E X C H A N G E

# Making the Difference: Applying a Logic of Diversity

by Scott E. Page


## Executive Overview

Each year, corporations spend billions of dollars on diversity training, education, and outreach. In this article, I explain why these efforts make good business sense and why organizations with diverse employees often perform best. I do this by describing a logic of diversity that relies on simple frameworks. Within these frameworks, I demonstrate how collections of individuals with diverse tools can outperform collections of high "ability" individuals at problem solving and predictive tasks. In problem solving, these benefits come not through portfolio effects but from *superadditivity:* Combinations of tools can be more powerful than the tools themselves. In predictive tasks, diversity in predictive models reduces collective error. It's a mathematical fact that diversity matters just as much as highly accurate models when making collective predictions. This logic of diversity provides a foundation on which to construct practices that leverage differences to improve performance.


Along the moving sidewalks inside Paris' Charles de Gaulle airport, you cannot help but notice a sequence of HSBC advertisements meant to show diverse perspectives. One shows two identical pictures of a half-full glass of water. Across one glass, the caption reads *moitié vide*, under the other *moitié plein*. A second advertisement shows two identical pictures of an apple with a bite taken out. *Défendu* scrolls across one apple and *recommandé* across the other. These ads encourage us to think of HSBC as a firm that sees a problem from more than one perspective—and they also provide a welcome diversion from the inefficiencies of the airport. This multiple perspective taking allows HSBC to add value, or so we are intended to believe.

The HSBC ads reflect a broader trend. Each year, corporations spend billions related to promoting positive messages about diversity both internally and externally. Why profit-seeking businesses commit so many resources to constructing diverse workforces and creating welcoming organizational cultures stems from two trends. First, businesses have become more global and hence more ethnically diverse. Firms sell to diverse consumers and hire from a diverse pool of candidates. The world, as has been said, is now flat, and consequently, organizations must cope with diversity. Second, the practice of work has become more team focused. The fixed hierarchy has given way to the evolving matrix (Mannix & Neale, 2006). In the past, welders positioned two stations apart on an assembly line need not get along. They need not validate one another's worldview. The same cannot be said of a team of scriptwriters or oncologists, who must learn to understand the language and actions of one another.

This article is adapted from Scott E. Page's book *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, published in 2007 by Princeton University Press.

***Scott E. Page*** *(spage@umich.edu) is Professor of Complex Systems, Political Science, and Economics at the University of Michigan-Ann Arbor. He is also an external faculty member of the Santa Fe Institute.*

Copyright by the Academy of Management; all rights reserved. Contents may not be copied, e-mailed, posted to a listserv, or otherwise transmitted without the copyright holder's express written permission. Users may print, download, or e-mail articles for individual use only.

This coincident emergence of diverse workforces and team-based work makes leveraging diversity a central concern of most organizations. A first question to ask is whether it's a good thing from a business perspective. Does it hurt or help the bottom line? A substantial empirical literature addresses the question of whether diversity improves team performance (Williams & O'Reilly, 1998). A brief summary of that literature reveals that the answer depends on several factors. Particularly important is *what people believe* (Ely & Thomas, 2001). If people do not believe in the value of diversity, then when part of a diverse team they're not as likely to produce good outcomes. That expectations shape behavior and that behavior shapes outcomes should not come as a big surprise. How though to change expectations? How does an organization get its employees to believe that diversity leads to better outcomes? Taking out advertisements or printing up human resources documents with elaborate graphics and catchy tag lines won't make it so. Managers and employees need, to quote Springsteen, "a reason to believe."

Simple, clean logic can provide that reason. In this brief article, I derive links between cognitive differences among team members and better collective outcomes at specific tasks: *problem solving* and *prediction*. I build those links using conceptual frameworks that borrow from psychology, computer science, and economics. These links not only provide a foundation for understanding when, how, and if diversity produces benefits—the reason to believe—they also point toward specific policies and practices that can leverage the power of diversity.

The bottom line: Diversity can improve the bottom line. It may even matter as much as ability.

## Diverse Perspectives and Heuristics

I begin by formalizing the loose notion of a *perspective*. No end of brochures and advertisements sing the praises of diverse perspectives, but what are they? Here, I define a perspective to be a representation of the set of the possible: the set of the semiconductor designs, welfare policies, or fall leather coats. Two people possess diverse perspectives if they mentally represent the "set of the possible" differently. For example, one person might organize a collection of books by their authors' last names; another person might organize those same books by color and size. One professor might arrange students' names by class rank; another professor might order those same students alphabetically.

How a person represents the set of the possible determines "what is next to what." For example, *The Catcher in the Rye* may seem rather disconnected from Mao's *Little Red Book*, but they are adjacent in a perspective that organizes books by color and size. Perspectives matter because "what is next to what" determines how a person locates new solutions. The linkage between perspectives and locating solutions can be clarified with an example. Suppose you are making butternut squash soup. You've pureed the sautéed onion and added the cream and baked squash, but the result tastes bland. Arrayed before you is an enormous spice rack. You're thinking that perhaps you'll add cumin. You sniff the cumin. It smells fine, but next to it, you see curry. So you smell the curry and decide it will be wonderful. You only try the curry because it sits adjacent to the cumin. Had the spices been arranged differently, say by color, you might have added cinnamon instead. What is next to what—in this case curry is next to cumin—determines where you look.

This same logic extends to almost any problem-solving situation: *Two people with different perspectives test different potential improvements and increase the probability of an innovation*.

Diverse perspectives may be the cause of most breakthroughs, but this does not mean that all diverse perspectives prove helpful. Someone who sees a problem from a different perspective will notice different *candidate* solutions. But those candidate solutions need not improve the status quo. Diverse perspectives prove most valuable if they embed information relevant to the problem being solved. For example, in trying to increase fuel efficiency, a perspective that focuses on the weight of parts will likely yield good ideas. A perspective that considers their color probably won't. Therefore, while organizations should encourage bringing diverse perspectives to a prob-

lem, they must also have some method for identifying useful perspectives.

Perspectives describe how people see a problem, but they do not fully capture the act of problem solving. When solving problems, people also use heuristics. Heuristics are methods or tools to find solutions. In my description of searching for a spice to add to the soup, I've assumed that you looked at adjacent jars. This is an example of a *heuristic*. Heuristics take many forms and vary in their sophistication from simple rules of thumb to complicated algorithms. To give a flavor for how heuristics operate, I describe here a famous simple heuristic known as *do the opposite*. In a classic episode of the television show *Seinfeld*, Jerry's bumbling friend George Costanza comes to the realization that every decision he has made in his life has been the wrong one. This realization results in an epiphany: He should *do the opposite*. He should do the reverse of whatever he thinks is best. If the rules in his head tell him to be kind, he should be rude. If they tell him to arrive early, he should show up late. If they tell him to dress casually, he should dress formally. The irony, of course, is that doing the opposite of what he thinks is right is the only "right" thing George has ever done, and by the end of the show he achieves personal and professional success. Diverse heuristics, like diverse perspectives, improve problem solving, but they do so in a different way.

Whereas perspectives change "what is next to what," heuristics change how a person searches for solutions. Imagine two engineers working for a manufacturing company trying to improve the speed of an assembly line. The first engineer's heuristic might be to try to break down individual tasks into smaller tasks. The second engineer's heuristic may be to switch the order of the tasks. The two heuristics differ, and because they differ, they identify different candidate solutions, increasing the probability of a breakthrough.

This brief description of diverse perspectives and heuristics and how they operate reveals only part of the power of diversity. What I've shown is that by seeing problems differently (diverse perspectives) and by looking for solutions in different ways (diverse heuristics), teams, groups, and organizations can locate more potential innovations. I now show that these individual improvements can be combined, creating *superadditivity*. Superadditivity exists when the total exceeds the sum of the parts, when $1+1 = 3$.

The idea that $1+1 = 3$ may seem counterintuitive. Yet, when we add heuristics, either the two heuristics are the same (i.e., each points to the same solution, and therefore $1+1 = 1$) or the two heuristics differ (in which case $1+1 = 3$). Why three? Let's do the math. Let's go back to our assembly line problem. The first heuristic might advocate dividing a task that consists of six spot welds into two tasks. The second heuristic might advocate gluing on a piece of trim prior to the welding. The third heuristic comes from doing both—dividing the task *and* switching the order. Thus, any time you have two heuristics, you can create a third by combining the two heuristics. A similar logic shows that $1+1+1 = 7$. Far from being a meaningless buzzword, superadditivity can be real, but only if people bring diverse perspectives and solutions to a problem.

The logic that diversity creates superadditive benefits differs from the standard portfolio analogy for diversity. According to the portfolio analogy, a firm wants diversity so as to be able to respond to diverse situations just as a stock investor wants a diverse portfolio of stocks. Just as a diverse portfolio guarantees a good payoff regardless of the state of the world, a diverse set of employees ensures that someone exists within the firm to handle any situation that arises. The portfolio analogy, though accurate in some cases, breaks down when applied to team-based problem solving. There's no give and take between stocks in a portfolio. One stock doesn't say to another stock, "I never thought of the problem that way." Nor can stocks build on solutions thought of by existing stocks. That just doesn't make any sense.

I do not mean to imply that diversity does not provide insurance as suggested by the portfolio analogy. It does. However, the value of insurance against risk should not obscure the potentially larger superadditive benefits that accrue from having employees with diverse perspectives and heuristics.

Before moving on to more theoretical results, I

want to inject a brief comment about identity diversity. In the framework that I have described, diverse perspectives and heuristics underpin diversity's benefits. These more cognitively based notions of differences are distinct from identity-based distinctions such as race, gender, age, ethnicity, and so on. Though conceptually distinct, cognitive and identity diversity often correlate empirically. This correlation arises because perspectives and heuristics that people apply to problems do not come from thin air. They are the product of training, practice, and life experiences. How we see the world is informed and influenced by our values, our identities, and our cultures. People often reason by analogy. Each person's unique set of life experiences provides the engine for these analogies. Diverse identities, therefore, often translate into diverse perspectives and heuristics.

## Problem Solving: Diversity Trumps Ability

I have just outlined the basic logic for how diverse perspectives and heuristics can improve problem solving. I now want to push this logic a little further and touch on some formal results. First, I want to describe some experiments that I ran while an assistant professor at Caltech. For fun, I constructed a computer model of diverse problem solvers confronting a difficult problem. In my model, I represented diversity as differences in the ways problem solvers encoded the problem and searched for solutions, i.e., diverse *perspectives* and *heuristics*. I then stumbled upon a counterintuitive finding: Diverse groups of problem solvers—groups of people with diverse perspectives and heuristics—consistently outperformed groups composed of the best individual performers. So, if I formed two groups—one random (and therefore diverse) and one consisting of the best individual performers—the first group almost always did better. In other words, *diversity trumped ability*.

This counterintuitive finding led me to try to identify sufficient conditions for this to be true. What assumptions did I have to make for diverse groups, on average, to outperform groups of the best individuals? That turned out to be a rather difficult task. So, following the logic of my own model, I enlisted the help of someone else, Lu Hong, a person with a different set of perspectives and heuristics than my own, to help me identify those conditions. Together, we found a set of conditions that, if they hold, imply that diversity trumps ability.

To show what these conditions are and why they matter, I will describe a simple model. Suppose that I begin with an initial pool of problem solvers from which I draw a random (e.g., diverse) team and a team of the best individual problem solvers. Each of these teams will have some moderate number of people, whereas the initial pool of people could be quite large. It could consist of everyone who works for a firm or every faculty member at a university. I then compare the collective performance of the team of the best problem solvers against the collective performance of the randomly selected problem solvers.

Before I go too far, I want to remind you of the goal. Keep in mind that the diversity-trumps-ability result won't always hold. It holds given certain conditions. If, for example, the teams have only a single member, the team of the best problem solvers will consist of the best individual, and the team of random problem solvers will consist of a random person. Therefore, the first team will outperform the second. Of course, in this case ability doesn't trump diversity because the second team isn't diverse. It has only one person. Thus, having the teams have more than one person will be a condition for the result to hold.

The question Lu and I asked was, what other conditions are needed? If those conditions are unrealistic, then we should not expect diversity to trump ability in practice. If those conditions seem mild, then perhaps we should. That's one reason that we "do the math," so that we can see when logic holds and when it doesn't. Doing the math has other benefits as well, not the least of which is that we better understand how diversity produces benefits, which better enables us to leverage it in practice.

The first condition we identified relates to the difficulty of the problem. Easy problems don't require diverse approaches.

**Condition 1: The Problem Is Difficult:** *No individual problem solver always locates the best solution.*[1]

Without this condition, diversity cannot trump ability. If any individual problem solver always finds the best solution, then the collection of the best problem solvers (which by definition contains *the* best problem solver) always locates the best solution. For example, if we need to find the answer to a standard engineering problem, we can just ask an engineer who can give us the correct answer. We have no need to put together an interdisciplinary team. For harder problems, like designing an aircraft engine, we need a team. And that team needs diverse thinkers.

**Condition 2: The Calculus Condition:** *The local optima of every problem solver can be written down in a list. In other words, all problem solvers are smart.*

The second condition concerns the ability of the problem solvers. All of the possible problem solvers must have some ability to solve the problem. We cannot set loose a bunch of anthropologists and economists in the physics lab and hope they produce cold fusion. To formalize the idea that the problem solvers must have relevant cognitive tools, Lu and I assumed that the problem solvers got stuck in only a reasonable number of places. In the language of mathematics, such points are called *local optima*. We decided to call this restriction the *Calculus Condition*. We did this because people who know calculus can take derivatives, and therefore have a reasonable number of local optima. Here's why. Think of a problem as creating a mathematical function in which high values are good solutions. The derivative equals the slope of that function, which like the slope of a mountain is either positive (uphill), negative (downhill), or zero (on a peak or a plateau). On a peak the derivative equals zero; the slope goes neither up nor down. Calculus enables a person to find points with derivatives equal to zero. Therefore, people who know calculus can find peaks. Economists don't know calculus when it comes to physics, but they probably do know calculus when asked about tax policy.

**Condition 3: The Diversity Condition:** *Any solution other than the global optimum is not a local optimum for some non-zero percentage of problem solvers.*

The third condition requires that for any proposed solution other than the global optimum, some problem solver can find an improvement on that solution. In formal terms, this means that the intersection of the problem solvers' local optima contains only the global optimum. We call this the *Diversity Condition*, as it assumes diversity among the problem solvers. This condition does *not* say that given any solution some problem solver can immediately jump to the global optimum. That assumption would be much stronger and would rarely be the case. The assumption says, instead, that some problem exists who can find an improvement. That improvement need not be large. It need only be an improvement.

**Condition 4: Reasonably Sized Teams Drawn from Lots of Potential Problem Solvers:** *The initial population of problem solvers must be large, and the teams of problem solvers working together must consist of more than a handful of problem solvers.*

The final condition requires that the initial pool of problem solvers must be reasonably large and that the set of problem solvers who form the teams must not be too small. The logic behind this condition becomes clear in extreme cases. If the initial set consists of only 15 problem solvers, then the best ten should outperform a random ten. With so few problem solvers, the best ten cannot help but be diverse and therefore have different local optima. At the same time, the teams that work together must be large enough that the random collection can be sufficiently diverse. Think of it this way: We need to be selecting people from a big pool, and we need to be constructing teams that are big enough for diversity to come into play.

These four conditions—(a) the problem has to be hard, (b) the people have to be smart, (c) the people have to be diverse, and (d) the teams have to be reasonably big and chosen from a large pool—prove sufficient for diversity to trump abil-

[1] If the best problem solver finds the optimal solution 99.9% of the time, the collection of randomly selected problem solvers will not outperform the group of the best.

ity. They're not the only conditions under which the result holds, but if they're satisfied, diversity will almost always trump ability.[2]

**The Diversity Trumps Ability Theorem:** *Given conditions 1 to 4, a randomly selected collection of problem solvers almost always outperforms a collection of the best individual problem solvers.*

This theorem is no mere metaphor, cute empirical anecdote, or small-sample empirical effect that may or may not be true with more trials. It's a logical truth like the Pythagorean Theorem (Hong & Page, 2004). The reason diversity trumps ability is not deep: The best problem solvers likely have similar perspectives and heuristics. The random problem solvers bring diverse ways of thinking. Therefore, the best problem solvers all get stuck in the same places. The random problem solvers don't.

The Diversity Trumps Ability Theorem implies that hiring people of high individual abilities may be less important than hiring people with diverse skills if those employees will work as part of a team. The logic of the theorem does not imply the irrelevance of ability. People need not remove those "my child is an honor student at Neil Armstrong Junior High" bumper stickers from their minivans, nor should universities randomly allocate admission slots. Ability still matters, but so does diversity. And, as the theorem shows, once an ability threshold has been met, diversity matters more than ability.

These comparisons between diversity and ability require some care. Comparing ability to diversity is not unlike comparing a shiny apple to a fruit basket. Ability is a property of an individual—a nice shiny apple. Diversity is a property of a collection of people—a basket with many kinds of fruit. Rather than think of them as opposing concerns, we should see diversity and ability as complementary: The better the individual fruits, the better the fruit basket.

Problem solving is a central task for many organizations, but it is not the only task for which diversity improves performance. Diversity also improves collective predictions. As I now show, the so called "wisdom of crowds" comes not from having just smart people in the crowd, but from having smart people with diverse predictive models.

## Predictive Models and the Wisdom of Crowds

When a company decides which product to launch, when venture capitalists decide where to invest, and when stock analysts decide when to buy or sell, they're making predictions. In putting together teams of people to make predictions, we might think that we want smart people, i.e., accurate individual predictors. And that's true. But it's also true that we want diverse predictors. We want people who differ in how they make predictions. Diversity should not be a second-order concern—multicolored sprinkles on the cake of ability; it merits equal billing. As in problem solving, diversity matters just as much as ability, and ideally, an organization or team would have an abundance of both.

To show the value of diversity in prediction, I need to define formally what I mean by a *predictive model*. Predictive models rely on *interpretations*. Interpretations are the mappings we make from the real world into categories. Categories, in turn, are conceptual boxes, or placeholders. For example, if I see a restaurant called Del Churro, I place it in the category *Mexican restaurants*. That may be true or it may not (most likely it is). I then predict that I'll enjoy eating there because I like Mexican food.

These interpretations underpin statistical predictive models as well. When stock analysts run regressions they restrict what dimensions, or attributes, they consider. In doing so, they create boxes. They use these boxes to construct a predictive model. In the best-selling book *Blink*, Malcolm Gladwell describes several instances in which experts' predictive models use very simple interpretations (2005). Gladwell loads his book with examples, including the story of an expert who instantly recognized a multimillion-dollar sculpture as fake even though scientific analysis had found otherwise, and one of an expert who can predict (using lots of analysis) whether a married couple will stay together just by looking at a

[2] In mathematics, the phrase *almost always* means with probability one.

few dimensions of their relationship. Gladwell's work shows the value of simple heuristics, an idea that receives a more formal treatment in the work of Gerd Gigerenzer and Peter Todd (1999).

As my example of the Mexican restaurant suggests, we need not think of predictive models as applying to just important events like stock market price changes or the causes of a disease. We apply predictive models almost every time we think. And our predictive models rely on interpretations. A popular predictive model for when a television show has reached its peak relies on categorizing episodes based on specific events. "Jumping the shark" (a reference to Fonzie jumping over a shark tank on his motorcycle, which signified the long decline of *Happy Days*, can take many forms. It could be a wedding that resolves building tension, or a death. It could be the appearance of a special guest star. Nancy Reagan showing up on the television program *Diff'rent Strokes* was clearly jumping the shark.[3]

In making these predictions, be they about television shows or IPOs, people rely on predictive models that in turn rely on interpretations. Note that predictive models and interpretations differ from heuristics and perspectives. An interpretation categorizes part of the world. It's the mapping of that big gray cloud into the category "rain cloud." A predictive model tells us what we think will happen: "It looks like rain." Predictive models are thoughts. Heuristics are courses of action. A heuristic tells us what to do: "It's raining—let's run for cover," or what not to do: "We get just as wet by running, so let's walk." A perspective is not an action. It is a representation of the world. Each person possesses all of these: perspectives, heuristics, interpretations, and predictive models. And each of us differs in the particular collection of these tools that we hold inside our heads.

[3] See www.jumptheshark.com.

## The Diversity Prediction Theorem

Having constructed a model of how people make predictions, I can now turn to analyzing the role that diversity plays in the ability of a team, group, or crowd to make a prediction. I am going to consider some real-world data to show the importance of diversity. If we are going to look at some data, we might as well look at something important. So let's look at NFL draft predictions. (The actual reasons for considering this example are that the predictors have a stake in being correct and sufficient variance exists in the predictions to make the case interesting.)

The table below shows predictions for the top dozen picks in the 2005 NFL draft from seven

**Table 1**
**Experts' Predictions of 2005 NFL Draft**

| Player\Expert | Wright | Adler | Fanball | SNews | Zimm | Prisco | Judge | Crowd |
|---|---|---|---|---|---|---|---|---|
| Smith | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1.0 |
| Brown | 2 | 2 | 4 | 2 | 2 | 5 | 2 | 2.7 |
| Edwards | 3 | 3 | 2 | 7 | 3 | 2 | 3 | 3.3 |
| Benson | 4 | 4 | 13 | 4 | 8 | 4 | 8 | 5.9 |
| Williams | 8 | 5 | 5 | 5 | 4 | 13 | 4 | 6.4 |
| Jones | 16 | 9 | 6 | 8 | 6 | 6 | 9 | 8.1 |
| Williamson | 13 | 14 | 12 | 12 | 13 | 7 | 7 | 9.7 |
| Rolle | 6 | 6 | 8 | 10 | 9 | 8 | 6 | 7.9 |
| Rogers | 9 | 8 | 9 | 9 | 16 | 9 | 9 | 9.9 |
| Williams | 7 | 7 | 7 | 6 | 7 | 12 | 12 | 8.0 |
| Ware | 11 | 15 | 14 | 24 | 11 | 11 | 13 | 13.9 |
| Merriman | 12 | 11 | 3 | 11 | 12 | 10 | 11 | 10.1 |
| **Sq Error** | 158 | 89 | 210 | 235 | 112 | 82 | 75 | 34.4 |

prognosticators.[4] The players are listed in the order that they were selected. Each predictor provides a ranking of the draftees. The names in the columns identify the predictors. These experts' predictions came from detailed analyses. They don't call them draft experts for nothing. These people—er, men—devote long days and nights to evaluating team needs, player skills, and a host of other factors.

The last row of the table shows the total squared error for each predictor. I calculated this number by summing the squared errors of the prediction for each player. To calculate that number, I take the actual draft position of each player, subtract the predicted position, and square the difference. So, when the *Sporting News* (SNews) predicted Braylon Edwards to go seventh, their error on that pick was $(3 - 7)^2$, which equals 16. If we look at the errors across the predictors, we see that they differ in their accuracy. The best has an error of 75. The worst predictor has an error of 235. The average of the individual errors equals 137.3. Comparing the accuracy of the individual predictors to the accuracy of the crowd reveals that the crowd is more accurate than any of its members.[5] That won't always be true, but the crowd will always be more accurate than its average member. The wisdom of crowds, therefore, does exist, but the brilliance of crowds does only every once in a while.

By itself, this example doesn't prove that diversity is valuable. It just shows that in this one instance a diverse group of predictors was more accurate than any member of the group. To show why diversity deserves the credit for the crowd's success, I need to introduce one more statistical term, which I call the *prediction diversity*. Prediction diversity is nothing more than the variance of the experts' predictions. A little math shows that in the NFL example, the prediction diversity equals 102.9. Notice the relationship between the crowd's error (34.4), the average individual error (137.3), and the prediction diversity (102.9): *Collective error equals average error minus prediction diversity*. This equality is not an artifact of our example. It is always true. I call this the *Diversity Prediction Theorem*.

**The Diversity Prediction Theorem:** *Collective Error = Average Individual Error − Prediction Diversity*

Let's think for a moment what this theorem means. It says that prediction diversity matters just as much as individual prediction accuracy when putting together a crowd of predictors. Equations such as this move us from some loose intuition that diverse points of view might be useful to an explicit characterization of how useful. Diversity isn't just something of marginal value. *Diversity matters just as much as individual ability*. That is not a feel-good statement. It's a mathematical fact.

## Putting the Logic to Work

These two theorems—the Diversity Trumps Ability Theorem and the Diversity Prediction Theorem—provide a foundation for claims that diversity provides benefits. That is, as they say, a good thing. But feeling good is not enough. Organizations can use this logic as more than a justification for policies that seek out diverse employees. Organizations can leverage this logic to be more innovative and productive. In what follows, I describe some direction for how this might be done.

### Lesson #1: Move Beyond the Portfolio Analogy and Promote Interactions

As I mention above, many arguments for diversity lean on the portfolio analogy. For the same reason that a financial adviser advocates building a diverse portfolio of stocks, a firm should have diverse employees. The portfolio analogy sees diversity as a form of insurance. And it's true that diversity often performs that function. However, the portfolio analogy misses a key part of the logic: *the "superadditivity" of diverse tools*. People have perspectives, heuristics, interpretations, and predictive models.

When a collection of people work together to

[4] The analysts are Scott Wright from *NFL Countdown*, James Adler from About.com, the Fanball Staff at Fanball.com, the *Sporting News*, Paul Zimmerman from *Sports Illustrated*, and Pete Prisco and Clark Judge from *CBS Sportsline*.

[5] The crowd, in this case, is the sum of all of the analysts.

solve a problem, and one person makes an improvement, the others can often improve on this new solution even further. Problem solving is not the realization of a state but a process of innovation in which improvements build on improvements. This superadditivity can be found in many real-world examples. For example, attendees at the 1904 St. Louis World's Fair could choose from a wide array of food choices: ice cream, cookies, cakes, waffles, and so on. One hot day during the fair, an ice cream vendor ran out of cups. A Syrian waffle vendor in the booth next door named Ernest Hami improvised by rolling up waffles to make cones. The rest, as they say, is history.[6] The parts of the portfolio—the waffles and the ice cream—combined to create something new, and better: the ice cream cone. The key here is that the waffles and the ice cream interacted, and through that interaction produced a superadditive benefit. Diverse teams of people can produce similar gains, but they need to interact in order to do so.

### Lesson # 2: Contain Multitudes

The logic I have presented implies that rather than having a single perspective, interpretation, heuristic, or predictive model, people and organizations should have many. We must become Whitmanesque and contain multitudes. The advantages of containing multitudes should be clear. Diverse perspectives and heuristics improve problem solving. Diverse interpretations and predictive models lead to more accurate predictions. Crowds are not wise, but crowds of models are.

One way to maintain this diversity is to mimic evolution. From evolution, we know that diversity together with crude selective pressures can solve hard problems. In evolution, genetic mutation maintains diversity. Those mutations that increase fitness survive; those that do not fall by the wayside. The same effects occur within groups of people. Good attempts survive. Bad ones don't. Experimentation can lead to a better "best" individual performer. More important, it can result in better collective performance. Increasing diversity improves collective performance at prediction and problem solving.

Consider the following thought experiment. Suppose that we have to predict the amount of leather produced by a cow. This requires knowing the surface area of a cow. Even the complicated surfaces from calculus class are far more regular than your average cow. Fortunately, a book by John Harte (1985) on modeling offers a solution to this problem. *We can imagine a spherical cow*. I am going to ignore how we'd milk this spherical cow. Calculating the surface area of the spherical cow requires high school level math. That number won't be correct. It's an estimate, a prediction.

Someone else might decide to construct a different predictive model. She might imagine cows shaped like boxes. She might even tape together a few Gateway computer boxes until she reached cow-like proportions. Someone else might imagine elliptical cows. Either of these two other models may prove more accurate than the spherical cow model. If so, that's great. That doesn't mean that we should toss out the spherical cow model. The greatest benefit may well come from having multiple models that can be averaged into a crowd. The crowd of cow models may well be better than the best.

The amount of experimenting that makes sense depends upon the circumstances. Clearly, the lower the costs of experimenting and the more important the problem, the more we should experiment. We should also err on the side of more searching when a problem is connected to other problems. If we can understand how proteins fold, we can make headway on lots of other problems. Cognitive tools flow freely between domains. And by combing tools, we can find even larger breakthroughs (Axelrod & Cohen, 2000).

### Lesson #3: Look Outside: Consult Dissenters

When an organization confronts problems, it may lock in on a particular perspective. In an organization, common perspectives facilitate communication and the development of more advanced heuristics, but they also create common local optima. Thus, if one person gets stuck and if everyone sees the problem the same way, then everyone is stuck. Now, it could be that an organization's

[6] Unbeknownst to Hami, Italo Marchiony, a recent Italian immigrant to New York, had patented the ice cream cone in 1903.

shared perspective makes every problem easy so that solutions are always optimal. Experience suggests otherwise. The only organizations that always make optimal decisions lie nestled within the pages of introductory economics books. Given that organizations make mistakes and sometimes do so systematically, every so often they bring in people from the outside. These outside consultants are not necessarily smarter than the people who already work there. They're just different. They bring with them different perspectives and heuristics that can improve outcomes.

An example from within the academy provides empirical grounding. Most universities organize themselves into departments. This hinders diversity for obvious reasons. These departments largely monitor themselves, so deans and provosts look for signs of external validation to see how these departments perform, such as the placement of graduate students, the number of publications by faculty in top journals, and the frequency of attempts by other schools to hire away faculty. These signals indicate if a department performs well, but they provide almost no clue for *how* a department could perform better. For this reason, universities periodically invite committees composed of scholars from other schools who work in the same discipline or a closely related one to provide suggestions for how the department might improve. How do they do this? They gather information about the current state of the department and advocate certain changes. Are these visitors more able than the people in the department? Probably not. But they are different. And they leverage those differences to make improvements.

These visiting committees can be thought of as a type of consultant. In fact, they are consultants (they just do not get paid as much as real consultants). And, clearly, this same line of thinking explains a benefit of consultants: They're able to provide diversity to help departments or companies improve. Sure, some companies trot out highly paid consultants in fancy suits to add credibility to decisions that directors have already made—"Look, McKinsey agrees with me!" And yes, some consulting companies perform services that firms do not have the capacity or ability to do themselves, but many consultants do consult. And when doing so they make improvements. (No, really.) Otherwise, there would not be so many consulting companies, and consultants wouldn't be paid so much money. But the fact that these consultants add value does not mean that they are giants of the earth, smarter and more capable than others. A freshly minted MBA need not know more about dog food than Purina or more about manufacturing processes than General Motors or Toyota to add value. She need only be moderately capable and different. In difference lies value.

The careful reader will notice the subversive nature of this logic. I might have described visiting committees and consultants as experts, but I did not. Instead, I've described them as people who think differently. Visiting committees and consultants challenge the status quo. They are what Cass Sunstein (2003) might call "dissenters." In politics dissenters identify new policy dimensions, and they force us to abandon our existing predictive models. Dissent is useful. Without it societies would falter. Organizational consultants—whether academic, nonprofit, or for profit—are dissenters too, paid dissenters.

### Lesson # 4: Create Prediction Markets

Given the potential wisdom of crowds, an organization might benefit from creating internal information markets to make predictions. Information markets have substantial appeal to businesses and organizations. They can be highly accurate and low cost. Currently, most large companies and organizations hire people to construct models to predict future demands, sales, or, in the case of political parties, votes. Without these predictions long-range planning becomes difficult, if not impossible. By creating an internal prediction market, an organization can leverage the wisdom of its own crowd. This prediction market could supplement or even replace the experts' predictions. Some companies, such as Hewlett Packard and Google, have already done this. Chen and Plott (2002), for example, report that Hewlett Packard used managers to predict printer sales. The managers' predictions proved to be as good as, and in some cases better than, the experts'.

Consider an auto company that wants to predict what types of cars will sell best in the coming

five years—a prediction that auto manufacturers address regularly, and one that they often get wrong. They could set up an internal market that includes all of the company's engineers. This probably would not work. Their engineers probably do not have much knowledge about consumer trends—they're engineers! They may as well ask them to predict the Oscars. These information markets require that participants have reasonable models. For this reason, if these same engineers were asked to predict which of two vehicle designs would prove more durable, the information market would perform well. For this task, the engineers possess diverse and reasonable models (they understand different parts of the vehicle). Owing to that diversity, they can collectively predict well.

### Lesson # 5: Focus on Relevant Diversity

For organizations, what counts is relevant diversity, and again how much they should weigh diversity relative to individual performance depends on the context. A firm that is hiring people for a job for which they have a well-defined ability measure may not reap many benefits from diversity. This would be true for a company hiring people to paint houses or to deliver messages by bicycle.

In contrast, consider a firm hiring people to design web pages. These potential hires would have to work together either directly or indirectly. In this case, the firm would want to consider diversity as much as ability. The firm should look for able people with diverse training, experiences, and identities. Unfortunately, human resources professionals can't just look at someone and see her perspectives or heuristics.

The opaqueness of cognitive differences explains why firms interview, administer tests, ask for recommendations, and sample previous work. They're trying to make inferences about the tools applicants possess. Someone with a computer science degree probably knows more programming heuristics than someone with a degree in biology. And someone who has worked for five years selling cars probably brings finer and more interesting interpretations of consumer types than someone who has been confined to a cubicle writing manuals for DVD players. But what of the undergrad riding the skateboard with all of those tattoos and piercings? Many people in corporate America would think, "He looks different from us." Does that mean that they should hire him? The answer depends. If the kid on the skateboard knows the equivalent of calculus for the problem—if the job at hand is, say, designing bowling shirts or tennis shoes—they might want to think about doing so. But if the firm invests in derivatives and the skateboarder stopped taking math classes in fifth grade, then the firm would do better to look elsewhere.

Enlightened employers seek out diversity. Owing to their success, Google can hire almost anyone they like. Google could just hire the top students from the engineering schools like MIT, Caltech, Stanford, Illinois, Michigan, Georgia Tech, and Cal-Berkeley. These people would all be smart, but they might be trained similarly. They also might have had similar college experiences. And they might be far from representative in the identity groups to which they belong.

Given that Google organizes itself in work teams that solve problems, success depends on both ability and diversity. That's why Google doesn't pursue a strategy of hiring only the people with the best grades from the best schools. In their own description of "who we're looking for," Google's first criterion is diversity—"people with broad knowledge and expertise in many different areas of computer science and mathematics"—as is their last: "people with diverse interests and skills." People who think alike get stuck. So Google samples widely. They look for diversity in training, experience, and identity. Computer science graduates from Santa Clara work alongside former math professors. But Google is also aware of the Calculus Condition. They seek diverse people with knowledge in mathematics and computer science. They're not seeking poets. That said, if a good mathematical epidemiologist showed up at their door, they'd hire her.

Many identity attributes correlate with or influence how we think. Leveraging diversity requires more than greater racial and gender balance. Forgetting this can result in lost opportunities. The United States Army has substantial

identity diversity at every rank. But because of the hierarchical nature of the military, they do not have much age diversity within a rank, so the people making the same kinds of decisions and giving advice to the same people are likely all about the same age. This reduces perspective and predictive model diversity. Some of the strongest evidence in all of the empirical diversity literature relates to demographic diversity. Those who arrive at the same time think the same way (Pfeffer, 1982). Therefore, maintaining age diversity can be crucial to success.

Firms might also test applicants for cognitive diversity relative to one another and to their current employees. Testing for diversity isn't as hard as it sounds. One consulting company asks job applicants to predict the annual sales for a standard household product, something like rubber bands, peanut butter, lug nuts, or size C batteries. This company wants to identify applicants who understand that total demand equals the sum of individual demands (recall the Calculus Condition in the Diversity Trumps Ability theorem). They also want to identify people who think differently. They achieve both goals by learning if and how the applicants segment the market of consumers. Among those applicants who get the question about the C batteries, those who parse households in interesting ways, such as households with male children, have a good chance of getting hired. Those applicants who divide the country into regions probably do not. And yet the company wouldn't want all people who identified young boys as big users of batteries. They'd want some people who identified other market segments, like campers. Asking silly questions doesn't just get silly answers, it reveals diversity in thinking. That's why Google asks prospective employees how many golf balls fit in a school bus.

### Lesson #6: The Samuel Paul Bowie Caveat

We can go too far in pursuing skill difference. In our pursuit of diversity, we must keep in mind the need to balance diversity with ability. We need only recall the 1984 NBA draft, in which the Portland Trail Blazers picked Samuel Paul Bowie, a seven-foot center from Kentucky, over a forward from North Carolina named Michael Jordan, perhaps the greatest basketball player of all time. Reasons for the Bowie pick vary. Some claim that Jordan's talents had been obscured by North Carolina's team-oriented style of play.

I'm willing to cut the Blazers some slack. Portland's error could have resulted from having the wrong predictive model. Portland executives had reason to believe in the value of a good center. They had won a title just a few years earlier with an injury-prone Bill Walton at the pivot. Further supporting their case, only one team from 1959 to 1984 had won an NBA title without an all-star center. Add to this the fact that Portland already had an excellent tandem at small forward and shooting guard—Clyde Drexler and Jim Paxson—and the Bowie decision looks reasonable. But with the benefit of hindsight, choosing Bowie, an example of choosing diversity over ability, looks silly. If Michael Jordan's available, draft him. Sometimes ability trumps diversity.

### Lesson #7: Avoid Lumping by Identity and Stereotyping

Employers often use identity as a crude proxy for cognitive diversity. And it's true that the types of cognitive diversity that I've discussed correlate with identity. Even so, organizations probably can do better than to rely on coarse identity classifications to categorize people. People are multifaceted and multi-tooled. We all have different experiences and training as well as different identities. Those experiential and training differences also translate into diverse toolboxes.

Mapping people into identity groups often over-lumps. Placing a recent immigrant from Nairobi, Kenya; a grandson of a sharecropper from the Mississippi Delta; and the daughter of a dentist from Barrington, Illinois, into the same category—African Americans—obscures differences, as does placing the granddaughter of a miner from Copper Harbor, Michigan; a son of Gloria Vanderbilt (that would be Anderson Cooper); and a recently married former au pair from Lithuania into the box labeled "non-Hispanic white." Similarly, having an Asian American box that lumps together people whose ancestors came from Singapore, Malaysia, China, Japan, and Korea bunches together diverse cultural identities. Each

of these lumps of people, if unpacked, would prove cognitively diverse.

This lumping also ignores combinations of identities. A group consisting of five French men, three Korean men, two Kenyan women, and a woman from Singapore contains a man and Kenyans but it does not contain Kenyan men and therefore may not be able to look at the problem in the same way that a Kenyan man might. And again, there is no single way in which a Kenyan man would look at a problem.

This insight also can be used to temper our enthusiasm for pipelines used to recruit minorities. These programs nurture potential employees or students from underrepresented groups. They may improve numbers, but they can limit the amount of cognitive diversity that a firm gets. By hiring only African American engineers who graduated from Berkeley and attended the same summer internship program, a company like Cisco sacrifices cognitive diversity on the altar of identity accounting. Their employees look different, but they may not think differently. Thus, the use of pipelines probably has a negative effect on the benefits of diversity. It probably reduces the performance of identity-diverse firms. The greater identity diversity gained through the pipeline could be more than offset by their lack of experiential, demographic, or training diversity. Far better that Cisco forms a consortium of companies to create multiple pipelines to obtain what might be called within-lump diversity. Or even better, perhaps society might be structured in such a way that those pipelines are not needed.

Lumping people by identity group creates stereotypes and stigmatization. Many people think men are smarter than women, that people who grew up on farms work harder, and that Italians can cook better than the English. We describe people as typical Europeans or as fraternity boys. These stereotypes are predictive models. They place people in categories and make predictions based on those categorizations. If informed by lots of experience, these predictive models may be more accurate than not, provided we've lumped correctly. It is probably empirically true that on average, Italians probably are better cooks than English people, and frat boys do eat a lot of food (and drink a lot of beer). But some do not. No evidence suggests that men are on average smarter than women.

Stereotypes, therefore, are to be avoided. In addition to being crude predictive models, they create three other problems. First, because stereotypes are predictive models about people, and not about physical phenomena, they can influence behavior and become self-reinforcing (Jackson & Fryer, 2002). People may evaluate women as less effective than men at task performance, even if by objective standards the women perform as well. This might happen if enough people carry around a crude predictive model that says that men are better workers than women. This can reduce incentives for women to work hard, and thus the stereotypes become self-fulfilling. Any stereotype—that Asians do better in math, that Indians are better spellers, that British people are wittier, or that African Americans are more creative—can induce self-fulfilling behavior. If we make stereotypical inferences about people who belong to an identity group, we reduce their incentives to accumulate tools outside these stereotypes. We limit opportunity. To use Glenn Loury's phrasing (2000; 2002), stereotypic predictive models stigmatize.

Second, stereotypes hinder collective performance by restricting how people think. People feel compelled to represent their identity groups when they are underrepresented in a group. If a person totes along the fundamental preferences of his identity group, he may lose track of the organization's goals (Brewer & Brown, 1998; McGrath, Arrow, & Berdahl, 2000; Tafjel & Turner, 1986). At the societal level, identity-based concerns with justice and equality of opportunity have a place, a central one. Organizations should not, like Mussolini, concern themselves only with having the trains run on time. Trying to act and think as a woman or an Asian or a black male, instead of as oneself, hurts group performance. People need not strip themselves of their identities, but they shouldn't let their core identities confine them.

To use the language of Kwame Appiah (2005), identities root us. They should provide us with meaning and purpose, but they should not limit

us. People should be allowed to be different, to possess multiple identities, and to pursue a range of experiences and training. As Toni Morrison put it, "In *Tar Baby*, the classic concept of the individual with a solid, coherent identity is eschewed for a model of identity which sees the individual as a kaleidoscope of heterogeneous impulses and desires, constructed from multiple forms of interaction with the world as a play of difference that cannot be completely comprehended."[7]

Finally, the use of stereotypes limits predictive diversity when evaluating people. By definition, stereotypes are widely shared predictive models. By applying stereotypes, people are not thinking differently at the individual level, and collectively won't make accurate predictions. This logic underpins the Diversity Prediction Theorem. Good collective predictions require abandoning stereotypes. People should look for attributes or categories that differ from those of others but still make sense.

## Lesson #8: Maintain Humility in the Face of Mystery

My final piece of advice returns to the question of why organizations should promote diversity. The frameworks that I have presented can be used to support proactive diversity policies. They show that individual diversity contributes to collective benefits. These results are theoretical, not empirical. They do not necessarily imply that companies, organizations, and universities that hire and admit diverse people can expect instant results, and that ability should be sacrificed on some altar of difference. To the contrary, the analysis suggests a need to balance the two. In the Diversity Trumps Ability Theorem, this takes the form of a threshold the problem solvers must satisfy (the Calculus Condition). In the Diversity Prediction Theorem, the tradeoff is more direct. Sometimes, organizations should trade some ability for diversity, but ideally, they should seek more of both.

I previously discussed why high-tech firms wouldn't want only freshly minted graduates from MIT and Caltech. People with different training and experiences may add more to the conversation than people who score better according to traditional measuring sticks. Employers and universities need to understand this logic and hire and admit accordingly.

The link between identity diversity and cognitive diversity is more subtle and mysterious. Nevertheless, for similar reasons, leading companies and universities shouldn't want all white men or all Asian women. Identity diversity often correlates with cognitive diversity and often does so strongly. The extent of that correlation depends upon the problem. As life experiences often frame how people see social issues, for public policy problems identity differences can translate directly into diverse perspectives. On more scientific and technical problems, the linkages are less direct. Yet this does not mean that they do not exist. The sources of innovation remain mysterious; life experiences can serendipitously provide insights. By building diverse teams of employees, organizations increase their chances of making a breakthrough.

Organizations should, therefore, continue to pursue pro-diversity policies, but those policies should reflect the potential for diversity to improve outcomes. Diversity matters not just because it is the right thing to do. Diversity matters because it can increase the bottom line by introducing more perspectives, heuristics, interpretations, and predictive models. These diverse cognitive tools can in turn improve an organization's ability to solve problems and make accurate predictions.

[7] See womenshistory.about.com/od/quotes/a/toni-morrison.htm

## References

Appiah, K. A. (2005). *The ethics of identity*. Princeton, NJ: Princeton University Press.

Axelrod, R., & Cohen, M. D. (2000). *Harnessing complexity: Organizational implications of a scientific frontier*. New York: Free Press.

Brewer, M.B., & Brown, R.J. (1998). Intergroup relations. In D.T. Gilbert, S.T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology*. Boston: McGraw-Hill.

Chen, K-Y., & Plott, C. R. (2002). Information aggregation mechanisms: Concept, design, and implementation for a sales forecasting problem. Caltech Working Paper #1131.

Ely, R. J., & Thomas, D. A. (2001). Cultural diversity at work: The moderating effects of work group perspectives

on diversity. *Administrative Science Quarterly, 46,* 229–273.

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Gladwell, M. (2005). *Blink: The power of thinking without thinking*. New York: Little, Brown, and Co.

Harte, J. (1985). *Consider a spherical cow: A course in environmental problem solving*. Los Altos, CA: W. Kaufmann.

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.

Jackson, M. O., & Fryer, R. G., Jr. (2002). Categorical cognition: A psychological model of categories and identification in decision making. Working Paper 1144, California Institute of Technology, Division of the Humanities and Social Sciences.

Lager, F. (1994). *Ben & Jerry's: The inside scoop: How two real guys built a business with a social conscience and a sense of humor*. New York: Crown Publishers.

Loury, G. C. (2000). Racial stereotypes. Lecture delivered at Harvard University as part of the DuBois Lectures, Cambridge, MA.

Loury, G. C. (2002). *The anatomy of racial inequality*. Cambridge, MA: Harvard University Press.

Mannix, E., & Neale, M. (2006). What differences make a difference? The promise and reality of diverse teams in organizations. *Psychology in the Public Interest*, forthcoming.

McGrath, J. E., Arrow, H., & Berdahl, J. L. (2000). The study of groups: Past, present, and future. *Personality and Social Psychology Review, 4*(1), 95–105.

Nalebuff, B., & Ayres, I. (2003). *Why not?: How to use everyday ingenuity to solve problems big and small*. Cambridge, MA: Harvard Business School Press.

Pfeffer, J. (1982). *Organizations and Organization Theory*. Marshfield, MA: Pitman.

Sunstein, C. R. (2003). *Why societies need dissent*. Cambridge, MA: Harvard University Press.

Tafjel, H., & Turner, J.C. (1986). The social identity theory of intergroup behavior. In S. Austin & W.G. Austin (Eds.), *Psychology of intergroup relations*. Chicago: Nelson Hall.

Williams, K. Y., & O'Reilly, C. A., III. (1998). Demography and diversity in organizations: A review of 40 years of research. *Research in Organizational Behavior, 20,* 77–140.