to think about the statistical properties of the predictions. In the formal literature, there exist two distinct strands. Lu Hong and I refer to these as the *generated* and *interpreted* frameworks (Hong & Page, 2009). In the generated framework, predictions are modeled as the truth plus a disturbance. Imagine that people are predicting the height of a tree. Each might have a slightly different vantage point. Each can then be thought of as seeing the true height plus some error term. If those error terms are independent, then they will be diverse, and the crowd error will be small. In this framework, diversity comes from error or noise from an external source. In this framework, crowds do better because the individual random errors cancel.

The incredible success of statisticians in predicting recent electoral outcomes relies on the generated framework logic. Each poll represents a signal of the truth plus some error. By averaging across many polls, those errors cancel. Of course, if the statistician possesses more information about the accuracy and the biases of those polls, he or she can use weighting procedures in which more accurate and more diverse predictions get greater weight that perform even better (Lamberson & Page, 2011; Silver, 2012).

This notion of predictions as the truth plus an error term makes for a tractable model but has limited empirical plausibility in many contexts and equates diversity with noise. It does not connect to the previous models in which a population of people have diverse ideas or models that they carry around in their head. Fortunately, this way of thinking about predictions—as coming from models in our heads, can be captured by the interpretive signal framework.

In the interpretive framework, people use different models and therefore make different mistakes. Psychologists differ in how they conceptualize these mental models. Johnson-Laird (2009) relied on a truth table framework. My own attempts to construct models of predictive models borrow more from Smith and Medin (1981) who think in terms of predictions based on categories. In this approach, an individual assigns events to categories and then constructs causal or correlative frameworks between those categories. Thus, people who have had different life experiences, come from different cultural groups, or have different formal training may well construct different categories and rely on different models to make their

¹⁷One can explore these same sorts of questions empirically by looking at the heuristics that individuals use to choose their type. This requires good data and the ability to characterize heuristics formally. For one remarkable application in the context of scientific research, see Evans and Foster (2011) who analyzed the heuristics that scientists use to decide on which molecules to study.

¹⁸See Page (2007) for a more elaborate presentation of these models.

¹⁹See Page (2007), Armstrong (2001), and Batchelor and Dua (1995) for more complete treatments.

²⁰Statisticians square errors to punish larger errors more and because if you summed signed errors negative and positive errors would cancel one another out.

predictions. If one takes the Johnson-Laird approach, people could construct different truth tables based on different experiences or influences. In either case, one can think of the diverse models as being produced at least partly by the kinds of social processes that I described in the previous section.

The diversity in categorizations that underpins interpretative signals has been the focus of substantial research in several fields. Social psychologists have identified differences in how people categorize the world, including how we categorize others (Linville, Fischer, & Salovey, 1989). Cognitive anthropologists have developed a number of techniques to characterize the homogeneity, or what they call the *consensus*, within a culture (Romney, Weller, & Batchelder, 1986). These include pile sorts, in which individuals are asked to place objects or events in piles based on similarity. These categorization tasks test to see if individuals make sense of the world in the same way. In a parallel research agenda, brain scientists have been studying how people classify entities by their features (see Shafto, Kemp, Mansinghka, Gordon, & Tenenbaum, 2006).

Notice that the logic for how interpretative signals aggregate provides more nuance. If people use different categories and therefore construct different models, they can collectively make accurate predictions not because their errors cancel—the logic from the generated framework—but because they are filtering the world differently and therefore each seeing different approximations.

Holding accuracy fixed, more diverse predictions produce more accurate crowds (Ashton, 1986). True, as in the generated case, the errors are canceling out, but the reason they cancel out is not because the predictions are statistically independent draws from some external source. They are different because the people rely on distinct models.²¹

Let us now link the logic of the Diversity Prediction Theorem using the interpretive framework to the social process model. It is not difficult to rewrite a version of the diversity prediction theorem using the diversity index (Economo et al., 2012). Thus, the diversity index together with average individual accuracy can be a proxy for crowd accuracy. A social process that produces more diversity in how people think will result in more accurate collective prediction—provided of course that those additional models are not too much less accurate. What this tells us is that we would like for social processes to drive out the bad models—as their reduction in accuracy probably outweighs whatever advantages they provide in terms of diversity—but to keep the models that are diverse but relatively high accuracy.

The value of these diverse models becomes more important as the world becomes more complex, that is, harder to predict. The logic goes as follows: we would like the crowd error to be small. However, the environment is complex, that is, hard to predict. Any individual model is likely to produce a large error. Now, invoke the Diversity Prediction Theorem. It says, *crowd error equals average error minus diversity*. If average error is large (because of the complexity of the world), the only way that the crowd error can be small is if the diversity is large. *Therefore, diversity is necessary in the face of complexity.*

²¹Of course, for both interpreted and generated signal structures, one can do better than simple averaging—a subject of much attention within the ensemble learning community (Breiman, 1996)

This intuition can be made more formal by thinking in terms of the number of relevant variables and interdependencies. Suppose that we are predicting outcomes in a system that has only eight relevant variables. This might be a task like predicting the number of pounds of bananas or boxes of crackers sold during a given month by a large grocery store. Assuming sufficient data, one can probably construct a model that is highly accurate. There may not be much room for improvement on the best models and, as a result, at most a small benefit from diversity.

Next, suppose that we are predicting the result of a much more complex process—perhaps the weather next weekend or the price of a stock 2 years from now. This may be difficult for any one person or model to get right because the system contains many interacting parts that can be categorized in numerous ways. In other words, there will exist many possible models. In such cases, aggregating diverse models will be more likely to add substantial value. And in fact, in weather forecasting, human forecasters, although less accurate, can improve on computer predictions owing to humans looking at the data differently (Silver, 2012). A similar phenomenon occurred in the Netflix Prize competition in which teams competed for a million dollars to predict movie preferences. Here as well, aggregates of models far outperformed the best individual model (Bell & Koren, 2007).

Even with diverse models, we have no guarantee of being able to predict the future, especially not in the modern complex world Tetlock (2005). Complex systems, by definition, are difficult to predict. This does not mean that they are random. Whatever did happen will be logically consistent, which will compel us to seek better theories and models (Watts, 2011). Complex systems, because they contain interdependent parts, can also produce large events (Taleb, 2007). Often, these large events are hard to predict. Almost no one predicted the dismantling of the Soviet Bloc (Kalyvas, 1999), the Arab spring (Lamberson & Page, 2012), or the financial market collapse. The latter case may be partly delusional; no one wants the bubble to be a bubble (Reinhart & Rogoff, 2009). Nevertheless, the bubble did exist, and we all suffered. Greater diversity in the ideas that people carry around in their heads would not necessarily reduce the number and extent of such events, but it should lead to less collective surprise and therefore better preparation for the consequences.

Problem Solving

I next turn to problem solving and the role that diversity of cognitive tools plays in finding better solutions to problems. A substantial part of the modern economy consists of people who solve problems for a living (Autor, Katz, & Kearney, 2006; Wolff, 2006). In problem solving, the value added of an individual is contextual. It depends on the tools the person brings as well as the tools of the other people working on the problem.

The tool framework problematizes the notion of individual ability. To understand why the concept of ability may have less purchase when people solve problems rather than drive spikes, consider an economy in which people engage primarily in physical labor. The ability of an individual might represent how many table legs he or she can make in a day. Assuming no crowding effects, the total number of table legs made by

10 people equals the sum of their individual production levels. In that world, production is additive. The best team consists of the best individuals. The same would be true if people were laying rails or applying rivets. In each of these cases, if I wanted to build the best possible team of 10 members, I would hire the 10 individuals of highest ability—the ones who could produce the most table legs.

This additive logic does not hold for problem solving, which is why diversity becomes important. Suppose that I use IQ as proxy for ability and I hire people based solely on IQ. Generally, that would not give the best team because team problem-solving ability is not additive in the abilities of the team members (Hong & Page, 2001). Evidence using IQ test scores from British schoolchildren on the Raven's Progressive Matrix test shows that someone's contribution to group may depend on the others in the group. In other words, given two people with the same IQ, one may be more likely to get problems right than other people in the group tend to get wrong (Bachrach, Graepel, Kasneci, Kosinski, & Van Gael, 2012). Another study shows that a team may have something like an IQ and that a team's IQ does not correlate strongly with the IQs of its members (Woolley, Chabris, Pentland, Hashmi, & Malone, 2010).

How then does one think about team performance on problem solving tasks? Lu Hong and I (Hong & Page, 2001, 2004) constructed a model based on what we call *perspectives* and *heuristics*. A perspective is a representation of a problem, a mapping of the set of all possible solutions. Suppose that individuals are tasked with finding a minimal route that visits 10 cities, a traveling salesperson problem. One perspective would be to write the cities in a list according to the order that they will be visited. Another perspective would be to list the cities alphabetically but assign a number to each one depending on where it is in the order. A third perspective would be to draw the cities on a map with the route represented by lines.

Heuristics are the tricks, tools, and techniques that an individual uses to locate new solutions within his or her perspective. Heuristics for the first perspective would include flipping adjacent cities, flipping random cities, and flipping cities separated by a single city.

In our model, individuals can differ in both their perspectives and their heuristics. In this framework, a person's ability equals how well he or she performs on the problem on his or her own. The ability of a group of people equals how well they do collectively. As mentioned earlier, ability is not additive. The ability of the group does not equal the sum of the abilities of the individuals. More germane to the present argument, it is also true, under some rather mild assumptions, that diversity often trumps ability, by that I mean that choosing team members based on diversity of perspectives and heuristics will lead to better outcomes than choosing team members based on ability (Hong & Page, 2004).

Diversity trumps ability because the best problem solvers often tend to be similar: if you choose the two best problem solvers from a large set, then they are likely to use similar perspectives and heuristics. As a result, their two heads will not be much better than one. In contrast, if two people rely on diverse perspectives, then a solution that appears natural to one problem solver may be innovative to another. I should emphasize that these results are theoretical, not experimental or empirical.

That said, empirical analyses covering tens of thousands of cases from disparate domains align with the intuition produced

by the model. Studies of online problem-solving sites involving hundreds of thousands of problem solvers show that writing questions to attract more diverse problem solvers leads to higher probabilities of correct solutions (Jeppesen & Lakhani, 2010; Lakhani & Jeppesen, 2007). Studies of high impact papers in the academy show the value of citing diverse sources (Schilling & Green, 2011) and even the value of having co-authors from different institutions (Jones, Wuchty, & Uzzi, 2008). Studies of patents show that the proximity of researchers results in fewer citations (Shi, Adamic, Tseng, & Gavin, 2009). Finally, as mentioned in the introduction, individuals in larger cities have long been known to be more productive. Many explain this as partly being due to exposure to more ideas, to cities maintaining more diversity (Glaeser, 2010).

These findings contrast with the literally hundreds of studies of small teams and groups that show little effect of diversity on average that I mentioned in the introduction. Further, meta-analyses of the experimental studies demonstrate that one cannot separate the sorts of functional diversity from other types of diversity such as identity diversity and personality type diversity (see van Knippenberg & Schippers, 2007). This conflation of concepts muddles empirical tests of any one effect. Also, even though there exist many studies, the sample size within each one is small. By way of comparison, the aforementioned research on patents and academic papers covers tens of thousands of cases, and the research on cities implicitly includes millions of people. Overall then, the data supporting the value of cognitive diversity are, by the usual standards of evidence, quite strong for small groups and exceptionally strong for the large populations of scientists.

Resilience

Last, I will briefly describe the theoretical link between diversity and resilience. Among the first to make this connections was Ashby, who helped to develop the field of cybernetics. Ashby's *Law of Requisite Variety* (Ashby, 1956) states that for every disturbance, there must exist a response. Thus, if a society consists of people with diverse sets of skills, they are more likely to withstand some external disturbance. The logic is rather obvious: more tools imply more responses. Ashby's logic can be extended by allowing for recombination of responses. If responses can be combined say by attaching a saw to an electric motor or a chain to a bicycle, then the number of types of responses explodes combinatorially. This of course ignores any cost of maintaining that diversity.

The second, and more subtle link, between diversity and resilience relies on the Price Equation and what biologists call *Fisher's Fundamental Law* (Frank, 1997). This states that the rate of adaptation correlates with the degree of diversity. The mathematics of the Price Equation is a bit involved (see Page, 2010), but the logic is easy to grasp. Represent a problem as a mountain with a single peak, something like Mt. Fuji. Now, imagine that the value or utility of a solution equals the height of its location on the mountain, so the best solution is at the peak.

Imagine first that there exists no diversity, and that everyone has found the optimal solution: everyone is standing on the peak. Suppose that there is a shifting of the plates of the earth and what was a peak is now the side of a larger mountain. In a biological system, if there was no diversity, there would

be no ability to evolve. If, though, diversity had been maintained, so that some members of the population were not at the peak, then after the plate shift, some of them would now be on higher ground. Their DNA could lead the way to the higher peak. This is the logic of the Price Equation and Fisher's Fundamental Theorem: diversity is a form of exploration.

Exploration produces resilience to change. Here again, we see an important difference between complex systems and equilibrium systems. In equilibrium system, you should stay at the peak. Diversity is variance. In a complex system that constantly shifts and adapts, diversity enables the population to move from poor locations to better ones. Hence, more diversity implies more robustness.

DISCUSSION

In this article, I have described how theoretical analyses of collective predictions, problem solving, and resilience to change all demonstrate the value of cognitive diversity and how models of dynamic social processes provide some ideas about how that diversity might arise and be maintained. In particular, models with multiple processes seem capable of maintaining perhaps even more diversity. These findings suggest a possible value in considering multiple social psychological processes simultaneously.

The models I described can provide insights into how processes can maintain diversity within and across populations, but such models cannot tell us about the relative contributions to cognitive difference that result from identity, culture, experience, education, etc. Some differences or predispositions may be partly innate. Many may be due to context and social psychological influences. And probably, all are influenced in ways big and small by the interactions between who we are and who we are with. These questions can be seen as reformulations of the age old agency versus structure discussion. But that does not make them any less important.

In that my interests have been primarily whether societies produce too little or too much cognitive diversity (Hong et al., 2012), I have been drawn toward research in social psychology. Several of the social and individual level influences that lie at the core of the discipline, conformity, coordination, and dissonance reduction, all work to reduce diversity, as does emulating those who appear to be performing best.

Does this damping down lead to too little diversity? Being trained as an economist I cannot help but ask such questions. The optimal level, if such a thing exists, surely depends on context. We do not necessarily want a lot of diversity when people choose which side of the road on which to drive. Nor do we want diversity when people work as part of a coordinated team on a routine task. This logic underpins the enormous success of scientific management techniques, six sigma, and checklist procedures (Gawande, 2009; Tennant, 2001). These all work to reduce diversity—to do a job the same way, the right way, every time. And within a manufacturing process, that may be fine. But within a political or social system, diversity reduction to that extent would be dangerous (Le Bon, 1895).

When we look out at the world, we do not see pure conformity because individuals recognize and celebrate the value of

thinking differently. We take delight and even pride in being unique. These interests push back against the force of conformity (Becker et al., 2012). The balance between the desire to be different and the appeal of conforming produces the mangle of difference and sameness that we see in the world. But again, this in no way guarantees that the level of diversity produced necessarily aligns with what is desired.

To my knowledge, there exists no invisible hand argument that diversity necessarily equilibrates as some ideal level. One might argue that in complex environments, we need more diversity. Hence, business strategists often invoke the phrase “innovate or die” or advocate “quick and nimble” corporate cultures to borrow the title of a recent book (Bryant, 2014). Many social observers also feel that the collapse of distinct cultures may lead to less robustness. Ecologists have similar concerns—thus the transdisciplinary nature of many of the questions that link diversity to complexity.

That is not to say that more diversity is necessarily better. Perhaps, we need less diversity. Is not diversity one cause of complexity? Certainly, it is. And this suggests that there may well be a sweet spot, or at least a preferred region. And given the substantial functional contributions from cognitive diversity, questions of how much and what types of diversity would create a more robust, innovative, and fair society merit deeper thinking, and especially thinking by social psychologists. Why do people choose to learn what they learn? Why do people think about problems in particular ways? What social psychological processes drive the creation and spread of diverse perspectives, categories, analogies, mental models, and heuristics? All questions worthy of further study.

ACKNOWLEDGEMENTS

I would like to thank the editors for this fantastic opportunity and for yeoman editorial effort on this paper. In an earlier version, I was culling deer from sandwich eating teams (undoing subgroups). I owe enormous debts of gratitude to Evan Economo, Jenna Bednar, and Lu Hong for their contributions to the ideas contained within this paper. Andrea Jones-Rooy, Rich Gonzalez, and Craig McGarty gave extensive comments on earlier drafts. A portion of the research described in this paper was funded by Army Office of Research ARO-MURI grant W911NF1010379.

REFERENCES

- Adamic, L. A., Wei, X., Yang, J., Gerrish, S., Nam, K. K., & Clarkson, G. S. (2010). Individual focus and knowledge contribution. *First Monday*, 15(3). Available: <http://firstmonday.org/article/view/1057/977>
- Ajzen, I., & Fishbein, M. (1980). *Understanding attitudes and predicting social behavior*. Englewood Cliffs, NJ: Prentice-Hall.
- Alesina, A., & La Ferrara, E. (2005). Ethnic diversity and economic performance. *Journal of Economic Literature*, XLIII, 762–800.
- Alesina, A., Harnoss, J., & Rapoport, H. (2013). Birthplace diversity and economic prosperity. NBER working paper 18699.
- Aligica, P. D. (2014). *Institutional diversity and political economy: The Ostroms and beyond*. Oxford: University Press.
- Anderson, P. W. (1972). More is different. *Science*, 77(4047), 393–396.
- Armstrong, J. S. (2001). *Combining forecasts. Principles of forecasting: A handbook for researchers and practitioners*. Norwell, MA: Kluwer Academic Publishers.

- Asch, S. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership, and men* (pp. 177–190). Pittsburgh, PA: Carnegie Press.
- Ashby, W. R. (1956). *Introduction to cybernetics*. London: Chapman and Hall.
- Ashraf, Q., & Galor, O. (2013). The out of Africa hypothesis, human genetic diversity, and comparative economic development. *The American Economic Review*, 103, 1–46.
- Ashton, R. H. (1986). Combining the judgments of experts: How many and which ones? *Organizational Behavior and Human Decision Processes*, 38(3), 405–414.
- Autor, D. H., Katz, L. F., & Kearney, M. S. (2006). The polarization of the U. S. labor market. NBER Working Paper No. 11986.
- Axelrod, R., & Cohen, M. (2001). *Harnessing complexity: Organizational implications of a scientific frontier*. New York: Basic Books.
- Bachrach, Y., Graepel, T., Kasneci, G., Kosinski, M., & Van Gael, J. (2012). Crowd IQ—Aggregating opinions to boost performance. AAMAS 2012.
- Batchelor, R., & Dua, P. (1995). Forecaster diversity and the benefits of combining forecasts. *Management Science*, 41(1), 68–75.
- Becker, M., Vignoles, V. L., Owe, E., Brown, R., Smith, P. B., Easterbrook, M., & Yamakoglu, N. (2012). Culture and the distinctiveness motive: Constructing identity in individualistic and collectivistic contexts. *Journal of Personality and Social Psychology*, 102(4), 833.
- Bednar, J. (2009). *The robust federation*. Cambridge UK: Cambridge University Press.
- Bednar, J., & Page, S. E. (2007). Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. *Rationality and Society*, 19(1), 65–97.
- Bednar, J., Bramson, A., Jones-Rooy, A., & Page, S. E. (2010). Emergent cultural signatures and persistent diversity: A model of conformity and consistency. *Rationality and Society*, 22(4), 407–444.
- Bednar, J., Chen, Y., Liu, T. X., & Page, S. E. (2012). Behavioral spillovers and cognitive load in multiple games: An experimental study. *Games and Economic Behavior*, 74(1), 12–31.
- Bell, R., & Koren, Y. (2007). Lessons from the Netflix prize challenge. *ACM SIGKDD Explorations Newsletter—Special issue on visual analytics archive*, 9(2), 75–79.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2), 123–140.
- Brewer, M. B. (1991). The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin*, 17, 475–482.
- Brown, D. G., Page, S. E., Riolo, R. L., Zellner, M., & Rand, W. M. (2005). Path dependence and the validation of agent-based spatial models of land-use. *International Journal of Geographical Information Science*, 19(2), 153–174.
- Bryant, A. (2014). *Quick and nimble: Lessons from leading CEOs on how to create a culture of innovation*. New York, NY: Times Books.
- Cason, T., Savikhin, A., & Sheremeta, R. (2012). Behavioral spillovers in coordination games. *European Economic Review*, 56, 233–245.
- Darmofal, D. (2005). Elite cues and citizen disagreement with expert opinion. *Political Research Quarterly*, 58(3), 381–395.
- Dell, M. (2010). The persistent effects of Peru's mining mita. *Econometrica*, 78(6), 1863–1903.
- Economio, E., Hong, L., & Page, S. E. (2012). *Cognitive ecologies, social structure, and collective wisdom (mimeo)*. University of Michigan.
- Evans J. A., & Foster, J. G. (2011). Metaknowledge. *Science*, 332, 721–725.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817–868.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Festinger, L., & Carlsmith, M. (1959). Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology*, 58, 203–210.
- Florida, R. (2002). *The rise of the creative class: And how it's transforming work, leisure, community and everyday life*. New York: Basic Books.
- Frank, S. (1997). The Price equation, Fisher's fundamental theorem, kin selection, and causal analysis. *Evolution*, 51(6), 1712–1729.
- Gawande, A. (2009). *The checklist manifesto: How to get things right*. NY: Metropolitan Books.
- Gigerenzer, G. (2000). *Simple heuristics that make us smart*. USA: Oxford University Press.
- Glaeser, E. (2010). *Triumph of the city: How our greatest invention makes us richer, smarter, greener, healthier, and happier*. New York: Penguin.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., & McElreath, R. (2001). In search of homo economicus: Behavioral experiments in 15 small-scale societies. *American Economic Review*, 91, 73–78.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world. *Behavioral and Brain Sciences*, 33(2-3), 61–83.
- Hong, L., & Page, S. E. (2001). Problem solving by heterogeneous agents. *Journal of Economic Theory*, 97, 123–163.
- Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46), 16385–16389.
- Hong, L., & Page, S. E. (2009). Interpreted and generated signals. *Journal of Economic Theory*, 144, 2174–2196.
- Hong, L., Riolo, M., & Page, S. E. (2012). Incentives, information, and emergent collective accuracy. *Managerial and Decision Economics*, 33(5), 323–334.
- Huckfeldt, R. R., Johnson, P. E., & Sprague, J. D. (2004). *Political disagreement: The survival of diverse opinions within communication networks*. Cambridge, England: Cambridge University Press.
- Jeppesen, L. B., & Lakhani, K. R. (2010). Marginality and problem-solving effectiveness in broadcast search. *Organization Science*, 21, 1016–1033.
- Johnson, J., & Knight, J. (2011). *The priority of democracy*. Princeton, NJ: Princeton University Press.
- Johnson-Laird, P. N. (2009). *How we reason*. USA: Oxford University Press.
- Jones B., Wuchty, S., & Uzzi, B. (2008). Multi-university research teams: Shifting impact, geography, and stratification in science. *Science*, 322, 1259.
- Kalyvas, S. (1999). The decay and breakdown of communist one-party systems. *Annual Review of Political Science*, 2, 323–343.
- Kitayama, S., Snibbe, A., Markus, H., & Suzuki, T. (2004). Is there any free choice? Self and dissonance in two cultures. *Psychological Science*, 15, 527.
- van Knippenberg, D., & Schippers, M. C. (2007). Work group diversity. *Annual Review of Psychology*, 58, 515–541.
- Laitin, D., & Jeon, S. (forthcoming). Opportunities in cultural diversity. In *Emerging trends in the social and behavioral sciences*. New York: Sage Publications.
- Lakhani, K. R., & Jeppesen, L. B. (2007). Getting unusual suspects to solve R&D puzzles. Forethought. *Harvard Business Review*, 85(5), 30–32.
- Lamberson, P. J., & Page, S. E. (2011). Optimal forecasting groups. *Management Science*, 58, 791–804.
- Lamberson, P. J., & Page, S. E. (2012). Tipping points. *Quarterly Journal of Political Science*, 7(2), 175–208.
- Landemore, H. (2013). *Democratic reason*. Princeton, NJ: Princeton University Press.
- Le Bon, G. (1895). *The crowd: A study of the popular mind*. Southampton: Sparkling books edition.
- Linville, P., Fischer, G., & Salovey, P. (1989). Perceived distributions of the characteristics of in-group and out-group members: Empirical evidence and a computer simulation. *Journal of Personality and Social Psychology*, 57(2), 165–188.
- Medin, D., Bennis, W., & Chandler, M. (2010). The home-field disadvantage. *Perspectives on Psychological Science*, 5(6), 708–713.
- Miller, J., & Page, S. E. (2007). *Complex adaptive systems: An introduction to computational models of social life*. Princeton, NJ: Princeton University Press.
- Mitchell, M. (2009). *Complexity: A guided tour*. Oxford: University Press.
- Myerson, R. (1997). *Game theory: Analysis of conflict*. Cambridge, MA: Harvard University Press.
- Nowak, A., & Vallacher, R. R. (1998). *Dynamical social psychology*. New York: Guilford Press.
- Ostrom, E. (2005). *Understanding institutional diversity*. Princeton NJ: Princeton University Press.
- Page, S. E. (2006). Essay: Path dependence. *Quarterly Journal of Political Science*, 1, 87–115.
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton, NJ: Princeton University Press.
- Page S. E. (2008a). *Understanding complexity (video)*. Chantilly, VA: The Great Courses.
- Page, S. E. (2008b). Uncertainty, difficulty, and complexity. *Journal of Theoretical Politics*, 20, 115–149.
- Page, S. E. (2010). *Diversity and complexity*. Princeton, NJ: Princeton University Press.
- Page, S. E., Sander, L., & Schneider-Mizell, C. (2007). Conformity and dissonance in generalized voter models. *Journal of Statistical Physics*, 128, 1279–1297.
- Pelled, L. H., Eisenhardt, K. M., & Xin, K. R. (1999). Exploring the black box: An analysis of work group diversity, conflict, and performance. *Administrative Science Quarterly*, 44, 1–28.
- Putnam, R. (2007). E pluribus unum: Diversity and community in the twenty-first century: The 2006 Johan Skytte prize lecture. *Scandinavian Political Studies*, 30(2), 137–174.
- Reinhart, C., & Rogoff, K. (2009). *This time is different: Eight centuries of financial folly*. Princeton, NJ: Princeton University Press.
- Romney, A. K., Weller, S. C., & Batchelder, W. H. (1986). Culture as consensus: A theory of culture and informant accuracy. *American Anthropologist*, 88(2), 313–338.
- Ross, L., & Ward, A. (1995). Psychological barriers to dispute resolution. In M. Zanna (Ed.), *Advances in experimental social psychology* (vol. 27, pp. 255–304). San Diego: Academic Press.
- Ross, L., Bierbrauer, G., & Hoffman, S. (1976). The role of attribution processes in conformity and dissent. *American Psychologist*, 31, 148–157.

- Samuelson, L. (2001). Analogies, adaptations, and anomalies. *Journal of Economic Theory*, 97, 320–367.
- Schelling, T. C. (1978). *Micromotives and macrobehavior*. New York, NY: W.W. Norton.
- Schilling, M. A., & Green, E. (2011). Recombinant search and breakthrough idea generation: An analysis of high impact papers in the social sciences. *Research Policy*, 40(10), 1321–1331.
- Shafto, P., Kemp, C., Mansinghka, V., Gordon, M., & Tenenbaum, J. B. (2006). Learning cross-cutting systems of categories. *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society*.
- Sherif, M., & Harvey, O. J. (1952). A study in ego-functioning: Elimination of stable anchorages in individual and group situations. *Sociometry*, 15, 272–305.
- Shi, X., Adamic, L., Tseng, B., & Gavin, C. (2009). The impact of boundary spanning scholarly publications and patents. *PLoS ONE*, 4(8), e6547.
- Silver, N. (2012). *The signal and the noise: Why so many predictions fail but some don't*. Princeton, New York: Penguin Press.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge MA: Harvard University Press.
- Surowiecki, J. (2005). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. New York, NY: Little, Brown.
- Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. New York: Random House.
- Tennant, G. (2001). *Six sigma: SPC and TQM in manufacturing and services*. Surrey, UK: Gower Publishing.
- Tetlock, P. (2005). *Expert political judgment: How good is it? How can we know?* Princeton, New Jersey: Princeton University Press.
- Tilman, D. (1999). The ecological consequences of changes in biodiversity: A search for general principles. *Ecology*, 80, 1455–1474.
- Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342, 468–472.
- Van Overwalle, F., & Jordens, K. (2002). An adaptive connectionist model of cognitive dissonance. *Personality and Social Psychology Review*, 6(3), 204–231.
- Watts, D. (2011). *Everything is obvious: Once you know the answer*. New York, NY: Crown Business Press.
- Wolff, E. N. (2006). The growth of information workers in the us economy, 1950–2000: The role of technological change, computerization, and structural change. *Economic Systems Research*, 18(3), 221–255.
- Wolfram, S. (2002). *A new kind of science*. Champaign, IL: Wolfram Media.
- Woolley, A., Chabris, C., Pentland, A., Hashmi, N., & Malone, T. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–688.