



Peter Cappelli
Professor,
the Wharton School

Spotlight

Your Approach to Hiring Is All Wrong

Outsourcing and algorithms won't get you the people you need.



BUSINESSES HAVE NEVER done as much hiring as they do today. They've never spent as much money doing it. And they've never done a worse job of it.

For most of the post-World War II era, large corporations went about hiring this way: Human resources experts prepared a detailed *job analysis* to determine what tasks the job required and what attributes a good candidate should have. Next they did a *job evaluation* to determine how the job fit into the organizational chart and how much it should pay, especially compared with other jobs. Ads were posted, and applicants applied. Then came the task of sorting through the applicants. That included skills tests, reference checks, maybe personality and IQ tests, and extensive interviews to learn more about them as people. William H. Whyte, in *The Organization Man*, described this process as going on for as long as a week before

Spotlight

the winning candidate was offered the job. The vast majority of non-entry-level openings were filled from within.

Today's approach couldn't be more different. Census data shows, for example, that the majority of people who took a new job last year weren't searching for one: Somebody came and got them. Companies seek to fill their recruiting funnel with as many candidates as possible, especially "passive candidates," who aren't looking to move. Often employers advertise jobs that don't exist, hoping to find people who might be useful later on or in a different context.

The recruiting and hiring function has been eviscerated. Many U.S. companies—about 40%, according to research by Korn Ferry—have outsourced much if not all of the hiring process to familiar organizations such as Randstad, Manpower, and Adecco, which in turn use subcontractors, typically in India and the Philippines. The subcontractors scour LinkedIn and social media to find potential candidates. They sometimes contact them directly to see whether they can be persuaded to apply for a position and negotiate the salary they're willing to accept. (The recruiters get incentive pay if they negotiate the amount down.) To hire programmers, for example, these

subcontractors can scan websites that programmers might visit, trace their "digital exhaust" from cookies and other user-tracking measures to identify who they are, and then examine their curricula vitae.

At companies that still do their own recruitment and hiring, managers trying to fill open positions are largely left to figure out what the jobs require and what the ads should say. When applications come—always electronically—applicant-tracking software sifts through them for key words that the hiring managers want to see. Then the process moves into the Wild West, where a new industry of vendors offer an astonishing array of smart-sounding tools that claim to predict who will be a good hire. They use voice recognition, body language, clues on social media, and especially machine learning algorithms—everything but tea leaves. Entire publications are devoted to what these vendors are doing.

The big problem with all these new practices is that we don't know whether they actually produce satisfactory hires. Only about a third of U.S. companies report that they monitor whether their hiring practices lead to good employees; few of them do so carefully, and only a minority even track cost per hire and time to hire. Imagine if the CEO asked

how an advertising campaign had gone, and the response was "We have a good idea how long it took to roll out and what it cost, but we haven't looked to see whether we're selling more."

Hiring talent remains the number one concern of CEOs in the most recent Conference Board Annual Survey; it's also the top concern of the entire executive suite. PwC's 2017 CEO survey reports that chief executives view the unavailability of talent and skills as the biggest threat to their business. Employers also spend an enormous amount on hiring—an average of \$4,129 per job in the United States, according to Society for Human Resource Management estimates, and many times that amount for managerial roles—and the United States fills a staggering 66 million jobs a year. Most of the \$20 billion that companies spend on human resources vendors goes to hiring.

Why do employers spend so much on something so important while knowing so little about whether it works?

WHERE THE PROBLEM STARTS

Survey after survey finds employers complaining about how difficult hiring is. There may be many explanations, such as their having become very picky about candidates, especially in the

Idea in Brief

THE PROBLEM

Employers continue to hire at a high rate and spend enormous sums to do it. But they don't know whether their approaches are effective at finding and selecting good candidates.

THE ROOT CAUSES

Businesses focus on external candidates and don't track the results of their approaches. They often use outside vendors and high-tech tools that are unproven and have inherent flaws.

THE SOLUTION

Return to filling most positions by promoting from within. Measure the results produced by vendors and new tools, and be on the lookout for discrimination and privacy violations.

Protecting Against Discrimination

slack labor market of the Great Recession. But clearly they are hiring much more than at any other time in modern history, for two reasons.

The first is that openings are now filled more often by hiring from the outside than by promoting from within. In the era of lifetime employment, from the end of World War II through the 1970s, corporations filled roughly 90% of their vacancies through promotions and lateral assignments. Today the figure is a third or less. When they hire from outside, organizations don't have to pay to train and develop their employees. Since the restructuring waves of the early 1980s, it has been relatively easy to find experienced talent outside. Only 28% of talent acquisition leaders today report that internal candidates are an important source of people to fill vacancies—presumably because of less internal development and fewer clear career ladders.

Less promotion internally means that hiring efforts are no longer concentrated on entry-level jobs and recent graduates. (If you doubt this, go to the “careers” link on any company website and look for a job opening that doesn't require prior experience.) Now companies must be good at hiring across most levels, because the candidates they want are already doing the job somewhere else. These people don't need training, so they may be ready to contribute right away, but they are much harder to find.

The second reason hiring is so difficult is that retention has become tough: Companies hire from their competitors and vice versa, so they have to keep replacing people who leave. Census and

Bureau of Labor Statistics data shows that 95% of hiring is done to fill existing positions. Most of those vacancies are caused by voluntary turnover. LinkedIn data indicates that the most common reason employees consider a position elsewhere is career advancement—which is surely related to employers' not promoting to fill vacancies.

The root cause of most hiring, therefore, is drastically poor retention. Here are some simple ways to fix that:

Track the percentage of openings filled from within. An adage of business is that we manage what we measure, but companies don't seem to be applying that maxim to tracking hires. Most are shocked to learn how few of their openings are filled from within—is it really the case that their people can't handle different and bigger roles?

Require that all openings be posted internally. Internal job boards were created during the dot-com boom to reduce turnover by making it easier for people to find new jobs within their existing employer. Managers weren't even allowed to know if a subordinate was looking to move within the company, for fear that they would try to block that person and he or she would leave. But during the Great Recession employees weren't quitting, and many companies slid back to the old model whereby managers could prevent their subordinates from moving internally. JR Keller, of Cornell University, has found that when managers could fill a vacancy with someone they already had in mind, they ended up with employees who performed more poorly than those hired when the job had been posted and anyone could apply. The

Finding out whether your practices result in good hires is not only basic to good management but the only real defense against claims of adverse impact and discrimination. Other than white males under age 40 with no disabilities or work-related health problems, workers have special protections under federal and state laws against hiring practices that may have an adverse impact on them. As a practical matter, that means if members of a particular group are less likely to be recruited or hired, the employer must show that the hiring process is not discriminatory.

The only defense against evidence of adverse impact is for the employer to show that its hiring practices are valid—that is, they predict who will be a good employee in meaningful and statistically significant ways—and that no alternative would predict as well with less adverse impact. That analysis must be conducted with data on the employer's own applicants and hires. The fact that the vendor that sold you the test you use has evidence that it was valid in other contexts is not sufficient.

commonsense explanation for this is that few enterprises really know what talent and capabilities they have.

Recognize the costs of outside hiring. In addition to the time and effort of hiring, my colleague Matthew Bidwell found, outside hires take three years to perform as well as internal hires in the same job, while internal hires take seven years to earn as much as outside hires are paid. Outside hiring also causes current employees to spend time and energy positioning themselves for jobs elsewhere. It disrupts the culture and burdens peers who must help new hires figure out how things work.



Obsessed with new technologies and driving down costs, employers largely ignore the ultimate goal: making the best possible hires.

None of this is to suggest that outside hiring is necessarily a bad idea. But unless your company is a Silicon Valley gazelle, adding new jobs at a furious pace, you should ask yourself some serious questions if most of your openings are being filled from outside.

A different approach for dealing with retention (which seems creepy to some) is to try to determine who is interested in leaving and then intervene. Vendors like Jobvite comb social media and public sites for clues, such as LinkedIn profile updates. Measuring “flight risk” is one of the most common goals of companies that do their own sophisticated HR analytics. This is reminiscent of the early days of job boards, when employers would try to find out who was posting résumés and either punish them or embrace them, depending on leadership’s mood.

Whether companies should be examining social media content in relation to hiring or any other employment action is a challenging ethical question. On one hand, the information is essentially public and may reveal relevant information. On the other hand, it is invasive, and candidates are rarely asked for permission to scrutinize their information. Hiring a private detective to shadow a candidate would also gather public information that might be relevant, yet most people would view it as an unacceptable invasion of privacy.

THE HIRING PROCESS

When we turn to hiring itself, we find that employers are missing the forest for the trees: Obsessed with new

technologies and driving down costs, they largely ignore the ultimate goal: making the best possible hires. Here’s how the process should be revamped:

Don’t post “phantom jobs.” It costs nothing to post job openings on a company website, which are then scooped up by Indeed and other online companies and pushed out to potential job seekers around the world. Thus it may be unsurprising that some of these jobs don’t really exist. Employers may simply be fishing for candidates. (“Let’s see if someone really great is out there, and if so, we’ll create a position for him or her.”) Often job ads stay up even after positions have been filled, to keep collecting candidates for future vacancies or just because it takes more effort to pull the ad down than to leave it up. Sometimes ads are posted by unscrupulous recruiters looking for résumés to pitch to clients elsewhere. Because these phantom jobs make the labor market look tighter than it really is, they are a problem for economic policy makers as well as for frustrated job seekers. Companies should take ads down when jobs are filled.

Design jobs with realistic requirements. Figuring out what the requirements of a job should be—and the corresponding attributes candidates must have—is a bigger challenge now, because so many companies have reduced the number of internal recruiters whose function, in part, is to push back on hiring managers’ wish lists. (“That job doesn’t require 10 years of experience,” or “No one with all those qualifications will be willing to accept the salary you’re proposing

to pay.”) My earlier research found that companies piled on job requirements, baked them into the applicant-tracking software that sorted résumés according to binary decisions (yes, it has the key word; no, it doesn’t), and then found that virtually no applicants met all the criteria. Trimming recruiters, who have expertise in hiring, and handing the process over to hiring managers is a prime example of being penny-wise and pound-foolish.

Reconsider your focus on passive candidates. The recruiting process begins with a search for experienced people who aren’t looking to move. This is based on the notion that something may be wrong with anyone who wants to leave his or her current job. (Of the more than 20,000 talent professionals who responded to a LinkedIn survey in 2015, 86% said their recruiting organizations focused “very much so” or “to some extent” on passive candidates; I suspect that if anything, that number has since grown.) Recruiters know that the vast majority of people are open to moving at the right price: Surveys of employees find that only about 15% are *not* open to moving. As the economist Harold Demsetz said when asked by a competing university if he was happy working where he was: “Make me unhappy.”

Fascinating evidence from the LinkedIn survey cited above shows that although self-identified “passive” job seekers are different from “active” job seekers, it’s not in the way we might think. The number one factor that would encourage the former to move is more money. For active candidates the top factor is better work and

Spotlight

career opportunities. More active than passive job seekers report that they are passionate about their work, engaged in improving their skills, and reasonably satisfied with their current jobs. They seem interested in moving because they are ambitious, not because they want higher pay.

Employers spend a vastly disproportionate amount of their budgets on recruiters who chase passive candidates, but on average they fill only 11% of their positions with individually targeted people, according to research by Gerry Crispin and Chris Hoyt, of CareerXroads. I know of no evidence that passive candidates become better employees, let alone that the process is cost-effective. If you focus on passive candidates, think carefully about what that actually gets you. Better yet, check your data to find out.

Understand the limits of referrals.

The most popular channel for finding new hires is through employee referrals; up to 48% come from them, according to LinkedIn research. It seems like a cheap way to go, but does it produce better hires? Many employers think so. It's hard to know whether that's true, however, given that they don't check. And research by Emilio Castilla and colleagues suggests otherwise: They find that when referrals work out better than other hires, it's because their referrers look after them and essentially onboard them. If a referrer leaves before the new hire begins, the latter's performance is no better than that of nonreferrals, which is why it makes sense to pay referral bonuses six months or so after the person is hired—if he or she is still there.

A downside to referrals, of course, is that they can lead to a homogeneous workforce, because the people we know tend to be like us. This matters greatly for organizations interested in diversity, since recruiting is the only avenue allowed under U.S. law to increase diversity in a workforce. The Supreme Court has ruled that demographic criteria cannot be used even to break ties among candidates.

Measure the results. Few employers know which channel produces the best candidates at the lowest cost because they don't track the outcomes. Tata is an exception: It has long done what I advocate. For college recruiting, for example, it calculates which schools send it employees who perform the best, stay the longest, and are paid the lowest starting wage. Other employers should follow suit and monitor recruiting channels and employees' performance to identify which sources produce the best results.

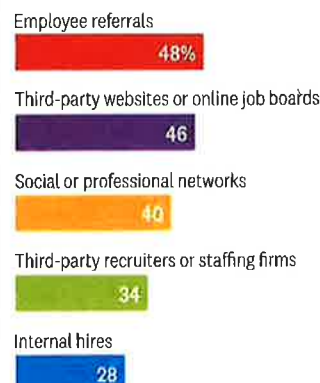
Persuade fewer people to apply.

The hiring industry pays a great deal of attention to "the funnel," whereby readers of a company's job postings become applicants, are interviewed, and ultimately are offered jobs. Contrary to the popular belief that the U.S. job market is extremely tight right now, most jobs still get lots of applicants. Recruiting and hiring consultants and vendors estimate that about 2% of applicants receive offers. Unfortunately, the main effort to improve hiring—virtually always aimed at making it faster and cheaper—has been to shovel more applicants into the funnel. Employers do that primarily through marketing, trying to get out the

The Grass Is Always Greener...

Organizations are much more interested in external talent than in their own employees to fill vacancies.

Top channels for quality hires



Source: LinkedIn Talent Solutions, 2012. Based on data from LinkedIn Talent Solutions, 2012. Based on data from LinkedIn Talent Solutions, 2012.

word that they are great places to work. Whether doing this is a misguided way of trying to attract better hires or just meant to make the organization feel more desirable isn't clear.

Much better to go in the other direction: Create a smaller but better-qualified applicant pool to improve the yield. Here's why: Every applicant costs you money—especially now, in a labor market where applicants have started to "ghost" employers, abandoning their applications midway through the process. Every application also exposes a company to legal risk, because the company has obligations to candidates (not to discriminate, for example) just as it does to employees. And collecting lots of applicants in a wide funnel means that a great many of them won't fit the job or the company, so employers have to rely on the next step of the hiring process—selection—to weed them out. As we will see, employers aren't good at that.

Once people are candidates, they may not be completely honest about their skills or interests—because they want to be hired—and employers’ ability to find out the truth is limited. More than a generation ago the psychologist John Wanous proposed giving applicants a realistic preview of what the job is like. That still makes sense as a way to head off those who would end up being unhappy in the job. It’s not surprising that Google has found a way to do this with gamification: Job seekers see what the work would be like by playing a game version of it. Marriott has done the same, even for low-level employees. Its My Marriott Hotel game targets young people in developing countries who may have had little experience in hotels to show them what it’s like and to steer them to the recruiting site if they score well on the game. The key for any company, though, is that the preview should make clear what is difficult and challenging about the work as well as why it’s fun so that candidates who don’t fit won’t apply.

It should be easy for candidates to learn about a company and a job, but making it really easy to apply, just to fill up that funnel, doesn’t make much sense. During the dot-com boom Texas Instruments cleverly introduced a preemployment test that allowed applicants to see their scores before they applied. If their scores weren’t high enough for the company to take their applications seriously, they tended not to proceed, and the company saved the cost of having to process their applications.

If the goal is to get better hires in a cost-effective manner, it’s more

important to scare away candidates who don’t fit than to jam more candidates into the recruiting funnel.

Test candidates’ standard skills.

How to determine which candidates to hire—what predicts who will be a good employee—has been rigorously studied at least since World War I. The personnel psychologists who investigated this have learned much about predicting good hires that contemporary organizations have since forgotten, such as that neither college grades nor unstructured sequential interviews (hopping from office to office) are a good predictor, whereas past performance is.

Since it can be difficult (if not impossible) to glean sufficient information about an outside applicant’s past performance, what other predictors are good? There is remarkably little consensus even among experts. That’s mainly because a typical job can have so many tasks and aspects, and different factors predict success at different tasks.

There is general agreement, however, that testing to see whether individuals have standard skills is about the best we can do. Can the candidate speak French? Can she do simple programming tasks? And so forth. But just doing the tests is not enough. The economists Mitchell Hoffman, Lisa B. Kahn, and Danielle Li found that even when companies conduct such tests, hiring managers often ignore them—and when they do, they get worse hires. The psychologist Nathan Kuncel and colleagues discovered that even when hiring managers use objective criteria and tests, applying their own weights and judgment to those criteria leads them to pick worse candidates

than if they had used a standard formula. Only 40% of employers, however, do any tests of skills or general abilities, including IQ. What are they doing instead? Seventy-four percent do drug tests, including for marijuana use; even employers in states where recreational use is now legal still seem to do so.

Be wary of vendors bearing high-tech gifts. Into the testing void has come a new group of entrepreneurs who either are data scientists or have them in tow. They bring a fresh approach to the hiring process—but often with little understanding of how hiring actually works. John Sumser, of HRExaminer, an online newsletter that focuses on HR technology, estimates that on average, companies get five to seven pitches *every day*—almost all of them about hiring—from vendors using data science to address HR issues. These vendors have all sorts of cool-sounding assessments, such as computer games that can be scored to predict who will be a good hire. We don’t know whether any of these actually lead to better hires, because few of them are validated against actual job performance. That aside, these assessments have spawned a counterwave of vendors who help candidates learn how to score well on them. Lloyds Bank, for example, developed a virtual-reality-based assessment of candidate potential, and JobTestPrep offers to teach potential candidates how to do well on it. Especially for IT and technical jobs, cheating on skills tests and even video interviews (where colleagues off camera give help) is such a concern that eTeki and other

specialized vendors help employers figure out who is cheating in real time.

Revamp your interviewing process. The amount of time employers spend on interviews has almost doubled since 2009, according to research from Glassdoor. How much of that increase represents delays in setting up those interviews is impossible to tell, but it provides at least a partial explanation for why it takes longer to fill jobs now. Interviews are arguably the most difficult technique to get right, because interviewers should stick to questions that predict good hires—mainly about past behavior or performance that’s relevant to the tasks of the job—and ask them consistently across candidates.

Just winging it and asking whatever comes to mind is next to useless.

More important, interviews are where biases most easily show up, because interviewers do usually decide on the fly what to ask of whom and how to interpret the answer. Everyone knows some executive who is absolutely certain he knows the one question that will really predict good candidates (“If you were stranded on a desert island...”). The sociologist Lauren Rivera’s examination of interviews for elite positions, such as those in professional services firms, indicates that hobbies, particularly those associated with the rich, feature prominently as a selection criterion.

Interviews are most important for assessing “fit with our culture,” which is the number one hiring criterion employers report using, according to research from the Rockefeller Foundation. It’s also one of the squishiest attributes to measure, because few organizations have an accurate and consistent view of their own culture—and even if they do, understanding what attributes represent a good fit is not straightforward. For example, does the fact that an applicant belonged to a fraternity reflect experience working with others or elitism or bad attitudes toward women? Should it be completely irrelevant? Letting someone with no experience or training make such calls

Data Science Can't Fix Hiring (Yet)



Recruiting managers desperately need new tools; because the existing ones—unstructured interviews, personality tests, personal referrals—aren’t very effective. The newest development in hiring, which is both promising and worrying, is the rise of data science-driven algorithms to find and assess job candidates. By my count, more than 100 vendors are creating and selling these tools to companies.

Unfortunately, data science—which is still in its infancy when it comes to recruiting and hiring—is not yet the panacea employers hope for.

Vendors of these new tools promise they will help reduce the role that social bias plays in hiring. And the algorithms can indeed help identify good job candidates who would previously have been screened out for lack of a certain education or social pedigree. But these tools may also

identify and promote the use of predictive variables that are (or should be) troubling.

Because most data scientists seem to know so little about the context of employment, their tools are often worse than nothing. For instance, an astonishing percentage build their models by simply looking at attributes of the “best performers” in workplaces and then identifying which job candidates have the same attributes. They use anything that’s easy to measure: facial expressions, word choice, comments on social media, and so forth. But a failure to check for any real difference between high-performing and low-performing employees on these attributes limits their usefulness. Furthermore, scooping up data from social media or the websites

people have visited also raises important questions about privacy. True, the information can be accessed legally; but the individuals who created the postings didn’t intend or authorize them to be used for such purposes. Furthermore, is it fair that something you posted as an undergrad can end up driving your hiring algorithm a generation later?

Another problem with machine learning approaches is that few employers collect the large volumes of data—number of hires, performance appraisals, and so on—that the algorithms require to make accurate predictions. Although vendors can theoretically overcome that hurdle by aggregating data from many employers, they don’t really know whether individual



Interviews are where biases most easily show up, because interviewers usually decide on the fly what to ask of whom and how to interpret the answer.

is a recipe for bad hires and, of course, discriminatory behavior. Think hard about whether your interviewing protocols make any sense and resist the urge to bring even more managers into the interview process.

Recognize the strengths and weaknesses of machine learning models.

Culture fit is another area into which new vendors are swarming. Typically they collect data from current employees, create a machine learning model to predict the attributes of the best ones, and then use that model to hire candidates with the same attributes.

As with many other things in this new industry, that sounds good until you think about it; then it becomes

replete with problems. Given the best performers of the past, the algorithm will almost certainly include *white* and *male* as key variables. If it's restricted from using that category, it will come up with attributes associated with being a white male, such as playing rugby.

Machine learning models do have the potential to find important but previously unconsidered relationships. Psychologists, who have dominated research on hiring, have been keen to study attributes relevant to their interests, such as personality, rather than asking the broader question "What identifies a potential good hire?" Their results gloss over the fact that they often have only a trivial ability to

predict who will be a good performer, particularly when many factors are involved. Machine learning, in contrast, can come up with highly predictive factors. Research by Evolv, a workforce analytics pioneer (now part of Cornerstone OnDemand), found that expected commuting distance for the candidate predicted turnover very well. But that's not a question the psychological models thought to ask. (And even that question has problems.)

The advice on selection is straightforward: Test for skills. Ask assessments vendors to show evidence that they can actually predict who the good employees will be. Do fewer, more-consistent interviews.

company contexts are so distinct that predictions based on data from the many are inaccurate for the one.

Yet another issue is that all analytic approaches to picking candidates are backward looking, in the sense that they are based on outcomes that have already happened. (Algorithms are especially reliant on past experiences in part because building them requires lots and lots of observations—many years' worth of job performance data even for a large employer.) As Amazon learned, the past may be very different from the future you seek. It discovered that the hiring algorithm it had been working on since 2014 gave lower scores to women—even to attributes associated with women, such as participating in women's studies

programs—because historically the best performers in the company had disproportionately been men. So the algorithm looked for people just like them. Unable to fix that problem, the company stopped using the algorithm in 2017. Nonetheless, many other companies are pressing ahead.

The underlying challenge for data scientists is that hiring is simply not like trying to predict, say, when a ball bearing will fail—a question for which any predictive measure might do. Hiring is so consequential that it is governed not just by legal frameworks but by fundamental notions of fairness. The fact that some criterion is associated with good job performance is necessary but not sufficient for using it in hiring.

Take a variable that data scientists have found to have predictive value: commuting distance to the job. According to the data, people with longer commutes suffer higher rates of attrition. However, commuting distance is governed by where you live—which is governed by housing prices, relates to income, and also relates to race. Picking whom to hire on the basis of where they live most likely has an adverse impact on protected groups such as racial minorities.

Unless no other criterion predicts at least as well as the one being used—and that is extremely difficult to determine in machine learning algorithms—companies violate the law if they use hiring criteria that have adverse impacts. Even then, to stay on the right side of the law,

they must show why the criterion creates good performance. That might be possible in the case of commuting time, but—at least for the moment—it is not for facial expressions, social media postings, or other measures whose significance companies cannot demonstrate.

In the end, the drawback to using algorithms is that we're trying to use them on the cheap: building them by looking only at best performers rather than all performers, using only measures that are easy to gather, and relying on vendors' claims that the algorithms work elsewhere rather than observing the results with our own employees. Not only is there no free lunch here, but you might be better off skipping the cheap meal altogether. ☹

HBR Reprint R1903B



Dane E. Holmes
Head of
HR, Goldman
Sachs

THE WAY FORWARD

It's impossible to get better at hiring if you can't tell whether the candidates you select become good employees. If you don't know where you're going, any road will take you there. You must have a way to measure which employees are the best ones.

Why is that not getting through to companies? Surveyed employers say the main reason they don't examine whether their practices lead to better hires is that measuring employee performance is difficult. Surely this is a prime example of making the perfect the enemy of the good. Some aspects of performance are not difficult to measure: Do employees quit? Are they absent? Virtually all employers conduct performance appraisals. If you don't trust them, try something simpler. Ask supervisors, "Do you regret hiring this individual? Would you hire him again?"

Organizations that don't check to see how well their practices predict the quality of their hires are lacking in one of the most consequential aspects of modern business. ©

HBR Reprint R19038

PETER CAPPELLI is the George W. Taylor Professor of Management at the Wharton School and the director of its Center for Human Resources. His most recent book is *Will College Pay Off? A Guide to the Most Important Financial Decision You'll Ever Make* (PublicAffairs, 2015).

Expanding the Pool

How Goldman Sachs Changed the Way It Recruits



GOLDMAN SACHS IS A people-centric business—every day our employees engage with our clients to find solutions to their challenges. As a consequence, hiring extraordinary talent is vital to our success and can never be taken for granted. In the wake of the 2008 financial crisis we faced a challenge that was, frankly, relatively new to our now 150-year-old firm. For decades investment banking had been one of the most sought-after, exciting, and fast-growing industries in the world. That made sense—we were growing by double digits and had high returns, which meant that opportunity and reward were in great supply. However, the crash took some of the sheen off our industry; both growth and returns moderated. And simultaneously, the battle for talent intensified—within and outside our industry. Many of the

candidates we were pursuing were heading off to Silicon Valley, private equity, or start-ups. Furthermore, we were no longer principally looking for a specialized cadre of accounting, finance, and economics majors: New skills, especially coding, were in huge demand at Goldman Sachs—and pretty much everywhere else. The wind had shifted from our backs to our faces, and we needed to respond.

Not long ago the firm relied on a narrower set of factors for identifying "the best" students, such as school, GPA, major, leadership roles, and relevant experience—the classic résumé topics. No longer. We decided to replace our hiring playbook with emerging best practices for assessment and recruitment, so we put together a task force of senior business leaders, PhDs in industrial and organizational

psychology, data scientists, and experts in recruiting. Some people asked, “Why overhaul a recruiting process that has proved so successful?” and “Don’t you already have many more qualified applicants than available jobs?” These were reasonable questions. But often staying successful is about learning and changing rather than sticking to the tried-and-true.

Each year we hire up to 3,000 summer interns and nearly as many new analysts directly from campuses. In our eyes, these are the firm’s future leaders, so it made sense to focus our initial reforms there. They involved two major additions to our campus recruiting strategy—video interviews and structured interviewing.

Asynchronous video interviews. Traditionally we had flown recruiters and business professionals to universities for first-round interviews. The schools would give us a set date and number of time slots to meet with students. That is most definitely not a scalable model. It restricted us to a smaller number of campuses and only as many students as we could squeeze into a limited schedule. It also meant that we tended to focus on top-ranked schools. How many qualified candidates were at a school became more important than who were the most talented students regardless of their school. However, we knew that candidates didn’t have to attend Harvard, Princeton, or Oxford to excel at Goldman Sachs—our leadership ranks were already rich with people from other schools. What’s more, as we’ve built offices in new cities and geographic locations, we’ve needed to recruit at

more schools located in those areas. Video interviews allow us to do that.

At a time when companies were just beginning to experiment with digital interviewing, we decided to use “asynchronous” video interviews—in which candidates record their answers to interview questions—for all first-round interactions with candidates. Our recruiters record standardized questions and send them to students, who have three days to return videos of their answers. This can be done on a computer or a mobile device. Our recruiters and business professionals review the videos to narrow the pool and then invite the selected applicants to a Goldman Sachs office for final-round, in-person interviews. (To create the video platform, we partnered with a company and built our own digital solution around its product.)

This approach has had a meaningful impact in two ways. First, with limited effort, we can now spend more time getting to know the people who apply for jobs at Goldman Sachs. In 2015, the year before we rolled out this platform, we interviewed fewer than 20% of all our campus applicants; in 2018 almost 40% of the students who applied to the firm participated in a first-round interview. Second, we now encounter talent from places we previously didn’t get to. In 2015 we interviewed students from 798 schools around the world, compared with 1,268 for our most recent incoming class. In the United States, where the majority of our student hires historically came from “target schools,” the opposite is now true. The top of our recruiting funnel is wider, and the output is more diverse.

Being a people-driven business, we have worked hard to ensure that the video interviews don’t feel cold and impersonal. They are only one component of a broader process that makes up the Goldman Sachs recruitment experience. We still regularly send Goldman professionals to campuses to engage directly with students at informational sessions, “coffee chats,” and other recruiting events. But now our goal is much more to share information than to assess candidates, because we want people to understand the firm and what it offers before they tell us why they want an internship or a job.

We also want them to be as well prepared as possible for our interview process. Our goal is a level playing field. To help achieve it, we’ve created tip sheets and instructions on preparing for a video interview. Because the platform doesn’t allow videos to be edited once they’ve been recorded, we offer a practice question before the interview begins and a countdown before the questions are asked. We also give students a formal channel for escalating issues should technical problems arise, though that rarely occurs.

We’re confident that this approach has created a better experience for recruits. It uses a medium they’ve grown up with (video), and most important, they can do their interviews when they feel fresh and at a time that works with their schedule. (Our data shows that they prefer Thursday or Sunday night—whereas our previous practice was to interview during working hours.) We suspected that if the process was a turnoff for applicants, we would see a dip in the percentage who



We are focused less on past achievements and more on understanding whether a candidate has qualities that will positively affect our firm and our culture. Our structured interview questions are designed to assess 10 core competencies.

accepted our interviews and our offers. That hasn't happened.

Structured questioning and assessments. How can you create an assessment process that not only helps select top talent but focuses on specific characteristics associated with success? Define it, structure it, and don't deviate from it. Research shows that structured interviews are effective at assessing candidates and helping predict job performance. So we ask candidates about specific experiences they've had that are similar to situations they may face at Goldman Sachs ("Tell me about a time when you were working on a project with someone who was not completing his or her tasks") and pose hypothetical scenarios they might encounter in the future ("In an elevator, you overhear confidential information about a coworker who is also a friend. The friend approaches you and asks if you've heard anything negative about him recently. What do you do?").

Essentially, we are focused less on past achievements and more on understanding whether a candidate has qualities that will positively affect our firm and our culture. Our structured interview questions are designed to assess candidates on 10 core competencies, including analytical thinking and integrity, which we know correlate with long-term success at the firm. They are evaluated on six competencies in the first round; if they progress, they're assessed on the remaining four during in-person interviews.

We have a rotating library of questions for each competency, along with a rubric for interviewers that explains how to rate responses on a five-point

scale from "outstanding" to "poor." We also train our interviewers to conduct structured interviews, provide them with prep materials immediately before they interview a candidate, and run detailed calibration meetings using all the candidate data we've gathered throughout the recruiting process to ensure that certain interviewers aren't introducing grade inflation (or deflation). We're experimenting with prehire assessment tests to be paired with these interviews; we already offer a technical coding and math exam for applicants to our engineering organization.

We decided not to pilot these changes and instead rolled them out en masse, because we realized that buy-in would come from being able to show results quickly—and because we know that no process is perfect. Indeed, what I love most about our new approach is that we've turned our recruiting department into a laboratory for continuous learning and refinement. With more than 50,000 candidate video recordings, we're now sitting on a treasure trove of data that will help us conduct insightful analyses and answer questions necessary to run our business: Are we measuring the right competencies? Should some be weighted more heavily than others? What about the candidates' backgrounds? Which interviewers are most effective? Does a top-ranked student at a state school create more value for us than an average student from the Ivy League? We already have indications that students recruited from the new schools in our pool perform just as well as students from our traditional

ones—and in some cases are more likely to stay longer at the firm.

What's next for our recruiting efforts? We receive almost 500,000 applications each year. From this pool we hire approximately 3%. We believe that many of the other 97% could be very successful at Goldman Sachs. As a result, picking the right 3% is less about just the individual and increasingly about matching the right person to the right role. That match may be made straight out of college or years later. We're experimenting with résumé-reading algorithms that will help candidates identify the business departments best suited to their skills and interests. We're looking at how virtual reality might help us better educate students about working in our offices and in our industry. And we're evaluating various tools and tests to bring even more data into the hiring decision process. Can I imagine a future in which companies rely exclusively on machines and algorithms to rate résumés and interviews? Maybe, for some. But I don't see us ever eliminating the human element at Goldman Sachs; it's too deeply embedded in our culture, in the work we do, and in what we believe drives success.

I'm excited to see where this journey takes us. Our 2019 campus class is shaping up to be the most diverse ever—and it's composed entirely of people who were selected through rigorous, objective assessments. There's no way we aren't better off as a result. ☺

HBR Reprint R1903B



DANE E. HOLMES is the global head of human capital management at Goldman Sachs.