

“Numbers Take Us Only So Far”

Facebook's global director of diversity explains why stats alone won't solve the problem of organizational bias. **by Maxine Williams**

■ WAS ONCE EVICTED from an apartment because I was black. I had secured a lovely place on the banks of Lake Geneva through an agent and therefore hadn't met the owner in person before signing the lease.

Once my family and I moved in and the color of my skin was clear to see, the landlady asked us to leave. If she had known that I was black, I was told, she would never have rented to me.

Terrible as it felt at the time, her directness was useful to me. It meant I didn't have to scour the facts looking for some other, nonracist rationale for her sudden rejection.

Many people have been denied housing, bank loans, jobs, promotions, and more because of their race. But they're rarely told that's the reason, as I was—particularly in the workplace. For one thing, such discrimination is illegal. For another, executives tend to think—and have a strong desire to believe—that they're hiring and promoting people fairly when they aren't. (Research shows that individuals who view themselves as



Originally published in
November–December 2007

objective are often the ones who apply the most unconscious bias.) Though managers don't cite or (usually) even perceive race as a factor in their decisions, they use ambiguous assessment criteria to filter out people who aren't like them, research by Kellogg professor Lauren Rivera shows. People in marginalized racial and ethnic groups are deemed more often than whites to be "not the right cultural fit" or "not ready" for high-level roles; they're taken out of the running because their "communication style" is somehow off the mark.

They're left only with lingering suspicions that their identity is the real issue, especially when decision makers' bias is masked by good intentions. I work in the field of diversity. I've also been black my whole life. So I know that under-

represented people in the workplace yearn for two things: The first is to hear that they're not crazy to suspect, at times, that there's a connection between negative treatment and bias. The second is to be offered institutional support.

The first need has a clear path to fulfillment. When we encounter colleagues or friends who have been mistreated and who believe that their identity may be the reason, we should acknowledge that it's fair to be suspicious. There's no leap of faith here—numerous studies show how pervasive such bias still is.

But how can we address the second need? In an effort to find valid, scalable ways to counteract or reverse bias and promote diversity, organizations are turning to people analytics—a relatively new field in business operations and talent management that replaces gut decisions with data-driven practices. People analytics aspires to be "evidence based." And for some HR issues—such as figuring out how many job interviews are needed to assess a candidate, or determining how employees' work commutes affect their job satisfaction—it is. Statistically significant findings have led to some big changes in organizations. Unfortunately, companies that try to apply analytics to the challenges of underrepresented groups at work often complain that the relevant data sets don't include enough people to produce reliable insights—the sample size, the *n*, is too small. Basically they're saying, "If only there were more of you, we could tell you why there are so few of you."

Companies have access to more data than they realize, however. To supplement a small *n*, they can venture out and look at the larger context in which they operate. But data volume alone won't

give leaders the insight they need to increase diversity in their organizations. They must also take a closer look at the individuals from underrepresented groups who work for them—those who barely register on the analytics radar.

Supplementing the *N*

Nonprofit research organizations are doing important work that sheds light on how bias shapes hiring and advancement in various industries and sectors. For example, a study by the Ascend Foundation showed that in 2013 white men and white women in five major Silicon Valley firms were 154% more likely to become executives than their Asian counterparts were. And though both race and gender were factors in the glass ceiling for Asians, race held 3.7 times the impact that gender did.

It took two more years of research and analysis—using data on several hundred thousand employees, drawn from the EEOC's aggregation of all Bay Area technology firms and from the individual reports of 13 U.S. tech companies—before Ascend determined how bias affected the prospects of blacks and Hispanics. Among those groups it again found that, overall, race had a greater negative impact than gender on advancement from the professional to the executive level. In the Bay Area white women fared worse than white men but much better than all Asians, Hispanics, and blacks. Minority women faced the biggest obstacle to entering the executive ranks. Black and Hispanic women were severely challenged by both their low numbers at the professional level and their lower chances of rising from professional to executive. Asian women, who had more representation at the professional level than other minorities, had the lowest chances of moving up from professional to executive. An analysis of national data found similar results.

By analyzing industry or sector data on underrepresented groups—and examining patterns in hiring, promotions, and other decisions about talent—we can better manage the problems and risks in our own organizations. Tech companies may look at the Ascend reports and say, "Hey, let's think about what's happening with our competitors' talent. There's a good chance it's happening here, too." Their HR teams might then add a layer of career tracking for women of color, for example, or create training programs for managing diverse teams.

Executives tend to think—and have a desire to believe—they're hiring and promoting people fairly when they're not.



Idea in Brief

Another approach is to extrapolate lessons from other companies' analyses. We might look, for instance, at Red Ventures, a Charlotte-based digital media company. Red Ventures is diverse by several measures. (It has a Latino CEO, and about 40% of its employees are people of color.) But that doesn't mean there aren't problems to solve. When I met with its top executives, they told me they had recently done an analysis of performance reviews at the firm and found that internalized stereotypes were having a negative effect on black and Latino employees' self-assessments. On average, members of those two groups rated their performance 30% lower than their managers did (whereas white male employees scored their performance 10% higher than their managers did). The study also uncovered a correlation between racial isolation and negative self-perception. For example, people of color who worked in engineering generally rated themselves lower than those who worked in sales, where there were more blacks and Latinos. These patterns were consistent at all levels, from junior to senior staff.

In response, the HR team at Red Ventures trained employees in how to do self-assessments, and that has started to close the gap for blacks and Latinos (who more recently rated themselves 22% lower than their managers did). Haile Corsetta, the company's VP of human capital, explained that the training "focused on the importance of completing quantitative and qualitative self-assessments honestly, in a way that shows how employees personally view their performance across our five key dimensions, rather than how they assume their manager or peers view their performance." She added: "We then shared tangible examples of what 'exceptional' versus 'solid' versus 'needs improvement' looks like in these dimensions to remove some of the subjectivity and help minority—and all—employees assess with greater direction and confidence."

Getting Personal

Once we've gone broader by supplementing the n, we can go deeper by examining individual cases. This is critical. Algorithms and statistics do not capture what it feels like to be the only black or Hispanic team member or the effect that marginalization has on individual employees and the group as a whole. We must talk openly with people, one-

on-one, to learn about their experiences with bias, and share our own stories to build trust and make the topic safe for discussion. What we discover through those conversations is every bit as important as what shows up in the aggregated data.

An industry colleague, who served as a lead on diversity at a tech company, broke it down for me like this: "When we do our employee surveys, the Latinos always say they are happy. But I'm Latino, and I know that we are often hesitant to rock the boat. Saying the truth is too risky, so we'll say what you want to hear—even if you sit us down in a focus group. I also know that those aggregated numbers where there are enough of us for the n to be significant don't reflect the heterogeneity in our community. Someone who is light-skinned and grew up in Latin America in an upper-middle-class family probably is very happy and comfortable indeed. Someone who is darker-skinned and grew up working-class in America is probably not feeling that same sense of belonging. I'm going to spend time and effort trying to build solutions for the ones I know are at a disadvantage, whether the data tells me that there's a problem with all Latinos or not."

This is a recurring theme. I spoke with 10 diversity and HR professionals at companies with head counts ranging from 50 to 300,000, all of whom are working on programs or interventions for the people who don't register as "big" in big data. They rely at least somewhat on their own intuition when exploring the impact of marginalization. This may seem counter to the mission of people analytics, which is to remove personal perspective and gut feelings from the talent equation entirely. But to discover the effects of bias in our organizations—and to identify complicating factors within groups, such as class and colorism among Latinos and others—we need to collect and analyze qualitative data, too. Intuition can help us find it. The diversity and HR folks described using their "spidey sense" or knowing there is "something in the water"—essentially, understanding that bias is probably a factor, even though people analytics doesn't always prove causes and predict outcomes. Through conversations with employees—and sometimes through focus groups, if the resources are there and participants feel it's safe to be honest—they reality-check what their instincts tell them, often drawing on their own experiences with bias. One colleague said, "The combination of qualitative and quantitative data

THE PROBLEM

Despite executives' belief that they hire and promote fairly, we know that people have been denied jobs and advancement because of their race. Managers often use ambiguous assessment criteria to filter out people who are not like them, and those in marginalized racial and ethnic groups are deemed more often than whites to be "not the right cultural fit" or "not ready" for high-level roles. Organizations need to find valid, scalable ways to counteract or reverse the biases behind these choices and to promote diversity.

PEOPLE ANALYTICS

People analytics offers some help by replacing gut decisions with data-driven practices. But not everyone registers on this radar—particularly not underrepresented groups. Hence organizations are left basically saying, "If only there were more of you, we could tell you why there are so few of you."

THE SOLUTION

When the numbers for an underrepresented group are too small to study in your organization, the author suggests including research on that population from the larger industry or context, learning from the experiences of other firms, and doing qualitative research into individual cases within your company. Because no single research method captures all the layers of bias, companies like Facebook are building cross-functional teams to face and address such challenges.



is ideal, but at the end of the day there is nothing that data will tell us that we don't already know as black people. I know what my experience was as an African-American man who worked for 15 years in roles that weren't related to improving diversity. It's as much heart as head in this work."

A Call to Action

The proposition at the heart of people analytics is sound—if you want to hire and manage fairly, gut-based decisions are not enough. However, we have to create a new approach, one that also works for small data sets—for the marginalized and the underrepresented.

Here are my recommendations:

First, analysts must challenge the traditional minimum confidence n , pushing themselves to look beyond the limited hard data. They don't have to prove that the difference in performance ratings between blacks and whites is "statistically significant" to help managers understand the impact of bias in performance reviews. We already know from the breadth and depth of social science research about bias that it is pervasive in the workplace and influences ratings, so we can combine those insights with what we hear and see on the ground and simply start operating as if bias exists in our companies. We may have to place a higher value on the experiences shared by five or 10 employees—or look more carefully at the descriptive data, such as head counts for underrepresented groups and average job satisfaction scores cut by race and gender—to examine the impact of bias at a more granular level.

In addition, analysts should frequently provide confidence intervals—that is, guidance on how much managers can trust the data if the n 's are too small to prove statistical significance. When managers get that information, they're more likely to make changes in their hiring and management practices, even if they believe—as most do—that they are already treating people fairly. Suppose, for example, that as Red Ventures began collecting data on self-assessments, analysts had a 75% confidence level that blacks and Latinos were under-rating themselves. The analysts could then have advised managers to go to their minority direct reports, examine the results from that performance period, and determine together whether the self-reviews truly reflected their contributions. It's a simple but collaborative way to address implicit

bias or stereotyping that you're reasonably sure is there while giving agency to each employee.

Second, companies also need to be more consistent and comprehensive in their qualitative analysis. Many already conduct interviews and focus groups to gain insights on the challenges of the underrepresented, some even do textual analysis of written performance reviews, exit interview notes, and hitting memos, looking for language that signals bias or negative stereotyping. But we have to go further. We need to find a viable way to create and process more-objective performance evaluations, given the internalized biases of both employees and managers, and to determine how those biases affect ratings.

This journey begins with educating all employees on the real-life impact of bias and negative stereotypes. At Facebook we offer a variety of training programs with an emphasis on spotting and countering bias, and we keep reinforcing key messages post-training, since we know these muscles take time to build. We issue reminders at critical points to shape decision making and behavior. For example, in our performance evaluation tool, we incorporate prompts for people to check word choice when writing reviews and self-assessments. We remind them, for instance, that terms like "cultural fit" can allow bias to creep in and that they should avoid describing women as "bossy" if they wouldn't describe men who demonstrated the same behaviors that way. We don't yet have data on how this is influencing the language used—it's a new intervention—but we will be examining patterns over time.

Perhaps above all, HR and analytics departments must value both qualitative and quantitative expertise and apply mixed-method approaches everywhere possible. At Facebook we're building cross-functional teams with both types of specialists, because no single research method can fully capture the complex layers of bias that everyone brings to the workplace. We view all research methods as trying to solve the same problem from different angles. Sometimes we approach challenges from a quantitative perspective first, to uncover the "what" before looking to the qualitative experts to dive into the "why" and "how." For instance, if the numbers showed that certain teams were losing or attracting minority employees at higher rates than others (the "what"), we might conduct interviews, run focus groups, or analyze text from company surveys to understand

Algorithms and statistics do not capture what it feels like to be the only black or Hispanic member of a team.

the "why," and pull out themes or lessons for other parts of the company. In other scenarios we might reverse the order of those steps. For example, if we repeatedly heard from members of one social group that they weren't seeing their peers getting recognized at the same rate as people in other groups, we could then investigate whether numerical trends confirmed those observations, or conduct statistical analyses to figure out which organizational circumstances were associated with employees' being more or less likely to get recognized.

Cross-functional teams also help us reap the benefits of cognitive diversity. Working together stretches everyone, challenging team members' own assumptions and biases. Getting to absolute "whys" and "hows" on any issue, from recruitment to engagement to performance, is always going to be tough. But we believe that with this approach, we stand the best chance of making improvements across the company. As we analyze the results of Facebook's Pulse survey, given twice a year to employees, and review Performance Summary Cycle inputs, we'll continue to look for signs of problems as well as progress.

EVIDENCE OF DISCRIMINATION or unfair outcomes may not be as certain or obvious in the workplace as it was for me the time I was evicted from my apartment. But we can increase our certainty, and it's essential that we do so. The underrepresented people at our companies are not crazy to perceive biases working against them, and they can get institutional support. ☺

HBR Reprint R706L